

The Hardness of LPN over Any Integer Ring and Field for PCG Applications

Hanlin Liu
Shanghai Jiao Tong University
hans1024@sjtu.edu.cn

Xiao Wang
Northwestern University
wangxiao@cs.northwestern.edu

Kang Yang
State Key Laboratory of Cryptology
yangk@sklc.org

Yu Yu
Shanghai Jiao Tong University
Shanghai Qizhi Institute
yuyu@yuyu.hk

Abstract

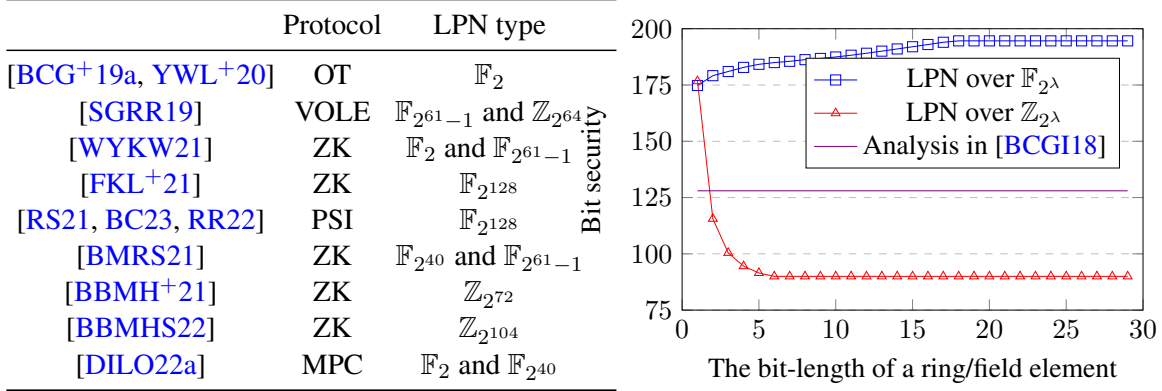
Learning parity with noise (LPN) has been widely studied and used in cryptography. It was recently brought to new prosperity since Boyle et al. (CCS'18), putting LPN to a central role in designing secure multi-party computation, zero-knowledge proofs, private set intersection, and many other protocols. In this paper, we thoroughly studied the concrete security of LPN problems in these settings. We found that many conclusions from classical LPN cryptanalysis do not apply to this new setting due to the low noise rates, extremely high dimensions, various types (in addition to \mathbb{F}_2) and noise distributions.

- For LPN over a field \mathbb{F}_q , we give a parameterized reduction from an exact noise distribution to a regular one that not only generalizes the recent result by Feneuil, Joux and Rivain (Crypto'22), but also significantly reduces the security loss by paying only an additive price in dimension and number of samples.
- We analyze the security of LPN over a ring \mathbb{Z}_{2^λ} . Although existing protocols based on LPN over integer rings use parameters as if they are over fields, we found an attack that effectively reduces the weight of a noise by half compared to LPN over fields. Consequently, prior works that use LPN over \mathbb{Z}_{2^λ} overestimate up to 40 bits of security.
- We provide a complete picture of the hardness of LPN over integer rings by showing: 1) the equivalence between its search and decisional versions; 2) an efficient reduction from LPN over \mathbb{F}_2 to LPN over \mathbb{Z}_{2^λ} ; and 3) generalization of our results to any integer ring.
- For LPN over finite fields, we found that prior analysis ignored some important differences between classical LPN cryptanalysis and the new setting, leading to overly conservative parameters. We show that even after bringing all classical LPN cryptanalysis, including the latest SD 2.0 analysis (Asiacrypt'22), to the setting over finite fields, much less weight of noises is needed for the same level of security.

To improve the use of LPN assumptions for a wide range of cryptographic protocols, we provide an open-sourced script that estimates the concrete security of LPN over integer rings and finite fields.

1 Introduction

The learning parity with noise (LPN) assumption states that it is hard to distinguish LPN samples ($\mathbf{A} \cdot \mathbf{s} + e$) from random samples, where \mathbf{A} is a public matrix, \mathbf{s} is a random secret and e is a noise vector sampled from a sparse distribution. The LPN assumption has been applied to build various primitives, e.g.,



(a) Prior works in the PCG framework and their required LPN variants over different fields and rings. (b) **The bit-security from our analysis for LPN over \mathbb{F}_{2^λ} and \mathbb{Z}_{2^λ} .** Parameters $N = 2^{10}$, $k = 652$, $t = 106$ are used.

Figure 1: **LPN assumptions in prior work and our analysis on one set of parameters.** For a set of parameters (N, k, t) , N is the number of samples, k is the dimension and t is the Hamming weight of a noise vector.

symmetric encryption and authentication (e.g., [HB01] and follow-up works), public key encryption [Ale03], commitment scheme [JKPT12], garbled circuits [App16], oblivious transfer [DDN14] and collision-resistant hash functions [BLVW19, YZW⁺19]. All these primitives adopt LPN over binary field \mathbb{F}_2 with moderate dimensions.

Recent works by Boyle et al. [BCGI18, BCG⁺19b] introduced the pseudorandom correlation generator (PCG) paradigm that can produce a large batch of correlated randomness, e.g., oblivious transfer (OT) and vector oblivious linear evaluation (VOLE), at a small cost. At the core of the PCG idea is to build a pseudorandom generator (PRG) with a simple internal structure from LPN-style assumptions and then privately evaluate such a PRG using function secret sharing [BGI15]. The sparsity of a noise e translates to communication efficiency, while the efficiency of LPN encoding translates to computational efficiency. Later, the PCG paradigm was used to build a series of concretely efficient protocols with sublinear communication [SGRR19, BCG⁺19b, YWL⁺20, WYKW21, CRR21, BBMHS22, BCG⁺22] for generating correlated OT (COT) or VOLE correlations. These PCG-like protocols have gained a lot of interests in designing various concretely efficient protocols, including secure multi-party computation (MPC) (e.g., [DPSZ12, NNOB12, KOS16, WRK17a, WRK17b, HSS20, CDE⁺18, YWZ20, DILO22a]), zero-knowledge (ZK) proofs (e.g., [WYKW21, BMRS21, DIO21, YSWW21, BBMH⁺21, DILO22b, BBMHS22, WYY⁺22]), privacy-preserving machine learning [SGRR19, WYX⁺21], private set intersection (PSI) [RS21, RR22, BC23], etc.

Although being widely used in a lot of cryptographic protocols and real-world applications, many of them use LPN variants that have not attracted much attention in cryptanalysis, especially compared to the classical LPN assumption over \mathbb{F}_2 [Ale03, FS09, HS13, TS16]. Furthermore, prior analysis on the classical LPN assumption does not directly cover the LPN variants used in the PCG setting because of their unique features:

- Protocols often require an LPN assumption over a ring other than \mathbb{F}_2 , including a finite field or even an integer ring¹ like \mathbb{Z}_{2^λ} .
- Most existing analyses focus on a Bernoulli noise distribution, but all PCG-like protocols adopt an exact noise distribution. Meanwhile, the difference in concrete security between the Bernoulli and exact noise

¹By integer ring we refer to \mathbb{Z}_N for any composite number N , which is used to distinguish from polynomial rings.

is unclear.

- Most applications require an LPN assumption with very high dimension (e.g., millions) and low noise rate (e.g., $1/10^5$), which is out of the typically reported range of parameters considered for coding theoretic primitives.

At this point, all implementations of PCG use the LPN parameters from the original work by Boyle et al. [BCGI18], who analyzed the concrete security of LPN over $\mathbb{F}_{2^{128}}$. However, as we summarize in Table 1a, follow-up works used the same analysis to choose parameters for many different variants of LPN over \mathbb{F}_2 , \mathbb{F}_p , and \mathbb{Z}_{2^λ} , many of which were not covered by the original analysis. It was not clear how large a gap in security when using LPN parameters over a field for LPN over another field or ring.

In addition, multiple PCG-like protocols [BCGI18, BCG+19b, BCG+19a, YWL+20, WYKW21, CRR21, BCG+22, BCG+22] together with some MPC protocols such as [HOSS18] adopt a variant of exact noise distributions called regular noise distributions to achieve significantly better efficiency, where a regular noise distribution further divides a noise vector into t consecutive sub-vectors of size $\lfloor N/t \rfloor$ and each sub-vector has a single noisy coordinate, where N is the number of samples. The recent reduction by Feneuil et al. [FJR22] from an exact noise distribution to a regular noise distribution incurs a large security loss, which makes the security gap between two noise distributions be still large.

1.1 Our Contribution

We conduct a systematic study on the hardness of the LPN problems used by the recent PCG-like protocols that can be in turn applied in MPC, ZK, PSI etc. For all relevant applications, each requiring a field or ring of different sizes, we provide concrete parameter recommendations by giving an open-sourced Python script² that estimates the bit-security of LPN problems over arbitrary integer rings and finite fields. The previous cryptanalysis provides a good coarse-level estimation but the exact bit-security can differ by 20 to 80 bits. In particular, the field size and the type of elements (finite field vs integer ring) hugely impact the security of the underlying LPN assumption. For example, as shown in Figure 1b³, for a set of parameters that [BCGI18] estimates to have 128-bit security, our analysis shows that the bit-security of LPN over finite fields is about 175 to 194 bits and its security over integer rings is mostly below 90 bits when the ring size is larger than 2^6 . These findings are a combination of correction/refinement from prior works as well as new attacks/reductions on LPN over rings, where we provide the upper bound and lower bound of their exact security.

We also give a tighter reduction from an exact noise distribution to a regular noise distribution, and thus significantly shrink the security gap between two noise distributions. We note that our reduction is not contradictory, but rather complementary, to the very recent work by Briaud and Øygarden [BØ23]. In particular, they proposed a new algebraic attack that can take advantage of regular noise distributions. They show that their algebraic attack on LPN with regular noise is more concretely efficient than other existing attacks for some primal-LPN parameter settings with small code rate, whereas our reduction establishes the asymptotic connection that LPN with regular noise could be as hard as that with exact-weight/Bernoulli noise up to some certain security loss. See Section 5.4 for detailed discussion.

The hardness of LPN under regular noise distributions. Recently, Feneuil et al. [FJR22] observed that an exact noise vector (of Hamming weight t) is also regular with some probability (estimated to e^{-t} in Section 3), and thus (T, ϵ) -hard LPN under an exact noise distribution implies $(T, e^t \cdot \epsilon)$ -hard LPN under a regular noise distribution. However, the security loss is sometimes unaffordable as LPN may not have

²This script can be found at <https://gist.github.com/hansliu1024/21c87609e75f6cc52decdc69981e1d5b>.

³In this figure and the following tables, the number of samples N and dimension k are adopted from [BCGI18], and the weight of a noise vector t in [BCGI18] is increased to make the parameters achieve 128-bit security according to the analysis [BCGI18].

LPN			This work					[BCGI18]
N	k	t	$\mathbb{F}_{2^{128}}$	\mathbb{F}_{2^8}	$\mathbb{Z}_{2^{128}}$	\mathbb{Z}_4	\mathbb{F}_2	Any field size
2^{10}	652	106	194	186	89	116	176	128
2^{12}	1589	172	155	146	76	95	131	128
2^{14}	3482	338	150	144	78	95	132	128
2^{16}	7391	667	151	148	82	99	135	128
2^{18}	15336	1312	153	153	87	104	139	128
2^{20}	32771	2467	155	157	92	108	143	128
2^{22}	67440	4788	156	160	97	113	147	128

Table 1: The comparison of our analysis and [BCGI18] about the bit-security of an LPN problem with dimension k , number of samples N and Hamming weight of noises t over different rings. Each noise vector is sampled uniformly at random such that it has a fixed Hamming weight t .

dual-LPN			This work					[BCGI18]
n	N	t	$\mathbb{F}_{2^{128}}$	\mathbb{F}_{2^8}	$\mathbb{Z}_{2^{128}}$	\mathbb{Z}_4	\mathbb{F}_2	Any field size
2^{10}	2^{12}	88	208	198	99	122	177	128
2^{12}	2^{14}	83	200	190	101	126	175	128
2^{14}	2^{16}	78	195	184	104	128	175	128
2^{16}	2^{18}	73	190	180	104	126	173	128
2^{18}	2^{20}	68	186	176	106	127	169	128
2^{20}	2^{22}	63	182	172	107	127	165	128
2^{22}	2^{24}	58	177	168	109	126	162	128

Table 2: The comparison of our analysis and [BCGI18] about the bit-security of a dual-LPN (a.k.a., syndrome decoding) problem with $n = N/4$ (corresponding to the number of COT/VOLE correlations), number of samples N , Hamming weight of noises t . Each noise vector is sampled uniformly at random such that it has a fixed weight t .

security beyond e^t in many practical settings. To reduce the security loss, we introduce a tunable parameter $\alpha \geq 2$ and divide a noise vector into αt blocks (each denoted by e_i). Furthermore, instead of hoping that every e_i has exact weight 1, we relax the condition to that the weight of e_i is *at most* 1. For each block, we add an extra sample with noise \tilde{e}_i such that vector (e_i, \tilde{e}_i) has the exact weight 1, which allows us to obtain a regular noise vector. As a result, we prove that if the LPN over a field \mathbb{F}_q with dimension k , sample number N , and exact weight- t noise is (T, ϵ) -hard, then the LPN over \mathbb{F}_q with dimension $(k + \alpha t)$, sample number $(N + \alpha t)$, and regular weight- (αt) noise is $(T - \text{poly}(k, N), 2^{\frac{t}{\alpha}} \cdot \epsilon)$ -hard, where the security loss is reduced by at least 2^α , while the dimension and number of samples are increased by only αt .

The hardness of LPN over integer rings. Although having been used in protocol design [SGRR19, BBMH⁺21, BBMHS22], LPN problems over integer rings (e.g., \mathbb{Z}_{2^k}) have only been studied in a handful of works. Akavia [Aka08] studied the hardness of learning characters with noise problem, which generalizes standard LPN to any integer ring and noises of bounded Hamming weight, in the random samples access model. It does not consider any reduction given in this paper.

All existing works select the parameters assuming that LPN over an integer ring is as secure as LPN over a finite field and then use the parameter recommendation from [BCGI18]. However, as our analysis shows in Figure 1b and in Tables 1 and 2, LPN over an integer ring is significantly more vulnerable to attacks than LPN over a finite field of similar size. *What's more, although LPN over a finite field becomes harder to attack as the field size increases, LPN over an integer ring becomes easier to attack as the ring*

size increases! We briefly discuss our results below and summarize their relationships in Figure 2.

1. Focusing on the most commonly used ring \mathbb{Z}_{2^λ} , we show a concrete attack that can solve a t -noise LPN over \mathbb{Z}_{2^λ} by solving a $\left(\frac{2^{(\lambda-1)}}{2^\lambda-1} \cdot t\right)$ -noise (which approximates to $t/2$) LPN over \mathbb{F}_2 . This means that LPN over an integer ring is concretely weaker than LPN over a finite field and we need to double the weight of noise vectors to cover this attack. The impact to existing cryptographic protocols is significant. It will lead to roughly $2 \times$ more communication and computation.
2. On the positive side, we provide an evidence that the LPN problem over an integer ring is generally hard. In particular, we show a reduction between t -noise LPN over \mathbb{F}_2 and $(\lambda \cdot t)$ -noise LPN over a ring \mathbb{Z}_{2^λ} , which means that LPN over an integer ring is asymptotically as hard as classical LPN. This “efficient” reduction requires a different noise distribution: instead of sampling t locations and putting a uniform non-zero entry from \mathbb{Z}_{2^λ} in each location, we need to sample λ weight- t noise vectors $e_0, \dots, e_{\lambda-1}$ over \mathbb{F}_2 , and define the final noise vector as $e = \sum_{i \in [\lambda]} 2^i \cdot e_i$ with weight $\leq \lambda \cdot t$. This noise distribution may be interesting, as it can be used in the design of PCG-like protocols by adopting the upper bound $\lambda \cdot t$ to run these protocols. This change of distributions is crucial: without such change, the best reduction that we can find goes from t -noise LPN over \mathbb{F}_2 to $(2^\lambda \cdot t)$ -noise LPN over \mathbb{Z}_{2^λ} , which is exponentially worse than the above. Another interesting fact is that the above reductions only require the code matrix \mathbf{A} to be Boolean, which eliminates the need for integer multiplication during LPN encoding. Prior work [CRR21] observed that using a Boolean code matrix is not vulnerable to existing linear-test attacks for LPN over finite fields; here we show that for LPN over integer rings, using a Boolean matrix is provably secure assuming that classical LPN over \mathbb{F}_2 is hard.
3. While the above reductions focus on the decisional version of LPN, we also give a reduction from computational LPN over \mathbb{Z}_{2^λ} to that over \mathbb{F}_2 . Thus, we show the equivalence between computational and decisional versions of LPN over \mathbb{Z}_{2^λ} as shown in Figure 2. We also generalize all the results to any integer ring. In particular, we show a concrete attack that can solve a t -noise LPN over ring $\mathbb{Z}_{p^{\lambda_1} q^{\lambda_2}}$ by solving either a $\left(\frac{p-1}{p} \cdot t\right)$ -noise LPN over \mathbb{F}_p or a $\left(\frac{q-1}{q} \cdot t\right)$ -noise LPN over \mathbb{F}_q , where p, q are two primes. This attack works for both computational and decisional versions of LPN. We also give a reduction from t -noise LPN over \mathbb{F}_p and t -noise LPN over \mathbb{F}_q to $((\lambda_1 + \lambda_2) \cdot t)$ -noise LPN over $\mathbb{Z}_{p^{\lambda_1} q^{\lambda_2}}$. Given these reductions over $\mathbb{Z}_{p^{\lambda_1} q^{\lambda_2}}$, one can easily generalize them to any integer ring.

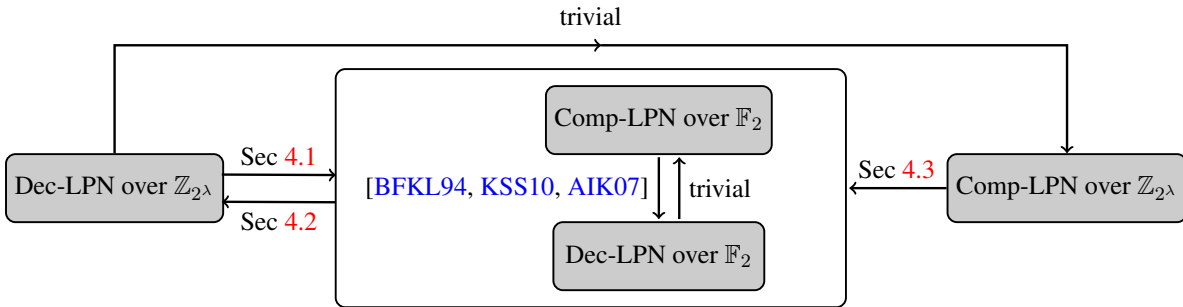


Figure 2: The reduction relations between computational and decisional versions of LPN over \mathbb{F}_2 and \mathbb{Z}_{2^λ} in the presence of Bernoulli and exact noise distribution.

The hardness of LPN over finite fields. When $|\mathbb{F}| > 2$, [BCGI18] is the main work that provides comprehensive cryptanalysis in this setting. They analyzed the concrete security of LPN instances over a large field by considering the Pooled Gauss (called Gaussian elimination in [BCGI18]), statistical decoding (SD,

a.k.a., low-weight parity check) and information set decoding (ISD) attacks, where an upper bound of the Prange’s ISD attack [Pra62] was adopted. Our new analysis shows that the parameters of LPN over finite fields given in [BCGI18] are overly conservative (see Table 1 and Table 2). Their analysis shows that SD is almost always the best attack and the attack cost does not change when changing the field size. We find that [BCGI18] ignored some important differences between classical LPN cryptanalysis and the new setting, and have the following results:

1. Concurrently, Carrier et al. [CDAMHT22] modified the framework of the SD attack and proposed a new variant, called the SD 2.0 algorithm. Under the Gilbert-Varshamov (GV) bound ⁴, they show that the attack even outperforms all ISD algorithms for the case that the code rate $k/N < 0.3$. However, we observe that the SD 2.0 algorithm [CDAMHT22] does not behave better when solving the low-noise LPN problems used in the PCG-like protocols. This is because the collision technique used by SD 2.0 takes exponential time $|\mathbb{F}|^{\theta(k)}$ that is much larger than the subexponential time $2^{O(k\mu)}$ to solve the low-noise LPN problem with ISD, where $\mu = 1/k^c$ is the noise rate (i.e., t/N) for constant $0 < c < 1$. Furthermore, we adapt the SD 2.0 algorithm to analyze low-noise LPN by replacing it with other collision techniques, e.g., [BCGI18].

For all the LPN parameters given in [BCGI18], we show that SD attack, even taking into account the adapted SD 2.0 algorithm, has the largest cost (see Table 3 and Table 4 of Section 5). Furthermore, we prove that both the optimal SD attack and SD 2.0 attack (adapted to the low-noise setting) require more cost than the original Prange’s ISD algorithm [Pra62] for a large set of commonly used parameters. Note that our results do not invalidate SD 2.0 due to the different parameter settings considered between [CDAMHT22] and us. This also shows the disparity of cryptanalysis between classical LPN problems with high noise rate and low-noise LPN problems used in PCG-like protocols.

2. All the attacks listed in [BCGI18] do not take into account the field sizes. As a result, follow-up works also use the same analysis formula for fields of other sizes. We show that ISD (including Pooled Gauss as its special case) is the best attack to solve the low-noise LPN problems over finite fields, and the cost of ISD attack increases as the field size increases. We also find that the advantage of the advanced ISD algorithm over Pooled Gauss attack gradually disappears as the field size increases (see Figure 3 of Section 5). Particularly, when the field size is large enough, ISD has the same cost as Pooled Gauss.

The very recent work [BØ23] proposed a new algebraic attack that is able to take advantage of regular noise distributions. ⁵ In this case, our analysis results will slightly change. In particular, for the LPN problems with regular noises and small code rates, their algebraic attack outperforms ISD attacks for many parameter sets (see Section 5.4 for details).

Compared to the analysis [BCGI18], our new analysis shows that one can reduce the weight parameter t by 20% – 40% when keeping the same security level. This directly allows us to obtain the efficiency improvement of PCG-like COT and VOLE protocols, which in turn improves the MPC, ZK and PSI applications.

2 Preliminary

2.1 Notation

We will use κ and ρ to denote the computational and statistical security parameters, respectively. We denote by \log the logarithm in base 2. For $a, b \in \mathbb{N}$ with $a \leq b$, we write $[a, b] = \{a, \dots, b\}$ and use $[n]$ to denote

⁴The GV bound decoding over \mathbb{F}_2 is to solve LPN instances that achieve the GV relative distance $t/N = \mathbf{H}^{-1}(1 - k/N)$, where $\mathbf{H}(\mu) = \mu \cdot \log(1/\mu) + (1 - \mu) \cdot \log(1/(1 - \mu))$ is the binary entropy function and \mathbf{H}^{-1} is the inverse of \mathbf{H} .

⁵The work [BØ23] was posted online one day before the Crypto’23 submission deadline.

$[0, n - 1]$ for simplicity. We use $x \leftarrow S$ to denote sampling x uniformly at random from a set S and $x \leftarrow \mathcal{D}$ to denote sampling x according to a distribution \mathcal{D} . For a ring \mathcal{R} , we denote by $|\mathcal{R}|$ the size of \mathcal{R} . By slightly abusing the notation, for a vector \mathbf{a} , we use $|\mathbf{a}|$ to denote the Hamming weight of \mathbf{a} , and denote by $\mathbf{a}[i]$ the i -th component of \mathbf{a} . For a vector $\mathbf{a} \in (\mathbb{Z}_{2^\lambda})^k$, we use $\text{BitDecomp}(\mathbf{a})$ to denote the bit-decomposition of \mathbf{a} , and the output of $\text{BitDecomp}(\mathbf{a})$ is denoted by $(\mathbf{a}^0, \mathbf{a}^1, \dots, \mathbf{a}^{\lambda-1})$ such that $\mathbf{a}^i \in \mathbb{F}_2^k$ for each $i \in [\lambda]$ and $(\mathbf{a}^0[j], \mathbf{a}^1[j], \dots, \mathbf{a}^{\lambda-1}[j])$ is the bit-decomposition of ring element $\mathbf{a}[j] \in \mathbb{Z}_{2^\lambda}$ for each $j \in [k]$. Let $\text{BitDecomp}^{-1}(\mathbf{a}^0, \mathbf{a}^1, \dots, \mathbf{a}^{\lambda-1}) = \sum_{i=0}^{\lambda-1} 2^i \cdot \mathbf{a}^i \in (\mathbb{Z}_{2^\lambda})^k$ be the inverse of $\text{BitDecomp}(\mathbf{a})$. We use $\text{poly}(\cdot)$ to denote a polynomial function. We will use the following lemma:

Lemma 1 (see, e.g., [YS16]). *For any $\mu \in (0, 1)$, if each coordinate of a vector $\mathbf{v} \in \mathbb{F}_2^t$ is independently set to 1 with probability μ , then the probability that $|\mathbf{v}| = \lceil \mu t \rceil$ is at least $\Omega(1/\sqrt{t})$.*

2.2 Learning Parity with Noise

Recently, variants of the Learning Parity with Noise (LPN) assumption [BFKL94] are used to build PCG-like protocols with sublinear communication for generating (C)OT and (V)OLE correlations. The LPN variants are defined over a general finite ring \mathcal{R} . The known LPN-based protocols mainly consider three cases for the choices of ring \mathcal{R} :

- Case 1 that $\mathcal{R} = \mathbb{F}_2$ is used to design the COT protocols [BCG⁺19b, BCG⁺19a, YWL⁺20, CRR21, BCG⁺22], which is in turn able to be transformed into standard OT protocols.
- Case 2 that \mathcal{R} is a large finite field \mathbb{F} with $|\mathbb{F}| \geq 2^\rho$ is used to construct the VOLE protocols [BCGI18, BCG⁺19a, SGRR19, WYKW21, CRR21, BCG⁺22] and the OLE protocol [BCG⁺19b].
- Case 3 that $\mathcal{R} = \mathbb{Z}_{2^\lambda}$ (e.g., $\lambda \in \{32, 64, 128\}$) is used to obtain the VOLE protocols [SGRR19, BBMH⁺21, BBMHS22].

For the other case that \mathcal{R} is a small finite field (e.g., $\mathcal{R} = \mathbb{F}_{2^8}$), it is still interesting to design some subfield VOLE protocols [SGRR19, WYKW21]. When considering some more general rings such as $\mathcal{R} = \mathbb{Z}_{p^\lambda}$ for a prime $p > 2$ and $\mathcal{R} = \mathbb{Z}_{p^{\lambda_1} q^{\lambda_2}}$ for two primes p, q , the LPN problems over such rings may be useful for future protocols.

Following the previous work (e.g., [BCG⁺19b, BCG⁺19a]), we define the (primal-)LPN and dual-LPN assumptions over a general ring \mathcal{R} as follows:

Definition 1 (LPN). *Let $\mathcal{D}(\mathcal{R}) = \{\mathcal{D}_{t,N}(\mathcal{R})\}_{t,N \in \mathbb{N}}$ denote a family of distributions over a ring \mathcal{R} such that for any $t, N \in \mathbb{N}$, $\text{Im}(\mathcal{D}_{t,N}(\mathcal{R})) \subseteq \mathcal{R}^N$. Let \mathbf{C} be a probabilistic code generation algorithm such that $\mathbf{C}(k, N, \mathcal{R})$ outputs a matrix $\mathbf{A} \in \mathcal{R}^{N \times k}$. For dimension $k = k(\kappa)$, number of samples $N = N(\kappa)$, Hamming weight of a noise vector $t = t(\kappa)$, and a ring \mathcal{R} , we say that the decisional $(\mathcal{D}, \mathbf{C}, \mathcal{R})$ -LPN(k, N, t) problem is (T, ϵ) -hard if for every probabilistic distinguisher \mathcal{B} running in time T , we have*

$$\left| \Pr_{\mathbf{A}, \mathbf{s}, \mathbf{e}} \left[\mathcal{B}(\mathbf{A}, \mathbf{b} \stackrel{\text{def}}{=} \mathbf{A} \cdot \mathbf{s} + \mathbf{e}) = 1 \right] - \Pr_{\mathbf{A}, \mathbf{u}} \left[\mathcal{B}(\mathbf{A}, \mathbf{u}) = 1 \right] \right| \leq \epsilon,$$

where $\mathbf{A} \leftarrow \mathbf{C}(k, N, \mathcal{R})$, $\mathbf{s} \leftarrow \mathcal{R}^k$, $\mathbf{e} \leftarrow \mathcal{D}_{t,N}(\mathcal{R})$ and $\mathbf{u} \leftarrow \mathcal{R}^N$. We say that the computational $(\mathcal{D}, \mathbf{C}, \mathcal{R})$ -LPN(k, N, t) problem is (T, ϵ) -hard if for every probabilistic algorithm \mathcal{B} running in time T , we have

$$\Pr_{\mathbf{A}, \mathbf{s}, \mathbf{e}} \left[\mathcal{B}(\mathbf{A}, \mathbf{b} \stackrel{\text{def}}{=} \mathbf{A} \cdot \mathbf{s} + \mathbf{e}) = (\mathbf{s}, \mathbf{e}) \right] \leq \epsilon,$$

where $\mathbf{A}, \mathbf{s}, \mathbf{e}$ are defined as above.

In the above definition, both T and ϵ are functions of computational security parameter κ . We mainly consider the following families of noise distributions:

- **Bernoulli.** Let $\text{Ber}(\mathcal{R}) = \{\text{Ber}_{\mu,N}(\mathcal{R})\}_{\mu,N}$ be the family of Bernoulli distributions. In particular, $\text{Ber}_{\mu,N}(\mathcal{R})$ is a Bernoulli distribution with parameters μ, N over a ring \mathcal{R} , such that each component in a vector sampled from $\text{Ber}_{\mu,N}(\mathcal{R})$ is a uniform element in \mathcal{R} with probability μ and 0 otherwise. The expected Hamming weight of the sampled vector is $t = \mu N(|\mathcal{R}| - 1)/|\mathcal{R}|$. Note that the definition is equivalent to sampling a uniform *non-zero* element in \mathcal{R} with probability $\mu(|\mathcal{R}| - 1)/|\mathcal{R}|$ for each component.
- **Exact.** Let $\text{HW}(\mathcal{R}) = \{\text{HW}_{t,N}(\mathcal{R})\}_{t,N}$ be the family of exact distributions. In particular, each component of a noise vector sampled from $\text{HW}_{t,N}(\mathcal{R})$ is a uniform non-zero element in t random positions and zero elsewhere.
- **Regular.** To achieve better efficiency, prior works [HOSS18, BCG⁺19b, BCG⁺19a, YWL⁺20, WYKW21, BCG⁺22] consider the family of *regular* noise distributions, denoted by $\text{RHW}(\mathcal{R}) = \{\text{RHW}_{t,N}(\mathcal{R})\}_{t,N}$. In addition to fixed Hamming weight, the noise vector further divides into t consecutive sub-vectors of size $\lfloor N/t \rfloor$, where each sub-vector has a single noisy coordinate.

The existing LPN-based PCG-like protocols adopt the latter two noise distributions, and the standard LPN assumption adopts the Bernoulli distribution. While the standard LPN assumption uses random linear codes to instantiate \mathbf{C} (i.e., sampling \mathbf{A} uniformly at random), these LPN-based protocols adopt other kinds of linear codes to obtain faster computation. The other linear codes to instantiate \mathbf{C} include: local linear codes [Ale03], LDPC codes and its variant [IKOS08, CRR21], quasi-cyclic codes [AMBD⁺18], MDPC codes [MTSB13], Druk-Ishai codes [DI14], etc. In this paper, we did not analyze the hardness of the LPN problems based on quasi-cyclic codes, which needs to take into account the effect of the DOOM attack [Sen11] that allows providing \sqrt{N} computational speedup. We are *not* aware that other kinds of linear codes listed above lead to significantly better attacks, compared to random linear codes. To simplify the notation, we often omit \mathbf{C} from the $(\mathcal{D}, \mathbf{C}, \mathcal{R})$ -LPN(k, N, t) problem, and only write $(\mathcal{D}, \mathcal{R})$ -LPN(k, N, t).

Below, we define the dual-LPN assumption over a general finite ring \mathcal{R} with a family \mathcal{D} of noise distributions, where both the decisional version and search version are described. The dual-LPN is also known as syndrome decoding.

Definition 2 (Dual LPN). *Let $\mathcal{D}(\mathcal{R})$ and \mathbf{C} be as in Definition 1. For two integers N, n with $N > n$, we define*

$$\mathbf{C}^\perp(N, n, \mathcal{R}) = \{\mathbf{H} \in \mathcal{R}^{n \times N} : \mathbf{H} \cdot \mathbf{A} = 0, \mathbf{A} \in \mathbf{C}(N - n, N, \mathcal{R}), \text{rank}(\mathbf{H}) = n\}.$$

For output length $n = n(\kappa)$, number of samples $N = N(\kappa)$, noise-vector Hamming weight $t = t(\kappa)$, we say that the decisional $(\mathcal{D}, \mathbf{C}^\perp, \mathcal{R})$ -dual-LPN(N, n, t) problem is (T, ϵ) -hard if for every PPT distinguisher \mathcal{B} running in time T :

$$\left| \Pr_{\mathbf{H}, \mathbf{e}} [\mathcal{B}(\mathbf{H}, \mathbf{H} \cdot \mathbf{e}) = 1] - \Pr_{\mathbf{H}, \mathbf{u}} [\mathcal{B}(\mathbf{H}, \mathbf{u}) = 1] \right| \leq \epsilon,$$

where $\mathbf{H} \leftarrow \mathbf{C}^\perp(N, n, \mathcal{R})$, $\mathbf{e} \leftarrow \mathcal{D}_{t,N}(\mathcal{R})$ and $\mathbf{u} \leftarrow \mathcal{R}^N$.

We say that the computational $(\mathcal{D}, \mathbf{C}^\perp, \mathcal{R})$ -dual-LPN(N, n, t) problem is (T, ϵ) -hard if for every probabilistic algorithm \mathcal{B} running in time T , we have

$$\Pr_{\mathbf{H}, \mathbf{e}} [\mathcal{B}(\mathbf{H}, \mathbf{H} \cdot \mathbf{e}) = \mathbf{e}] \leq \epsilon,$$

where \mathbf{H}, \mathbf{e} are defined as above.

For any fixed code generation algorithm \mathbf{C} and noise distribution \mathcal{D} , the dual-LPN problem defined as above is equivalent to the primal-LPN problem from Definition 1 with dimension $k = N - n$ and the number of samples N . The direction transforming an LPN instance into a dual-LPN instance directly follows the simple fact that $\mathbf{H} \cdot (\mathbf{A} \cdot \mathbf{s} + \mathbf{e}) = (\mathbf{H} \cdot \mathbf{A}) \cdot \mathbf{s} + \mathbf{H} \cdot \mathbf{e} = \mathbf{H} \cdot \mathbf{e}$, as \mathbf{H} is the parity-check matrix of the code generated by \mathbf{A} . The reverse direction can be obtained in a way similar to [MM11, Lemma 4.9]. As such, we often write the $(\mathcal{D}, \mathcal{R})$ -dual-LPN(N, n, t) problem for simplicity. We give an overview of some attacks on the LPN problem in Appendix C.

3 The Hardness of LPN with Regular Noise Distributions

Many MPC protocols [HOSS18, BCG⁺19b, BCG⁺19a, YWL⁺20, WYKW21] rely on the hardness of low-noise LPN under the regular noise distribution (RHW, \mathbb{F}_q)-LPN(k, N, t), where the noise vector is divided into t blocks, each of length $\lfloor N/t \rfloor$ and Hamming weight 1. Intuitively, LPN with regular noise may look easier than that with exact noise. For example, the LPN over \mathbb{F}_2 with regular noise simply leaks t parity bits of its secret and thus reduces the dimension by t .

Recently, Feneuil et al. [FJR22] introduced a reduction from the (dual)-LPN with a regular noise distribution to that with an exact noise distribution⁶.

Theorem 1 (Theorem 1 of [FJR22], adapted). *If the (HW, \mathbb{F}_q)-problem is (T, ϵ) -hard, then the (RHW, \mathbb{F}_q)-problem is $(T, \epsilon \cdot \binom{N}{t} / (\frac{N}{t})^t)$ -hard, where the statement holds for both versions of LPN, i.e., problem $\in \{\text{dual-LPN}(N, n, t), \text{LPN}(k, N, t)\}$.*

The reduction suffers significant security loss, i.e., the penalty factor

$$p_t = \binom{N}{t} / \left(\frac{N}{t}\right)^t = \frac{t^t}{t!} \cdot \prod_{i=1}^{t-1} \left(1 - \frac{i}{N}\right) = e^{t - \Theta(\ln t) - \Theta(t^2/N)} = e^{t(1-o(1))},$$

where the approximation is based on Stirling's $\ln(t!) = t \ln t - t + \Theta(\ln t)$, $t = o(N)$, and $4^{-x} \leq 1 - x \leq e^{-x}$ for $0 \leq x \leq 1/2$. Meanwhile, it is not hard to see that for many non-trivial parameter choices $\epsilon > e^{-t}$. Consider the dual LPN problem

$$[\mathbf{H}_1, \mathbf{H}_2] \cdot (\mathbf{e}_1 \| \mathbf{e}_2) = \mathbf{H}_1 \mathbf{e}_1 + \mathbf{H}_2 \mathbf{e}_2 = \mathbf{y}$$

where $\mathbf{H}_1 \in \mathbb{F}_q^{n \times n}$, $\mathbf{H}_2 \in \mathbb{F}_q^{n \times (N-n)}$, $\mathbf{e}_1 \in \mathbb{F}_q^n$ and $\mathbf{e}_2 \in \mathbb{F}_q^{N-n}$. A polynomial-time attack simply bets $\mathbf{e}_2 = \mathbf{0}$ and computes $\mathbf{e}_1 = \mathbf{H}_1^{-1} \mathbf{y}$ (assuming WLOG invertible \mathbf{H}_1), which succeeds with probability

$$\binom{n}{t} / \binom{N}{t} = \prod_{i=1}^{N-n} \left(\frac{n-t+i}{n+i}\right) > \left(1 - \frac{t}{n+1}\right)^{N-n} \approx e^{-\frac{t(N-n)}{n+1}}.$$

E.g., for $N \leq 2n$ the term $\epsilon > e^{-t}$ cannot afford multiplying with p_t , which renders the bound meaningless. This is just the ϵ for $T = \text{poly}(n, N)$, and for larger T it only becomes worse due to the known optimizations and trade-offs between T and ϵ (see, e.g., [TS16]). The case for (primal)-LPN is likewise.

In retrospect, [FJR22, Theorem 1] incurs significant loss because it simply uses $1/p_t$ to account for the probability that an exact noise vector is regular at the same time. We provide a new reduction with a new parameter α such that [FJR22, Theorem 1] can be seen as a special case of $\alpha = 1$. More importantly, with large α we are able to reduce the security loss dramatically by dividing the exponent by α , while paying only an additive price αt in dimension and number of samples.

⁶In particular, [FJR22] considers a d -split noise, which consists of d blocks of length N/d and each block has weight t/d . For $d = t$, it corresponds to the (most often used) case of regular noise.

3.0.1 The high-level intuition.

LPN samples $(\mathbf{A}, \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e})$ with regular noise of weight t can be divided into t equal-size blocks $(\mathbf{A}_i, \mathbf{b}_i = \mathbf{A}_i\mathbf{s} + \mathbf{e}_i)$ ($i \in [1, t]$), and for $q = 2$ the adversary can multiply every block by an all-ones row vector $\mathbf{1}^\top$ (recall $\mathbf{1}^\top \mathbf{e}_i = 1$) and therefore reduce the dimension by t (seen by applying a change of basis). However, we are going to prove the other direction, i.e., given LPN samples $(\mathbf{A}, \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e})$ with dimension k and noise of exact weight t , we divide them into αt blocks with a new parameter α , i.e., $(\mathbf{A}_i, \mathbf{b}_i = \mathbf{A}_i\mathbf{s} + \mathbf{e}_i)$ for $i \in [1, \alpha t]$. Instead of hoping that every \mathbf{e}_i has exact weight 1 (as Feneuil et al. did in [FJR22] that incurs significant security loss), we relax the condition to $|e_i| \leq 1$, which occurs with higher probability (and hence less security loss) especially for large α . For each block, we add an extra random sample (\mathbf{a}_i, v_i) with noise element \tilde{e}_i such that the vector $(\mathbf{e}_i^\top, \tilde{e}_i)$ has the exact weight 1 (i.e., the resulting noise is regular). This is possible if the dimension of the target LPN with regular noise is of dimension $k + \alpha t$. That is, the additional αt bits of uncertainty would help in simulating αt values of v_i almost perfectly.

Theorem 2. *Let $t, N \in \mathbb{N}$, and $\alpha \geq 2$ such that $\alpha t \in \mathbb{N}$ and $(\alpha t) | N$. If the $(\text{HW}, \mathbb{F}_q)$ -LPN(k, N, t) is (T, ϵ) -hard, then the $(\text{RHW}, \mathbb{F}_q)$ -LPN($k + \alpha t, N + \alpha t, \alpha t$) is $(T - \text{poly}(k, N), 2^{\frac{t}{\alpha}} \cdot \epsilon)$ -hard.*

Proof. Let $N = \alpha t m$ for $m \in \mathbb{N}$. On input of $(\text{HW}, \mathbb{F}_q)$ -LPN(k, N, t) samples parsed as αt blocks:

$$\mathbf{A} \stackrel{\text{def}}{=} \begin{bmatrix} \mathbf{A}_1 \in \mathbb{F}_q^{m \times k} \\ \vdots \\ \mathbf{A}_{\alpha t} \in \mathbb{F}_q^{m \times k} \end{bmatrix}, \mathbf{b} \stackrel{\text{def}}{=} \begin{bmatrix} \mathbf{b}_1 = (\mathbf{A}_1 \mathbf{s} + \mathbf{e}_1) \in \mathbb{F}_q^m \\ \vdots \\ \mathbf{b}_{\alpha t} = (\mathbf{A}_{\alpha t} \mathbf{s} + \mathbf{e}_{\alpha t}) \in \mathbb{F}_q^m \end{bmatrix}, \text{ where } \mathbf{s} \in \mathbb{F}_q^k.$$

Our analysis is conditioned on the event \mathcal{E} that the Hamming weight $|e_i| \leq 1$ for every $1 \leq i \leq \alpha t$, which occurs with probability

$$\Pr_{[\mathbf{e}_1^\top \dots \mathbf{e}_{\alpha t}^\top] \leftarrow \text{HW}_{t, N}(\mathbb{F}_q)}[\mathcal{E}] = \frac{\binom{\alpha t}{t} \cdot \left(\frac{N}{\alpha t}\right)^t}{\binom{N}{t}} = \prod_{i=1}^{t-1} \frac{(1 - \frac{i}{\alpha t})}{(1 - \frac{i}{N})} > \frac{4^{\sum_{i=1}^{t-1} -\frac{i}{\alpha t}}}{1} = 2^{\frac{1}{\alpha} - \frac{t}{\alpha}},$$

where the inequality is due to $1 - x \geq 4^{-x}$ for $0 \leq x \leq 1/2$, and $x = \frac{i}{\alpha t} < \frac{1}{\alpha} \leq 1/2$. Sample row vectors $\mathbf{r}_1^\top \stackrel{\$}{\leftarrow} \mathbb{F}_q^{k+\alpha t}, \dots, \mathbf{r}_{\alpha t}^\top \stackrel{\$}{\leftarrow} \mathbb{F}_q^{k+\alpha t}$. Condition on that they are linearly independent, which has probability more than $1 - 2^{-k+1}$ (see, e.g., [YS16, Fact A.2] for $q = 2$), pick any full-rank $\mathbf{B} \in \mathbb{F}_q^{k \times (k+\alpha t)}$ such that the square matrix \mathbf{M} has full rank over \mathbb{F}_q , where matrix \mathbf{M} is defined as follows:

$$\mathbf{M} \stackrel{\text{def}}{=} \begin{bmatrix} \mathbf{B} \\ \mathbf{r}_1^\top \\ \vdots \\ \mathbf{r}_{\alpha t}^\top \end{bmatrix} \in \mathbb{F}_q^{(k+\alpha t) \times (k+\alpha t)}.$$

$$\text{Let } \mathbf{C}_i \stackrel{\text{def}}{=} \begin{bmatrix} \mathbf{A}_i \mathbf{B} \\ \mathbf{r}_i^\top - \mathbf{1}^\top (\mathbf{A}_i \mathbf{B}) \end{bmatrix}, \mathbf{b}'_i \stackrel{\text{def}}{=} \begin{bmatrix} \mathbf{b}_i = \mathbf{A}_i \mathbf{B} \mathbf{x} + \mathbf{e}_i \\ \underbrace{\mathbf{r}_i^\top \mathbf{x} + u_i - \mathbf{1}^\top \mathbf{b}_i}_{=v_i} \end{bmatrix} \text{ for } i \in [1, \alpha t],$$

where $\mathbf{x} \stackrel{\$}{\leftarrow} \mathbb{F}_q^{k+\alpha t}$ subject to $\mathbf{s} = \mathbf{B}\mathbf{x}$, $\mathbf{1}^\top$ is the all-ones row vector (i.e., every coordinate is 1), $u_i \stackrel{\$}{\leftarrow} \mathbb{F}_q/\{0\}$ is the non-zero entry of \mathbf{e}_i if $|e_i| = 1$, and otherwise (i.e., $|e_i| = 0$) pick a fresh $u_i \stackrel{\$}{\leftarrow} \mathbb{F}_q/\{0\}$. We can verify $\mathbf{b}'_i = \mathbf{C}_i \mathbf{x} + \begin{bmatrix} \mathbf{e}_i \\ u_i - \mathbf{1}^\top \mathbf{e}_i \end{bmatrix}$, where $[\mathbf{e}_i^\top \quad u_i - \mathbf{1}^\top \mathbf{e}_i]$ is regular (of Hamming weight 1) as desired ⁷.

⁷Strictly speaking, the noise vector is ensured to have Hamming weight 1, but its coordinates may not take non-zero values with equal probability. The issue can be easily addressed by shuffling the matrices and samples accordingly.

By the definition of \mathbf{M} we have that $(\mathbf{r}_1^\top \mathbf{x}, \dots, \mathbf{r}_{\alpha t}^\top \mathbf{x})$ is uniformly random over $\mathbb{F}_q^{\alpha t}$ even conditioned on \mathbf{M} , $\mathbf{B}\mathbf{x}$ and other variable (e.g., all the \mathbf{A}_i 's, \mathbf{e}_i 's, u_i 's). Thus, with no knowledge about u_i and \mathbf{e}_i , the reduction can simply use random $v_i \xleftarrow{\$} \mathbb{F}_q$ and let $v_i = \mathbf{r}_i^\top \mathbf{x} + u_i - \mathbf{1}^\top \mathbf{b}_i$.

CASE: $q = 2$. u_i is always 1 (the only non-zero element in \mathbb{F}_2). Sample random $\mathbf{P}_i \xleftarrow{\$} \mathbb{F}_q^{m \times \alpha t}$ for $i \in [1, \alpha t]$. The following are LPN samples with the same regular noise $[\mathbf{e}_i^\top \quad u_i - \mathbf{1}^\top \mathbf{e}_i]$, which can be verified by comparing its difference with $(\mathbf{C}_i, \mathbf{b}'_i)$.

$$\left(\begin{array}{c} \left[\begin{array}{c} [\mathbf{A}_i \parallel \mathbf{P}_i] \cdot \mathbf{M} \\ \mathbf{r}_i^\top - \mathbf{1}^\top (\mathbf{A}_i \mathbf{B}) \end{array} \right], \left[\begin{array}{c} \mathbf{b}_i \\ v_i \end{array} \right] + \left[\mathbf{P}_i \cdot \left[\begin{array}{c} \mathbf{1}^\top \mathbf{b}_1 + v_1 - 1 \\ \vdots \\ \mathbf{1}^\top \mathbf{b}_{\alpha t} + v_{\alpha t} - 1 \\ \mathbf{0} \end{array} \right] \right] \end{array} \right), \quad (1)$$

where $\mathbf{r}_i^\top \mathbf{x} = \mathbf{1}^\top \mathbf{b}_i + v_i - 1$. Further, the above matrices $[\mathbf{A}_1 \parallel \mathbf{P}_1] \cdot \mathbf{M}, \dots, [\mathbf{A}_{\alpha t} \parallel \mathbf{P}_{\alpha t}] \cdot \mathbf{M}, \mathbf{r}_1^\top, \dots, \mathbf{r}_{\alpha t}^\top$ are 2^{-k+2} -close to uniform, as conditioned on any invertible \mathbf{M} (with probability more than $1 - 2^{-k+1}$) $[\mathbf{A}_i \parallel \mathbf{P}_i] \cdot \mathbf{M}$ is uniformly random if $[\mathbf{A}_i \parallel \mathbf{P}_i]$ is (see Lemma 4 in Appendix A). Therefore, the shuffled samples of the above is 2^{-k+1} -close to the i -th ($1 \leq i \leq \alpha t$) block of the $(\text{RHW}, \mathbb{F}_q)$ -LPN($k + \alpha t, N + \alpha t, \alpha t$) instance. Therefore, we just feed all αt blocks as per (1) (and shuffle the samples within each block) to the $(\text{RHW}, \mathbb{F}_q)$ -LPN solver, and if it returns \mathbf{x} then recover $\mathbf{s} := \mathbf{B}\mathbf{x}$, as the secret of $(\text{HW}, \mathbb{F}_q)$ -LPN. Quantitatively, if one breaks $(\text{RHW}, \mathbb{F}_q)$ -LPN with probability p , then it can also break $(\text{HW}, \mathbb{F}_q)$ -LPN with probability at least $2^{\frac{1}{\alpha} - \frac{t}{\alpha}}(p - 2^{-k+2}) \geq p \cdot 2^{-\frac{t}{\alpha}}$.

CASE: $q \geq 2$. u_i is uniform over $\{1, \dots, q-1\}$. We cannot assume any knowledge about u_i in the reduction, but the secret can absorb u_i . Define \mathbf{x}' such that $\mathbf{B}\mathbf{x}' \equiv \mathbf{B}\mathbf{x}$ and $\forall i \in [1, \alpha t] : \mathbf{r}_i^\top \mathbf{x}' \equiv \mathbf{r}_i^\top \mathbf{x} + u_i - 1$, i.e.,

$$\mathbf{M}\mathbf{x}' \equiv \mathbf{M}\mathbf{x} + \underbrace{[0, \dots, 0, u_1 - 1, \dots, u_{\alpha t} - 1]^\top}_{\mathbf{h} \stackrel{\text{def}}{=} \mathbf{h}},$$

which is always possible for any invertible \mathbf{M} , by letting $\mathbf{x}' \stackrel{\text{def}}{=} \mathbf{x} + \mathbf{M}^{-1}\mathbf{h}$. Therefore, the reduction in the $q = 2$ case works in the general case as well by considering \mathbf{x}' instead of \mathbf{x} , where $\mathbf{B}\mathbf{x}' = \mathbf{s}$ and $\mathbf{r}_i^\top \mathbf{x}' = \mathbf{1}^\top \mathbf{b}_i + v_i - 1$ just like in the $q = 2$ case. \square

We get similar result for dual-LPN in Corollary 1 via a simple reduction from LPN to dual-LPN, i.e., $(\mathbf{A}, \mathbf{b} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e})$ can be efficiently transformed into a dual-LPN instance $(\mathbf{H}, \mathbf{y} = \mathbf{H} \cdot \mathbf{b} = \mathbf{H} \cdot \mathbf{e})$ due to the duality between \mathbf{H} and \mathbf{A} , regardless of the distribution of \mathbf{e} .

Corollary 1. *Let $t, N \in \mathbb{N}$ and $\alpha \geq 2$ such that $\alpha t \in \mathbb{N}$ and $(\alpha t) | N$. If the $(\text{HW}, \mathbb{F}_q)$ -LPN(k, N, t) is (T, ϵ) -hard, then the $(\text{RHW}, \mathbb{F}_q)$ -dual-LPN($N + \alpha t, N - k, \alpha t$) is $(T - \text{poly}(k, N), 2^{\frac{t}{\alpha}} \cdot \epsilon)$ -hard.*

4 The Hardness of LPN over Integer Rings

LPN over an integer ring has been used in multiple prior works. For example, Schoppmann et al. [SGRR19] used it to construct VOLE over $\mathbb{Z}_{2^{64}}$, which can then be used for performing private machine learning tasks in the semi-honest setting. Recently, Baum et al. [BBMH⁺21, BBMHS22] proposed VOLE protocols based on LPN over a ring \mathbb{Z}_{2^λ} , and used them to construct concretely efficient ZK protocols on arithmetic circuits over \mathbb{Z}_{2^λ} . These VOLE protocols could also benefit other works that need VOLE over integer rings like the MPC protocol SPD \mathbb{Z}_{2^k} [CDE⁺18]. The current security estimate of LPN over \mathbb{Z}_{2^λ} in prior work is directly adapted from that for LPN over a field \mathbb{F} of size $|\mathbb{F}| \approx 2^\lambda$ [BCGI18]. As we will show in this section the hardness of LPN over \mathbb{Z}_{2^λ} is more related to that over \mathbb{F}_2 (rather than that over the λ -bit field). As depicted in Figure 2, we provide the following reductions between the hardness of LPN over \mathbb{Z}_{2^λ} and that over \mathbb{F}_2 .

- **Decisional LPN over $\mathbb{Z}_{2^\lambda} \rightarrow$ Decisional LPN over \mathbb{F}_2 .** We show that distinguishing LPN samples over \mathbb{Z}_{2^λ} with noise weight t from exact noise distribution is no harder than distinguishing LPN over \mathbb{F}_2 with noise weight $\frac{2^{(\lambda-1)}}{2^\lambda-1} \cdot t \approx t/2$. This reduction directly gives an attack that reduces the noise weight by half for an LPN instance over \mathbb{Z}_{2^λ} .
- **Decisional LPN over $\mathbb{F}_2 \rightarrow$ Decisional LPN over \mathbb{Z}_{2^λ} .** We show that distinguishing LPN over \mathbb{F}_2 with noise weight t is no harder than the distinguishing attack on LPN over \mathbb{Z}_{2^λ} with 1) non-standard Bernoulli-like integer noise of weight at most $\lambda \cdot t$; and 2) standard Bernoulli noise of weight $\approx 2^\lambda \cdot t$.
- **Computational LPN over $\mathbb{Z}_{2^\lambda} \rightarrow$ Computational LPN over \mathbb{F}_2 .** We show that a secret recovery attack on LPN over \mathbb{Z}_{2^λ} with noise weight t is no harder than that on LPN over \mathbb{F}_2 with noise weight roughly $t/2$. While a generic reduction requires $k^{\omega(\lambda)}$ -hardness for LPN over \mathbb{Z}_{2^λ} , we also give more efficient reductions for their weakly one-wayness that is more relevant to practical attacks and security estimates. We also discuss how to optimize the secret recovery attack on LPN over \mathbb{Z}_{2^λ} based on that over \mathbb{F}_2 in practice.

We give similar reductions for the LPN over $\mathbb{Z}_{p^{\lambda_1}q^{\lambda_2}}$ (for distinct primes p, q) in Appendix B. When we give the reductions between different computational LPN variants, we assume that LPN over a field in consideration mostly has a unique solution in the average case (except for a negligible fraction), which will simplify the analysis. Note that this is true for most interesting parameter regimes of LPN, which give rise to cryptographic applications (e.g., PCG and public-key encryption).

Lemma 2 (Unique decoding of LPN over any finite field \mathbb{F}). *For any $N > k + 4t \log N$, the following is bounded by $\frac{N^{2t}}{|\mathbb{F}|^{N-k-2t}}$:*

$$\Pr_{\mathbf{A} \leftarrow \mathbb{F}^{k \times N}} \left[\exists \mathbf{s}_1 \neq \mathbf{s}_2 \in \mathbb{F}^k, \mathbf{e}_1, \mathbf{e}_2 \in \mathbb{F}^N : |\mathbf{e}_1|, |\mathbf{e}_2| \leq t \wedge (\mathbf{A} \cdot \mathbf{s}_1 + \mathbf{e}_1 = \mathbf{A} \cdot \mathbf{s}_2 + \mathbf{e}_2) \right].$$

Proof. Let $\mathbf{s} \stackrel{\text{def}}{=} \mathbf{s}_1 - \mathbf{s}_2 \in \mathbb{F}^k$ and $\mathbf{e} \stackrel{\text{def}}{=} \mathbf{e}_2 - \mathbf{e}_1 \in \mathbb{F}^N$. For any $\mathbf{s} \neq \mathbf{0}$, $\mathbf{A} \cdot \mathbf{s}$ is uniform over \mathbb{F}^N . Together with $|\mathbf{e}| \leq 2t$, the probability that $\mathbf{A} \cdot \mathbf{s} = \mathbf{e}$ is at most $\sum_{i=0}^{2t} \binom{N}{i} / |\mathbb{F}|^{N-i} \leq \left(\sum_{i=0}^{2t} \binom{N}{i} \right) / |\mathbb{F}|^{N-2t} \leq N^{2t} / |\mathbb{F}|^{N-2t}$. By a union bound on all possible $\mathbf{s} \in \mathbb{F}^k$, we obtain the bound claimed in the lemma. \square

4.1 Reduction from Decisional LPN over \mathbb{Z}_{2^λ} to LPN over \mathbb{F}_2

We start with a simple observation that the distinguishing attack on LPN over \mathbb{Z}_{2^λ} can be based on that over \mathbb{F}_2 with roughly halved noise weight. Specifically, we have the following theorem.

Theorem 3. *If decisional $(\mathcal{D}, \mathbb{Z}_{2^\lambda})$ -LPN(k, N, w_1) is (T, ϵ) -hard, then decisional $(\mathcal{D}, \mathbb{F}_2)$ -LPN(k, N, w_2) is $(T - \text{poly}(k, N), O(d\epsilon))$ -hard, where $(\mathcal{D}, w_1, w_2, d) \in \{(\text{HW}, t, \frac{2^{(\lambda-1)}}{2^\lambda-1}t, \sqrt{t}), (\text{Ber}, \mu, \mu, 1)\}$.*

Proof. Given LPN samples over ring \mathbb{Z}_{2^λ} ($\mathbf{A}, \mathbf{b} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e} \pmod{2^\lambda}$), we observe that least significant bits (LSBs) of these samples ($\mathbf{A}^0 := \mathbf{A} \pmod{2}, \mathbf{b}^0 := \mathbf{b} \pmod{2}$) constitute exactly the LPN samples over \mathbb{F}_2 for noise $\mathbf{e}^0 = \mathbf{e} \pmod{2}$. In case that $\mathbf{e} \leftarrow \text{HW}_{t, N}(\mathbb{Z}_{2^\lambda})$, the noise vector \mathbf{e}^0 follows a Bernoulli-like distribution over \mathbb{F}_2^N , which is sampled by first picking t out of N coordinates at random and then filling in these t coordinates with random non-zero elements over \mathbb{Z}_{2^λ} (and the rest with zeros). Thus, overall \mathbf{e}^0 has expected weight $t' = \frac{2^{(\lambda-1)}}{2^\lambda-1} \cdot t$, where $\frac{2^{(\lambda-1)}}{2^\lambda-1}$ is the probability that a random non-zero element of \mathbb{Z}_{2^λ} is odd. By Lemma 1, this implies that with probability $\Omega(1/\sqrt{t})$, the noise vector \mathbf{e}^0 follows the exact noise distribution $\text{HW}_{t', N}(\mathbb{F}_2)$. On the other hand, the LSBs of (\mathbf{A}, \mathbf{u}) with an exact $\mathbf{u} \in \mathbb{Z}_{2^\lambda}$ are uniform as

well. Therefore, one can use $(\text{HW}, \mathbb{F}_2)\text{-LPN}(k, N, t')$ solver to distinguish $(\mathbf{A}^0, \mathbf{b}^0)$ from uniform samples. The proof for the second statement is likewise, except when taking the LSBs of $e \leftarrow \text{Ber}_{\mu, N}(\mathbb{Z}_{2^\lambda})$ we immediately get $e^0 \sim \text{Ber}_{\mu, N}(\mathbb{F}_2)$ as desired. \square

Despite the preserved noise probability μ in the case of Bernoulli distribution, we note that $\text{Ber}_{\mu, N}(\mathbb{Z}_{2^\lambda})$ has expected weight $(1 - 2^{-\lambda})\mu N$, while $\text{Ber}_{\mu, N}(\mathbb{F}_2)$ has expected weight $\mu N/2$ that is roughly $2\times$ smaller than $\text{Ber}_{\mu, N}(\mathbb{Z}_{2^\lambda})$.

4.2 Reduction from LPN over \mathbb{F}_2 to Decisional LPN over \mathbb{Z}_{2^λ}

We first show that the LPN assumption over \mathbb{F}_2 implies that over \mathbb{Z}_{2^λ} under the standard Bernoulli noise distribution. However, we achieve the goal by paying a price in the security loss due to the dependence among different noise vectors. As a result, we get the very conservative statement that decisional LPN over \mathbb{F}_2 with noise weight t is no harder than decisional LPN over \mathbb{Z}_{2^λ} with noise weight roughly $2^\lambda t$. We then introduce more useful Bernoulli-like noise distributions to enable more efficient reductions. In particular, we can reduce to an LPN over \mathbb{Z}_{2^λ} with noise weight λt (rather than $2^\lambda t$).

Theorem 4. *If decisional $(\text{Ber}, \mathbb{F}_2)\text{-LPN}(k, N, \mu/2^\lambda)$ is (T, ϵ) -hard, then decisional $(\text{Ber}, \mathbb{Z}_{2^\lambda})\text{-LPN}(k, N, \mu)$ is $(T - \text{poly}(k, N), \lambda\epsilon)$ -hard.*

Proof. Let $(\mathbf{A}, \mathbf{b} = \mathbf{A} \cdot \mathbf{s} + e \pmod{2^\lambda})$ be LPN samples over \mathbb{Z}_{2^λ} . Decompose the matrix and vectors into λ ones over \mathbb{F}_2 as follows: $(\mathbf{A}^0, \mathbf{A}^1, \dots, \mathbf{A}^{\lambda-1}) := \text{BitDecomp}(\mathbf{A})$, $(\mathbf{s}^0, \mathbf{s}^1, \dots, \mathbf{s}^{\lambda-1}) := \text{BitDecomp}(\mathbf{s})$, $(\mathbf{e}^0, \mathbf{e}^1, \dots, \mathbf{e}^{\lambda-1}) := \text{BitDecomp}(\mathbf{e})$ and $(\mathbf{b}^0, \mathbf{b}^1, \dots, \mathbf{b}^{\lambda-1}) := \text{BitDecomp}(\mathbf{b})$. Therefore, by the modular 2^λ arithmetic, \mathbf{b}^i depends only on \mathbf{A} , $(\mathbf{s}^i, \dots, \mathbf{s}^0)$, $(\mathbf{e}^i, \dots, \mathbf{e}^0)$, and we write it as

$$\mathbf{b}^i = \mathbf{A}^0 \mathbf{s}^i + \mathbf{e}^i + f_i(\mathbf{A}, \mathbf{S}(0, i-1), \mathbf{E}(0, i-1)) \pmod{2},$$

where $\mathbf{S}(0, i-1) \stackrel{\text{def}}{=} (\mathbf{s}^{i-1}, \dots, \mathbf{s}^0)$, and $\mathbf{E}(0, i-1) \stackrel{\text{def}}{=} (\mathbf{e}^{i-1}, \dots, \mathbf{e}^0)$, and f_i sums up the other terms not depending on \mathbf{s}^i and \mathbf{e}^i . Define the hybrid distributions

$$\begin{aligned} H_0 &= (\mathbf{A}, \mathbf{u}_0, \dots, \mathbf{u}_{i-1}, \mathbf{u}_i, \dots, \mathbf{u}_{\lambda-1}) \\ &\quad \vdots \\ H_i &= (\mathbf{A}, \mathbf{b}^0, \dots, \mathbf{b}^{i-1}, \mathbf{u}_i, \dots, \mathbf{u}_{\lambda-1}) \\ &\quad \vdots \\ H_\lambda &= (\mathbf{A}, \mathbf{b}^0, \dots, \mathbf{b}^{i-1}, \mathbf{b}^i, \dots, \mathbf{b}^{\lambda-1}) \end{aligned}$$

where $\mathbf{u}_j \leftarrow \mathbb{F}_2^N$ for $j \in [\lambda]$ is sampled independently at random. Note that all the \mathbf{s}^i 's are independent and uniformly random. Therefore, by the decisional $(\text{Ber}, \mathbb{F}_2)\text{-LPN}$ assumption

$$(\mathbf{A}^0, \mathbf{u}_i, \mathbf{S}(0, i-1), \mathbf{E}(0, i-1)) \sim_c (\mathbf{A}^0, \mathbf{A}^0 \mathbf{s}^i + \mathbf{e}^i \pmod{2}, \mathbf{S}(0, i-1), \mathbf{E}(0, i-1))$$

where $\mathbf{S}(0, i-1)$ is independent of any other variables, and the actual noise rate of LPN is that of \mathbf{e}^i conditioned on $\mathbf{E}(0, i-1)$ (see analysis below). This implies

$$(\mathbf{A}, \mathbf{b}^0, \dots, \mathbf{b}^{i-1}, \mathbf{u}_i + f_i(\mathbf{A}, \mathbf{S}(0, i-1), \mathbf{E}(0, i-1)) \pmod{2}) \sim_c (\mathbf{A}, \mathbf{b}^0, \dots, \mathbf{b}^{i-1}, \mathbf{b}^i)$$

which in turn implies $H_i \sim_c H_{i+1}$, where $\mathbf{b}^0, \dots, \mathbf{b}^{i-1}, f_i(\mathbf{A}, \mathbf{S}(0, i-1), \mathbf{E}(0, i-1))$ can be efficiently computed from $\mathbf{A}, \mathbf{S}(0, i-1), \mathbf{E}(0, i-1)$.

Therefore, if all the adjacent H_i and H_{i+1} are computationally indistinguishable, then H_0 and H_λ are computationally indistinguishable by a hybrid argument. It thus remains to estimate the noise rate need by the LPN assumption. Consider a single noise sample $(e^0[j], e^1[j], \dots, e^{\lambda-1}[j]) \leftarrow \text{Ber}_{\mu, N}(\mathbb{Z}_{2^\lambda})$, where $e^i[j]$ is the j -th bit of e^i . Conditioned on any non-zero $(e^0[j], \dots, e^{i-1}[j])$, $e^i[j]$ is uniformly random and thus unconditionally masks the corresponding $b^i[j]$. Otherwise, we have that

$$\Pr [e^i[j] = 1 \mid (e^0[j], \dots, e^{i-1}[j]) = 0^i] = \frac{\mu 2^{-(i+1)}}{1 - \mu + \mu 2^{-i}} \geq \mu 2^{-(i+1)}$$

is the noise rate needed to keep the computational indistinguishability between H_i and H_{i+1} , which reaches its minimum $\mu 2^{-\lambda}$ when $i = \lambda - 1$. \square

Based on the above theorem, we easily obtain the following corollary whose proof is given in Appendix A.

Corollary 2. *If decisional $(\text{Ber}, \mathbb{F}_2)$ -LPN($k, N, \mu/2^\lambda$) is hard, then computational $(\text{HW}, \mathbb{Z}_{2^\lambda})$ -LPN($k, N, t = (1 - 2^{-\lambda})\mu N$) is hard.*

In retrospect, the dependency among the noise vectors $\{e^i\}$ incurs significant losses during the reduction. This motivates us to introduce the specific noise distributions, $\text{IndBer}_{\mu, N}(\mathbb{Z}_{2^\lambda})$ and $\text{IndHW}_{t, N}(\mathbb{Z}_{2^\lambda})$, where Ind refers that the bit decomposition of the noise, $e^0, \dots, e^{\lambda-1}$ are independent and identically distributed, and parameter μ (resp., t) is noise rate (resp., weight) of each e^i .

- $\text{IndBer}_{\mu, N}(\mathbb{Z}_{2^\lambda})$ is bit-wise independent. By $e \leftarrow \text{IndBer}_{\mu, N}(\mathbb{Z}_{2^\lambda})$, we mean that $e := \sum_{i=0}^{\lambda-1} 2^i \cdot e^i$ with $e^i \leftarrow \text{Ber}_{\mu, N}(\mathbb{F}_2)$ for $i \in [\lambda]$. The noise rate of $\text{IndBer}_{\mu, N}(\mathbb{Z}_{2^\lambda})$ is the probability that each coordinate of e is non-zero, i.e., $1 - (1 - \mu/2)^\lambda \leq \lambda\mu/2$ by Bernoulli's inequality. Therefore, the expected Hamming weight of $e \leftarrow \text{IndBer}_{\mu, N}(\mathbb{Z}_{2^\lambda})$ is at most λt where $t = \mu N/2$.
- $\text{IndHW}_{t, N}(\mathbb{Z}_{2^\lambda})$ decomposes into λ independent vectors from $\text{HW}_{t, N}(\mathbb{F}_2)$. By $e \leftarrow \text{IndHW}_{t, N}(\mathbb{Z}_{2^\lambda})$, we mean that $e := \sum_{i=0}^{\lambda-1} 2^i \cdot e^i$ with $e^i \leftarrow \text{HW}_{t, N}(\mathbb{F}_2)$ for $i \in [\lambda]$. It is easy to see that the Hamming weight of e is at most λt .

Distributions $\text{IndBer}_{\mu, N}(\mathbb{Z}_{2^\lambda})$ and $\text{IndHW}_{t, N}(\mathbb{Z}_{2^\lambda})$ have not been used in existing protocols. However, LPN with such distributions can be used to design PCG-like protocols by running these protocols with maximum weight λt . Below, we show that decisional LPN over \mathbb{F}_2 with noise weight t is *tightly* equivalent to decisional LPN over \mathbb{Z}_{2^λ} with noise weight roughly λt under the new noise distributions.

Theorem 5. *Let $(\mathcal{D}_1, \mathcal{D}_2, w) \in \{(\text{Ber}, \text{IndBer}, \mu), (\text{HW}, \text{IndHW}, t)\}$ and we have:*

1. *If decisional $(\mathcal{D}_1, \mathbb{F}_2)$ -LPN(k, N, w) is (T, ϵ) -hard, then decisional $(\mathcal{D}_2, \mathbb{Z}_{2^\lambda})$ -LPN(k, N, w) is $(T - \text{poly}(k, N), \lambda\epsilon)$ -hard.*
2. *If decisional $(\mathcal{D}_2, \mathbb{Z}_{2^\lambda})$ -LPN(k, N, w) is (T, ϵ) -hard, then decisional $(\mathcal{D}_1, \mathbb{F}_2)$ -LPN(k, N, w) is $(T - \text{poly}(k, N), \epsilon)$ -hard.*

Proof. The proof of the first statement is similar to that of Theorem 4, except that now every e^i is independent of the previous e^0, \dots, e^{i-1} , where e^i follows $\text{Ber}_{\mu, N}(\mathbb{F}_2)$ by the definition of $e \leftarrow \text{IndBer}_{\mu, N}(\mathbb{Z}_{2^\lambda})$. The proof of the second statement can be trivially adapted from that of Theorem 3, i.e., $\mathbf{A}^0 = \mathbf{A} \pmod 2$, $\mathbf{A}^0 \cdot \mathbf{s}^0 + e^0 = \mathbf{A} \cdot \mathbf{s} + e \pmod 2$. \square

Algorithm 1: $\mathcal{A}_{\text{LPN}_{2^\lambda}}$, the secret recovery algorithm on LPN over \mathbb{Z}_{2^λ} ($\lambda \geq 2$) with oracle access to $\mathcal{A}_{\text{LPN}_2}$ (the solver for LPN over \mathbb{F}_2).

Input: $(\mathcal{D}, \mathbb{Z}_{2^\lambda})$ -LPN(k, N, t) samples $(\mathbf{A}, \mathbf{b} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e} \pmod{2^\lambda})$

Output: $\mathbf{s} \in \mathbb{Z}_{2^\lambda}$

- 1 $(\mathbf{A}^0, \mathbf{A}^1, \dots, \mathbf{A}^{\lambda-1}) := \text{BitDecomp}(\mathbf{A});$
 - 2 $(\mathbf{b}^0, \mathbf{b}^1, \dots, \mathbf{b}^{\lambda-1}) := \text{BitDecomp}(\mathbf{b});$
 - 3 $(\mathbf{s}^0, \mathbf{e}^0) \leftarrow \mathcal{A}_{\text{LPN}_2}(\mathbf{A}^0, \mathbf{b}^0);$
 - 4 $\mathbf{b}' := (\mathbf{b} - \mathbf{A} \cdot \mathbf{s}^0 - \mathbf{e}^0)/2;$
 - 5 **Return** $\mathbf{s} = \mathbf{s}^0 + 2 \cdot \mathcal{A}_{\text{LPN}_{2^{\lambda-1}}}(\mathbf{A}' := \sum_{i=0}^{\lambda-2} 2^i \cdot \mathbf{A}^i, \mathbf{b}')$;
-

On the choice of matrix \mathbf{A} . As we can see from the proofs of Theorem 4, Theorem 5 and Theorem 6 (shown in Section 4.3), all the reductions only rely on that \mathbf{A}^0 is uniformly distributed over $\mathbb{F}_2^{N \times k}$ while $\mathbf{A}^1, \dots, \mathbf{A}^{\lambda-1}$ can be arbitrary (or even zero matrix), where $(\mathbf{A}^0, \mathbf{A}^1, \dots, \mathbf{A}^{\lambda-1}) := \text{BitDecomp}(\mathbf{A})$. In other words, it suffices to use a Boolean matrix $\mathbf{A} = \mathbf{A}^0$, and the choices of $\mathbf{A}^1, \dots, \mathbf{A}^{\lambda-1}$ do not introduce any further hardness to the LPN problem over \mathbb{Z}_{2^λ} . Overall, we give a positive result that LPN over ring \mathbb{Z}_{2^λ} with Boolean matrices is secure if the corresponding LPN over binary field \mathbb{F}_2 is secure.

4.3 Reduction from Computational LPN over \mathbb{Z}_{2^λ} to LPN over \mathbb{F}_2

We show that an LPN instance over \mathbb{Z}_{2^λ} can be efficiently translated to λ instances of LPN over \mathbb{F}_2 , which are independent except that they share the same random matrix \mathbf{A}^0 over \mathbb{F}_2 and that the noise vectors of the λ instances are somehow correlated. We refer to the proof of Theorem 6 on how to address the correlation issue. Here we give a reduction from computational LPN over a ring \mathbb{Z}_{2^λ} to that over \mathbb{F}_2 by extending the corresponding reduction between their decisional versions shown in Section 4.1. Algorithm 1 shows how computational LPN over \mathbb{Z}_{2^λ} is reduced to that over $\mathbb{Z}_{2^{\lambda-1}}$ by recursion, whose correctness is analyzed in Lemma 3. Note that $\mathcal{A}_{\text{LPN}_{2^\lambda}}$ degenerates to secret recovery algorithm for LPN over \mathbb{F}_2 when $\lambda = 1$. Without loss of generality, we assume that $\mathcal{A}_{\text{LPN}_2}$ returns the noise vector in addition to the recovered secret.

Lemma 3. *Let $(\mathbf{A}, \mathbf{b} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e} \pmod{2^\lambda})$ be the LPN samples over \mathbb{Z}_{2^λ} , then $(\mathbf{A}', \mathbf{b}')$ as defined in Algorithm 1 constitute the LPN samples over $\mathbb{Z}_{2^{\lambda-1}}$, where $\mathbf{A}' = \sum_{i=0}^{\lambda-2} 2^i \cdot \mathbf{A}^i$, $\mathbf{b}' = \mathbf{A}' \cdot \mathbf{s}' + \mathbf{e}'$, $\mathbf{s}' = \sum_{i=1}^{\lambda-1} 2^{i-1} \cdot \mathbf{s}^i$ and $\mathbf{e}' = \sum_{i=1}^{\lambda-1} 2^{i-1} \cdot \mathbf{e}^i$*

Proof. Let $(\mathbf{A}^0, \mathbf{A}^1, \dots, \mathbf{A}^{\lambda-1})$ and $(\mathbf{b}^0, \mathbf{b}^1, \dots, \mathbf{b}^{\lambda-1})$ be the matrices and vectors defined in Algorithm 1. Let $(\mathbf{s}^0, \mathbf{s}^1, \dots, \mathbf{s}^{\lambda-1}) := \text{BitDecomp}(\mathbf{s})$ and $(\mathbf{e}^0, \mathbf{e}^1, \dots, \mathbf{e}^{\lambda-1}) := \text{BitDecomp}(\mathbf{e})$. Note that \mathbf{A}' is obtained from \mathbf{A} by truncating the most significant bits (MSBs), and thus follows the distribution in LPN over $\mathbb{Z}_{2^{\lambda-1}}$. It suffices to prove $\mathbf{b}' = \mathbf{A}' \cdot \mathbf{s}' + \mathbf{e}' \pmod{2^{\lambda-1}}$, where $\mathbf{s}' = \sum_{i=1}^{\lambda-1} 2^{i-1} \cdot \mathbf{s}^i$ and $\mathbf{e}' = \sum_{i=1}^{\lambda-1} 2^{i-1} \cdot \mathbf{e}^i$ are the secret and noise of LPN over $\mathbb{Z}_{2^{\lambda-1}}$ respectively. In particular, we have the following:

$$\begin{aligned} 2 \cdot \mathbf{b}' - 2 \cdot (\mathbf{A}' \cdot \mathbf{s}' + \mathbf{e}') &= \mathbf{b} - \mathbf{A} \cdot \mathbf{s}^0 - \mathbf{e}^0 - 2 \cdot (\mathbf{A}' \cdot \mathbf{s}' + \mathbf{e}') \\ &= \mathbf{b} - \mathbf{A} \cdot \mathbf{s}^0 - \mathbf{e}^0 - \mathbf{A} \cdot (\mathbf{s} - \mathbf{s}^0) - \mathbf{e} + \mathbf{e}^0 \\ &= \mathbf{b} - \mathbf{A} \cdot \mathbf{s} - \mathbf{e} = \mathbf{0} \pmod{2^\lambda}, \end{aligned}$$

where note that $2 \cdot \mathbf{e}' = \mathbf{e} - \mathbf{e}^0$. □

Below, we show that $(\epsilon^{\lambda+1})$ -hard computational LPN over \mathbb{Z}_{2^λ} implies $O(\epsilon)$ -hard LPN over \mathbb{F}_2 , where $\lambda = O(1)$ needs to be small in general (for polynomial hardness), and it can be up to $\lambda = k^{\Theta(1)}$ (for subexponential hardness). E.g., if the former is $2^{\sqrt{k}}$ -hard with $\lambda = k^{0.25}$, then the latter is $2^{\Theta(k^{0.25})}$ -hard.

Theorem 6. *If computational $(D_1, \mathbb{Z}_{2^\lambda})$ -LPN(k, N, w) is $(\lambda T + \text{poly}(k, N), \epsilon^{\lambda+1})$ -hard, then computational (D_2, \mathbb{F}_2) -LPN(k, N, w) is $(T, 2\epsilon)$ -hard, where $(D_1, D_2, w) \in \{(\text{Ber}, \text{Ber}, \mu), (\text{IndBer}, \text{Ber}, \mu), (\text{IndHW}, \text{HW}, t)\}$.*

Proof. For contradiction assume that there exists an algorithm $\mathcal{A}_{\text{LPN}_2}$ that recovers the secret of LPN over \mathbb{F}_2 with probability more than 2ϵ within time T . It suffices to prove the case $(D_1, D_2, w) = (\text{Ber}, \text{Ber}, \mu)$. For proof convenience⁸, consider distribution $\text{Ber}_{\mu, N}(\mathbb{F}_2)$ being sampled in two steps: first pick each coordinate with probability μ independently (and let the rest with 0's), and second assign the picked coordinates with uniform random bits. By a simple argument there exist at least an ϵ fraction of good $(\mathbf{A}^0, \text{coin}_1(e^i))$ for which $\mathcal{A}_{\text{LPN}_2}$ recovers s^i from $(\mathbf{A}^0, \mathbf{A}^0 s^i + e^i)$ with probability at least ϵ , where $\text{coin}_1(e^i)$ denotes the step-1 random coin for sampling e^i and the probability is taken over the s^i and the step-2 coin of e^i . Therefore, $\mathcal{A}_{\text{LPN}_{2^\lambda}}$ (see Algorithm 1) invokes \mathcal{A} on $(\mathbf{A}, \mathbf{b} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e} \pmod{2^\lambda})$ for λ times and recovers \mathbf{s} with an overall probability of at least $\epsilon^{\lambda+1}$, a contradiction to the assumption. The proofs for the other two cases are slightly simpler because $e^0, \dots, e^{\lambda-1}$ are independent and thus no two-step sampling is needed, i.e., $\text{coin}_1(e^i)$ is empty. \square

In Appendix A, we further prove the following theorem.

Theorem 7. *If computational $(\text{HW}, \mathbb{Z}_{2^\lambda})$ -LPN(k, N, t) is $(\lambda T + \text{poly}(k, N), \epsilon^{\lambda+1})$ -hard, then computational $(\text{HW}, \mathbb{F}_2)$ -LPN(k, N, t') is $(T, \frac{2\epsilon}{1 - \exp(-\delta^2 t/6)})$ -hard, where $t' = t \frac{2^{\lambda-1}}{2^\lambda - 1} (1 + \delta)$ for any constant $\delta > 0$.*

The above reduction suffers a significant security loss by exponent factor $1/(\lambda+1)$ since computationally intractable problems typically require a small success probability for efficient adversaries. In the setting of practical key recovery attacks, however, we often expect the success probability to be $(1 - 1/\text{poly}(k))$ or even overwhelming. In this case, we get more efficient reductions as below.

Theorem 8. *If the computational (D_1, \mathbb{F}_2) -LPN(k, N, w) problem can be broken (say by $\mathcal{A}_{\text{LPN}_2}$) in time T with success probability at least $(1 - \epsilon)$, then the computational $(D_2, \mathbb{Z}_{2^\lambda})$ -LPN(k, N, w) problem can be broken by $\mathcal{A}_{\text{LPN}_{2^\lambda}}$ (see Algorithm 1) in time $\lambda T + \text{poly}(k, N)$ with success probability at least $1 - (\lambda + 1)\sqrt{\epsilon}$, where $(D_1, D_2, w) \in \{(\text{Ber}, \text{Ber}, \mu), (\text{Ber}, \text{IndBer}, \mu), (\text{HW}, \text{IndHW}, t)\}$.*

Proof. Similar to the proof of Theorem 6, we have by a Markov inequality that for at least a $(1 - \sqrt{\epsilon})$ fraction of $(\mathbf{A}^0, \text{coin}_1(e^i))$ $\mathcal{A}_{\text{LPN}_2}$ recovers s^i from $(\mathbf{A}^0, \mathbf{A}^0 s^i + e^i)$ with probability at least $1 - \sqrt{\epsilon}$. Overall, $\mathcal{A}_{\text{LPN}_{2^\lambda}}$ succeeds with probability $(1 - \sqrt{\epsilon})(1 - \lambda\sqrt{\epsilon}) \geq 1 - (\lambda + 1)\sqrt{\epsilon}$ by a union bound. \square

In Appendix A, we also give a proof of the following theorem.

Theorem 9. *If the computational $(\text{HW}, \mathbb{F}_2)$ -LPN(k, N, t') problem can be broken (say by $\mathcal{A}_{\text{LPN}_2}$) in time T with success probability at least $(1 - \epsilon/2)$, then the computational $(\text{HW}, \mathbb{Z}_{2^\lambda})$ -LPN(k, N, t) problem can be broken by $\mathcal{A}_{\text{LPN}_{2^\lambda}}$ (see Algorithm 1) in time $\lambda T + \text{poly}(k, N)$ with success probability at least $1 - (\lambda + 1)\sqrt{\epsilon}$, where $t' = t \frac{2^{\lambda-1}}{2^\lambda - 1} (1 + \delta)$ for any δ and ϵ satisfying $\delta^2 t \geq 6 \ln(2/\epsilon)$.*

Optimized attacks on $(\text{Ber}/\text{HW}, \mathbb{Z}_{2^\lambda})$ -LPN. In practice, we optimize the attacks on $(\text{Ber}/\text{HW}, \mathbb{Z}_{2^\lambda})$ -LPN by exploiting the correlations among the error vectors of the λ instances, $e^0, \dots, e^{\lambda-1}$. In particular, Algorithm 1 recovers the corresponding secrets $s^0, s^1, \dots, s^{\lambda-1}$ sequentially. That means when the attacker works on the $(i + 1)$ -th LPN instance he already sees e^0, \dots, e^{i-1} from the previous i broken instances. As analyzed in the proof of Theorem 4, for any single noise sample $(e^0[j], e^1[j], \dots, e^{\lambda-1}[j]) \leftarrow \text{Ber}_{\mu, N}(\mathbb{Z}_{2^\lambda})$,

⁸The two-step sampling is defined to be in line with $\text{Ber}_{\mu, N}(\mathbb{Z}_{2^\lambda})$, and therefore captures the correlations among $e^0, \dots, e^{\lambda-1}$, which share the same step-1 randomness.

(HW, \mathbb{F})-LPN Parameters			This work ($\log \mathbb{F} = 128$)				This work ($\log \mathbb{F} = 1$)				[BCGI18] (Any field size)		
N	k	t	Gauss	SD	SD 2.0	ISD	Gauss	SD	SD 2.0	ISD	Gauss	SD	ISD
2^{10}	652	57	111	184	184	111	111	194	116	94	80	93	115
2^{12}	1589	98	100	151	151	100	100	154	137	83	85	80	104
2^{14}	3482	198	101	149	149	101	101	150	144	86	94	80	108
2^{16}	7391	389	103	147	147	103	103	148	146	91	99	80	112
2^{18}	15336	760	105	146	146	105	105	146	146	95	103	80	117
2^{20}	32771	1419	107	145	145	107	107	145	145	99	106	80	121
2^{22}	67440	2735	108	144	144	108	108	144	144	104	108	80	126

Table 3: The comparison of our analysis and [BCGI18] about the costs of Pooled Gauss, SD attack, our adapted SD 2.0 (SD 2.0) and ISD attacks to solve an LPN problem with dimension k , number of samples N , Hamming weight of noises t .

(HW, \mathbb{F})-dual-LPN Parameters			This work ($\log \mathbb{F} = 128$)				This work ($\log \mathbb{F} = 1$)				[BCGI18] (Any field size)		
n	N	t	Gauss	SD	SD 2.0	ISD	Gauss	SD	SD 2.0	ISD	Gauss	SD	ISD
2^{10}	2^{12}	44	117	189	189	117	117	191	170	97	80	100	117
2^{12}	2^{14}	39	111	170	169	111	111	170	166	95	80	92	112
2^{14}	2^{16}	34	107	151	151	107	107	151	151	93	80	84	107
2^{16}	2^{18}	32	108	145	145	108	108	145	145	95	84	82	109
2^{18}	2^{20}	31	112	143	143	112	112	143	143	99	88	82	112
2^{20}	2^{22}	30	116	141	141	116	116	141	141	103	93	82	116
2^{22}	2^{24}	29	119	139	139	119	119	139	139	107	97	82	120

Table 4: The comparison of our analysis and [BCGI18] about the costs of Pooled Gauss, SD attack, our adapted SD 2.0 (SD 2.0) and ISD attacks to solve a dual-LPN problem with $n = N/4$ (related to the number of COT/VOLE correlations), number of samples N , Hamming weight of noises t .

$e^i[j]$ is uniformly random conditioned on any non-zero $(e^0[j], \dots, e^{i-1}[j])$, and thus sample $b^i[j]$ is useless (encrypted by one-time pad) and should be discarded. In other words, the effective noise rate of the i -th LPN instance is roughly $\mu 2^{-(i+1)}$ given the attacker’s knowledge about e^0, \dots, e^{i-1} . Therefore, the success rate of solving the $(\text{Ber}, \mathbb{Z}_{2^\lambda})\text{-LPN}(k, N, \mu)$ is roughly the product of the λ instances of $(\text{Ber}, \mathbb{F}_2)\text{-LPN}$ with continuously halving noise rates $\mu, \mu/2, \dots, \mu/2^{\lambda-1}$. For instance, if solving these instances can succeed with probability $\epsilon, \epsilon^{2^{-1}}, \dots, \epsilon^{2^{-(\lambda-1)}}$ respectively, then it leads to a success probability of approximately ϵ^2 (instead of $\epsilon^{\lambda+1}$). The optimization for reducing $(\text{HW}, \mathbb{Z}_{2^\lambda})\text{-LPN}(k, N, t)$ to $(\text{HW}, \mathbb{F}_2)\text{-LPN}$ is likewise.

5 Concrete Analysis of Low-Noise LPN over Finite Fields

Recently, a series of works [BCGI18, BCG⁺19a, SGRR19, BCG⁺19b, YWL⁺20, WYKW21, CRR21] use the (dual-)LPN problem with very low noise rate over finite fields to construct concretely efficient PCG-like protocols, which extend a small number of correlations (e.g., COT, ROT, VOLE and OLE) to a large number of correlations with sublinear communication. These protocols can be used as building blocks to design a variety of MPC and ZK protocols. Therefore, the hardness of (dual-)LPN problems is crucial to guarantee the security of all the protocols.

Almost all of the known PCG-like protocols based on (dual-)LPN over a field adopt the formulas by Boyle et al. [BCGI18] to select the concrete parameters for some specified security level. Boyle et al. [BCGI18] obtained the formulas by analyzing three attacks: Pooled Gauss attack, ISD attack and SD attack. However, their analysis is oversimplified and the formulas are not accurate, which express in the following aspects:

- When analyzing the hardness of LPN with exact noise distribution $\text{HW}_{t,N}(\mathbb{F})$, the formula against Pooled Gauss attack is obtained by viewing $\text{HW}_{t,N}(\mathbb{F})$ as a Bernoulli distribution $\text{Ber}_{t/N,N}(\mathbb{F})$, which makes the formula not accurate.
- When analyzing the hardness of LPN against ISD attacks, the formula is obtained by an upper bound of the complexity of the Prange’s ISD algorithm [Pra62] to solve LPN problems over a large field. This does not cover the advanced ISD variants [Ste88, Dum91, MMT11, BJMM12]. Additionally, their analysis does not capture the impact of field sizes when calculating the ISD cost.
- When analyzing the hardness of LPN against SD attacks, each parity-check vector is assumed to be independently in compliance with a Bernoulli distribution, which is inaccurate [DT17]. In addition, their time estimate achieves a small successful probability, which does not match with that by the above two attacks.

In the following subsections, we will give more detailed analysis for the hardness of low-noise (dual-)LPN problems as well as more accurate formulas. When all known PCG-like protocols adopt exact or regular noise distributions, we focus on the case of exact noise distributions. Our formulas also cover the case of regular noise distributions, i.e., the attack costs are the same for two kinds of noise distributions. The only exception is that for the LPN problems with regular noises and small code rates (i.e., small k/N), the new algebraic attack in the very recent work [BØ23] can exploit the regular structure, and achieves lower cost than our formulas (see Section 5.4 for details).

According to our analysis, we show more accurate costs of Pooled Gauss, ISD, SD and SD 2.0 attacks in Table 3 and Table 4. In these tables, we also compare the attack costs with that given by Boyle et al. [BCGI18] and all the LPN parameters are adopted from [BCGI18]. Under the same LPN parameters, while Boyle et al. [BCGI18] showed that either Pooled Gauss attack or SD attack has the lowest cost, our more detailed analysis shows that ISD attack has the lowest cost. In other words, for PCG-like protocols, one can only consider the ISD attack to choose a set of LPN parameters in many cases. For large fields, our analysis shows that Pooled Gauss attack has the same cost as ISD attack, where recall that Pooled Gauss can be considered a special case of ISD [EKM17]. The formulas given by Boyle et al. [BCGI18] show that SD attack performs the best in most of LPN parameters. However, we point out that this is generally not true. In particular, our analysis results in Tables 3 and 4 show that SD attack has the highest cost among three kinds of attacks. Furthermore, we show in Section 5.3 that the SD 2.0 attack adapted to the low-noise setting (which recovers the optimal SD attack) require more cost than the ISD attack against LPN problems over large fields using random linear codes. Our proof does *not* require that the Hamming weight of noise vectors achieves the Gilbert-Varshamov bound, and instead allow all weights in the LPN parameters used in PCG-like protocols. Tables 3 and 4 also show that ISD attack has lower cost for smaller field size, which justifies that it is not accurate to use the same formulas for all field sizes as in [BCGI18].

Our analysis shows that most of LPN parameter sets used in know PCG-like protocols actually achieve better security. In other words, given a security parameter, smaller weight for noise vectors can be chosen, which allows for achieving higher efficiency in these PCG-like protocols. In particular, for 128-bit security level, Table 5 reports the comparison of Hamming weights of noise vectors between our analysis and previous work [BCGI18], given several sets of dimensions and the numbers of samples used in [BCGI18]. From this table, we can see that the Hamming weight of a noise vector is reduced by 19% – 36% for LPN and 41% – 43% for dual-LPN, respectively. While the COT/VOLE protocols in the PCG framework need to

(HW, \mathbb{F})-LPN(k, N, t)				(HW, \mathbb{F})-dual-LPN(N, n, t)			
N	k	t ([BCGI18])	t (ours)	n	N	t ([BCGI18])	t (ours)
2^{10}	652	106	68	2^{10}	2^{12}	88	50
2^{12}	1589	172	136	2^{12}	2^{14}	83	48
2^{14}	3482	338	274	2^{14}	2^{16}	78	45
2^{16}	7391	667	531	2^{16}	2^{18}	73	42
2^{18}	15336	1312	1023	2^{18}	2^{20}	68	39
2^{20}	32771	2467	1876	2^{20}	2^{22}	63	36
2^{22}	67440	4788	3552	2^{22}	2^{24}	58	34

Table 5: **The comparison of noise weights between our analysis and [BCGI18] for 128-bit security level.** Among the LPN parameters, N is the number of samples, k is the dimension, $n = N/4$ and t is the Hamming weight of a noise vector. We assume the field size is $\log |\mathbb{F}| = 128$, and an exact noise distribution $\text{HW}_{t,N}(\mathbb{F})$.

communicate $O(t\kappa \log \frac{N}{t})$ bits and compute t GGM trees for each iteration, the reduction of parameter t will be directly transferred to the improvements of communication and computation costs in these protocols.

5.1 The Hardness of LPN against Pooled Gauss Attack

For solving the problems of LPN and dual-LPN with exact noise distributions, Boyle et al. [BCGI18] gave a formula to compute the cost of Pooled Gauss attack by simplifying the underlying noise distribution $\text{HW}_{t,N}(\mathbb{F})$ to $\text{Ber}_{t/N,N}(\mathbb{F})$. Below, we discuss why the concrete security under $\text{Ber}_{t/N,N}(\mathbb{F})$ -noise is significantly lower than that under $\text{HW}_{t,N}(\mathbb{F})$ -noise. Moreover, we give a more accurate formula to estimate the cost under Pooled Gauss attack, and discuss the difference between the over-simplified analysis [BCGI18] and our analysis.

We may have the false impression that the concrete security of LPN under $\text{HW}_{t,N}(\mathbb{F})$ -noise is roughly the same as that under $\text{Ber}_{t/N,N}(\mathbb{F})$ -noise since for $\mathbf{e} \leftarrow \text{Ber}_{t/N,N}(\mathbb{F})$ the noise \mathbf{e} follows exactly the $\text{HW}_{t,N}(\mathbb{F})$ distribution conditioned on $|\mathbf{e}| = t$, which occurs with noticeable probability (see Lemma 1). However, this reduction gives only an upper bound on the security of $\text{Ber}_{t/N,N}(\mathbb{F})$ -LPN, which is not tight in the concrete setting. To see this, take the parameter ($n = 2^{10}$, $N = 2^{12}$, exact weight $t = 44$) for dual LPN considered in [17] as an example. We denote by $T(\text{Ber}, \mu, \text{SUC})$ and $T(\text{HW}, t, \text{SUC})$ the bit-security of LPN with Bernoulli noise of rate μ and noise of exact-weight- t respectively, with success probability $\geq \text{SUC}$. The above reduction translates to $T(\text{Ber}, t/N, 0.06) \leq T(\text{HW}, t = 44, 0.99)$, where $\Pr[|\mathbf{e}| = 44] \approx 6\%$ for $\mathbf{e} \leftarrow \text{Ber}_{\mu,N}$ and $\mu = t/N = 44/2^{12}$. For a tighter bound on $T(\text{Ber}, t/N, \text{SUC})$, consider collectively the weights up to a smaller threshold, e.g, $\Pr[|\mathbf{e}| \leq 35] \approx 10\%$, $\Pr[|\mathbf{e}| \leq 33] \approx 5\%$, and thus

$$T(\text{Ber}, t/N, 0.1) \leq T(\text{HW}, 35, 0.99), \quad T(\text{Ber}, t/N, 0.05) \leq T(\text{HW}, 33, 0.99) ,$$

where the rationale is any algorithm for LPN of exact-weight t implies (and in practice is) one for LPN of exact weight $t' < t$ with the same cost. Therefore, when evaluating the security for LPN with exact-weight $t = 44$, i.e., $T(\text{HW}, 44, 0.99)$, [17] simply used $\mathbf{e} \leftarrow \text{Ber}_{t/N,N}$ instead of $\mathbf{e} \leftarrow \text{HW}_{t,N}$, and hence what they obtained is more close to $T(\text{HW}, 35, 0.99)$ or even $T(\text{HW}, 33, 0.99)$. For reference, according to our security estimator the bit security $T(\text{HW}, w, 0.99)$ for $w = 44$, $w = 35$, and $w = 33$ is 117, 98, and 94 respectively. Hopefully, this explains the concrete security gap between using Bernoulli and exact-weight distributions from the reductionist perspective.

We first extend Pooled Gauss attack [EKM17] from solving a (dual-)LPN problem from a Bernoulli distribution to solving that with an exact noise distribution $\text{HW}_{t,N}(\mathbb{F})$. Specifically, the Pooled Gauss attack performs as follows:

- Given a (HW, \mathbb{F}) -LPN(k, N, t) instance $\mathbf{b} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e}$, for each iteration, guess k non-noisy coordinates of vector \mathbf{b} by sampling them at random, and obtain a length- k vector \mathbf{b}' and $k \times k$ matrix \mathbf{A}' . Then, compute $\mathbf{s}' := (\mathbf{A}')^{-1} \cdot \mathbf{b}'$ and verify whether \mathbf{s}' is correct or not by running the test algorithm in [EKM17].
- Given a (HW, \mathbb{F}) -dual-LPN(N, n, t) instance $\mathbf{b} = \mathbf{H} \cdot \mathbf{e}$, for each iteration, guess n coordinates of vector \mathbf{e} by sampling them uniformly at random such that these coordinates contain t noisy coordinates of \mathbf{e} , choose the corresponding $n \times n$ sub-matrix \mathbf{H}' from \mathbf{H} according to the n coordinates, compute $\mathbf{e}' := (\mathbf{H}')^{-1} \cdot \mathbf{b}$, and then checks if $|\mathbf{e}'| = t$.

The above attack uses the fixed N samples for both LPN and dual-LPN. For solving LPN, Pooled Gauss attack runs in time $\frac{\binom{N}{t}}{\binom{N-k}{t}} \cdot k^{2.8}$, where $\frac{\binom{N-k}{t}}{\binom{N}{t}}$ is the probability of guessing successfully in one iteration, and $k^{2.8}$ is the cost of inverting matrix \mathbf{A}' via Strassen's algorithm. For solving dual-LPN, Pooled Gauss attack runs in time $\frac{\binom{N}{t}}{\binom{N-k}{t}} \cdot (N-k)^{2.8}$, where $k = N - n$ is the dimension and $(N-k)^{2.8}$ is the cost of inverting matrix \mathbf{H}' via Strassen's algorithm. As LPN can be efficiently transformed into dual-LPN and vice versa, Pooled Gauss attack solves a (dual-)LPN problem with number of samples N , dimension k and weight of noises t in time $\frac{\binom{N}{t}}{\binom{N-k}{t}} \cdot \min(k^{2.8}, (N-k)^{2.8})$. Therefore, the bit-security of a (dual-)LPN instance with respect to Pooled Gauss attack is computed as

$$\log(\min(k^{2.8}, (N-k)^{2.8})) + \left(\log \binom{N}{t} - \log \binom{N-k}{t} \right).$$

For the LPN problem with an exact noise distribution, [BCGI18] simplified the cost of Pooled Gauss attack as $\left(\frac{1}{1-t/N}\right)^k \cdot k^{2.8}$. In the following, we compare the attack cost in [BCGI18] and our estimate by computing their ratio:

$$\frac{T_{\text{ours}}}{T_{\text{[BCGI18]}}} \approx \left(1 - \frac{t}{N}\right)^k \cdot \frac{\binom{N}{t}}{\binom{N-k}{t}} \geq \left(1 - \frac{t}{N}\right)^k \cdot \left(\frac{N}{N-k}\right)^t \approx e^{tk \cdot \left(\frac{1}{N-k} - \frac{1}{N}\right)} = e^{\frac{t \cdot k^2}{N \cdot (N-k)}},$$

where \approx denotes an approximate relation that omits a polynomial factor. For the (dual-)LPN parameters N, k used in known PCG-like protocols (e.g., $k = 0.75N$ considered in [BCGI18]), the cost of Pooled Gauss attack estimated by Boyle et al. [BCGI18] is about $2^{O(t)}$ times larger than our estimate.

5.2 The Hardness of LPN against ISD Attack

Based on the concrete analysis of ISD attack in [HOSS18], Boyle et al. [BCGI18] used an upper bound of the cost of the Prange's ISD algorithm [Pra62] to evaluate the hardness of LPN problems over any field against ISD attacks. As shown in Tables 3 and 4, the upper-bound formula cannot capture accurately the cost of more advanced ISD variants [Ste88, Dum91, MMT11, BJMM12]. In the following, we first summarize the known ISD variants, and then use the state-of-the-art ISD algorithm to evaluate the hardness of LPN problems over finite fields⁹. We distinguish two cases: binary field \mathbb{F}_2 and other larger fields.

Binary field \mathbb{F}_2 . For binary field \mathbb{F}_2 , we adopt the state-of-the-art BJMM-ISD algorithm [BJMM12] to analyze the concrete hardness of low-noise (dual-)LPN problems. From the expected time of BJMM-ISD shown in Theorem 18 of Appendix D, we can see that it is hard to give a succinct formula to compute the cost of BJMM-ISD. Thus, we choose to provide a script (see Footnote 2), which allows to automatically compute

⁹ Torres and Sendrier [TS16] show that known variants of ISD have the same asymptotic complexity in the sub-linear error weight regime, and the difference between the non-asymptotic exponent of some ISD variants is relatively small (but non-zero). We pick the more advanced ISD for a more accurate security estimate.

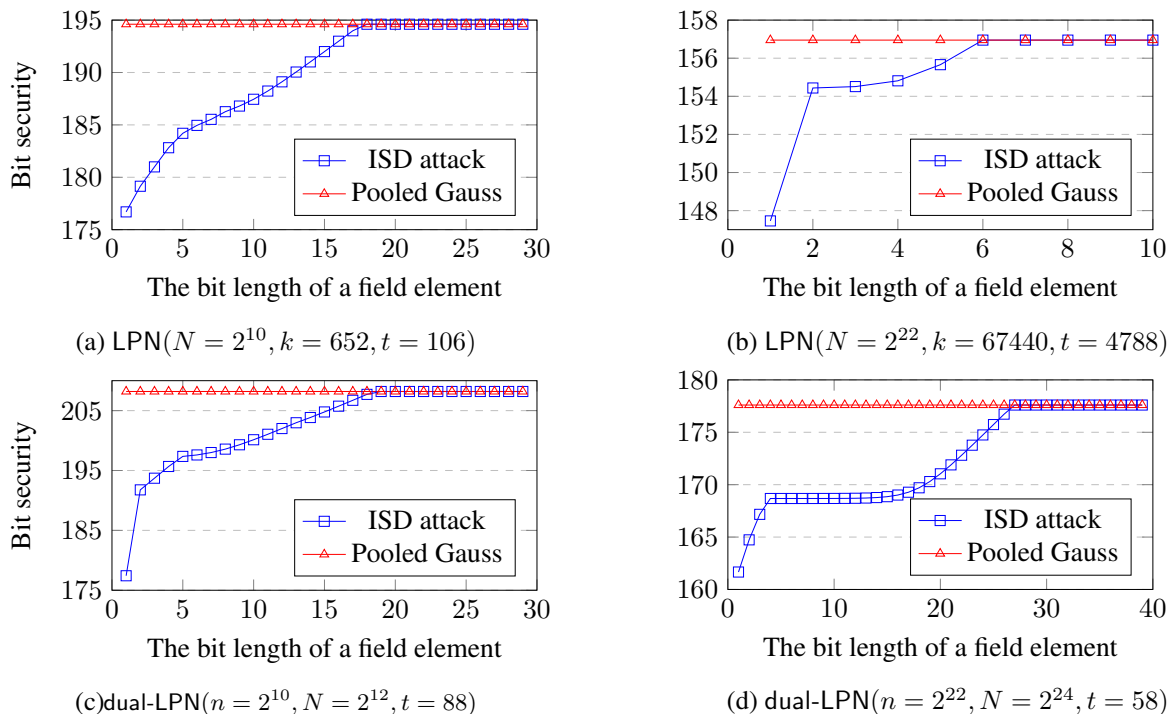


Figure 3: The costs of Pooled Gauss and SD-ISD attacks from our analysis to solve LPN and dual-LPN problems on different field sizes, where N is the number of samples, k is the dimension, $n = N/4$ and t is the weight of a noise vector.

the cost of the BJMM-ISD attack. Our analysis does not consider the known NN-ISD variants [MO15, BM18], because these ISD variants introduce quite large polynomial overheads, and make them less efficient when analyzing the concrete costs of low-noise LPN problems. Besides, our analysis does not cover the ISD algorithms [BBC⁺19, EB22, EMZ22] with more efficient space consumption, and always assume that sufficient memory is available which makes the LPN parameters more conservative.

Larger fields. For the case of non-binary fields, we focus on the hardness of (dual-)LPN problems over any finite field \mathbb{F} with $|\mathbb{F}| \geq 256$ (especially for large fields with $|\mathbb{F}| \geq 2^8$). For other field sizes (i.e., $|\mathbb{F}| \in \{4, 8, 16, 32, 64, 128\}$), low-noise (dual-)LPN seem to be less interesting for PCG, MPC and ZK applications. We adopt the generalized SD-ISD algorithm [Pet10] to analyze the cost of solving low-noise (dual-)LPN problems over field \mathbb{F} . As such, we provide a Python script to automatically compute the cost of the SD-ISD attack.

Compared to the SD-ISD algorithm by Peters [Pet10], the SD-ISD variant by Meurer [Meu12] has the performance advantage when the field size $|\mathbb{F}| < 128$, and the advantage becomes vanishingly small when $|\mathbb{F}| \geq 128$. Thus, for the case that $|\mathbb{F}| \geq 256$, it is enough to adopt the Peters’s SD-ISD algorithm to evaluate the hardness of low-noise (dual-)LPN problems. Note that it is unclear how to extend the generalization approaches [Pet10, Meu12] to more efficient MMT-ISD and BJMM-ISD algorithms for solving low-noise LPN over a larger field \mathbb{F} and make the resulting ISD algorithm be significantly more efficient than the SD-ISD variant [Pet10]. We leave it as an interesting future work. Even if these generalization approaches can be efficiently applied in the MMT-ISD and BJMM-ISD algorithms, the performance advantage of MMT-ISD and BJMM-ISD (compared to SD-ISD) decreases when the field size increases, and will diminish for a sufficiently large field. Our analysis does not consider the ISD variants [Hir16, GKH17] that generalize the May-Ozerov nearest-neighbor algorithm to solve dual-LPN problems any finite field, as these ISD variants

introduce quite large polynomial overheads and thus are less efficient when analyzing the concrete costs of solving low-noise LPN.

For any finite field \mathbb{F} , the generalized SD-ISD variant [Pet10] generates two sets \mathcal{S}_0 and \mathcal{S}_1 with size $\binom{(k+\ell)/2}{p/2} \cdot (|\mathbb{F}| - 1)^{p/2}$, where p and ℓ are two additional parameters for the SD-ISD variant. Note that the size of $|\mathcal{S}_0|$ and $|\mathcal{S}_1|$ increases exponentially with p . Following the work [Pet10], p should be quite small to minimize the cost of going through sets \mathcal{S}_0 and \mathcal{S}_1 , and ℓ is set as $\ell = \log_{|\mathbb{F}|} \binom{k/2}{p} + p \log_{|\mathbb{F}|} (|\mathbb{F}| - 1)$. If the even integer $p \neq 0$, then the cost of the generalized SD-ISD algorithm is at least $O((k + \ell)|\mathbb{F}|)$, which is the cost to just find identical elements in two sets \mathcal{S}_0 and \mathcal{S}_1 . When the field size is large enough, we will have to choose $p = 0$ and $\ell = 0$ to minimize the cost of the generalized SD-ISD attack, according to the above equation. In this case, the generalized SD-ISD attack actually becomes Pooled Gauss attack. In Figure 3, we give a comparison of costs between Pooled Gauss attack and the generalized SD-ISD attack for solving LPN and dual-LPN problems on different field sizes. From this figure, we can see that ISD attack has the same cost as Pooled Gauss attack when the field size is sufficiently large.

5.3 The Hardness of LPN against SD Attack

Boyle et al. [BCGI18] analyzed the cost of solving LPN problems against SD attack, and shown that this attack performs the best among three kinds of attacks for their parameters selection. In the following, we show two oversimplifications in their analysis [BCGI18], which lead to underestimate the cost of this attack. First, Boyle et al. [BCGI18] assumes that each parity-check vector $\mathbf{v} \in \mathcal{V}$ independently follows a Bernoulli distribution $\text{Ber}_{w/N, N}(\mathbb{F})$ where $w \in \mathbb{N}$ is the Hamming weight of \mathbf{v} , which is inaccurate¹⁰ [DT17] and leads to underestimated cost of SD attack. To obtain more accurate complexity of this attack, a weaker assumption was proposed by [DT17] where each parity-check vector $\mathbf{v} \in \mathcal{V}$ is assumed to be independently in compliance with an exact noise distribution $\text{HW}_{w, N}(\mathbb{F})$. Second, Boyle et al. [BCGI18] underestimated the cost of this attack as $T = T_1/\epsilon$, where T_1 is the time of finding one parity check vector $\mathbf{v} \in \mathcal{V}$ and ϵ is the distinguishing advantage for one vote. The SD attack solves the decisional LPN problem with negligible advantage in time T , while other attacks solve the LPN problem with constant advantage. Following the previous works [Al 01, Ove06, FKI07, DT17], this attack takes time $T = T_1/\epsilon^2$ to distinguish LPN samples from random samples with constant advantage.¹¹

The SD 2.0 algorithm [CDAMHT22] outperforms the ISD algorithms for the GV bound decoding over \mathbb{F}_2 with code rate $k/N < 0.3$. Given a $(\text{HW}, \mathbb{F})\text{-LPN}(k, N, t)$ instance $\mathbf{b} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e}$, the SD 2.0 algorithm, that generalizes from \mathbb{F}_2 to any finite field \mathbb{F} , introduces a new parameter s , and proceeds in two stages to recover \mathbf{e} .

1. Find a set of parity-check vectors

$$\mathcal{V} \subset \left\{ \mathbf{v} \stackrel{\text{def}}{=} [\mathbf{v}_1 \quad \mathbf{v}_2] \mid \mathbf{v}_1 \in \mathbb{F}^s, \mathbf{v} \cdot \mathbf{A} = \mathbf{0} \text{ and } |\mathbf{v}_2| = w \right\},$$

with a sufficiently small $w > 0$. Note that for each \mathbf{v} , $\langle \mathbf{v}, \mathbf{b} \rangle = \langle \mathbf{v}, \mathbf{A} \cdot \mathbf{s} + \mathbf{e} \rangle = \langle \mathbf{v}, \mathbf{e} \rangle = \langle \mathbf{v}_1, \mathbf{e}_1 \rangle + \langle \mathbf{v}_2, \mathbf{e}_2 \rangle$, where $\mathbf{e} \stackrel{\text{def}}{=} [\mathbf{e}_1 \quad \mathbf{e}_2]$ with $\mathbf{e}_2 \in \mathbb{F}^{N-s}$. The SD 2.0 algorithm [CDAMHT22] assumes that $\langle \mathbf{v}_2, \mathbf{e}_2 \rangle$ independently follows a Bernoulli distribution $\text{Ber}_{\epsilon, 1}(\mathbb{F})$ with some $\epsilon > 0$ for each $\mathbf{v} \in \mathcal{V}$.

¹⁰The Bernoulli distribution admits a slackness event that the weight of a parity-check vector \mathbf{v} goes much below the expected w leading to underestimated attack costs. However, for an optimal weight w , such low-weight vectors \mathbf{v} 's violating the Gilbert–Varshamov bound may not exist at all.

¹¹We shall distinguish the differences between $1/\epsilon$ and $1/\epsilon^2$. If a single *key-recovery* attack succeeds with probability ϵ , then repeating roughly $1/\epsilon$ independent instances achieves constant (or even overwhelming) success probability. In contrast, if a single *distinguishing attack* gains advantage ϵ , then the number of independent votes needed to amplify the advantage to constant is about $1/\epsilon^2$ by a Chernoff bound.

2. Recover $e_1 \in \mathbb{F}^s$ from

$$\{(\mathbf{v}_1, \langle \mathbf{v}_1, \mathbf{e}_1 \rangle + \Delta) \mid [\mathbf{v}_1 \ \mathbf{v}_2] \in \mathcal{V}, \mathbf{v}_1 \in \mathbb{F}^s \text{ and } \Delta \leftarrow \text{Ber}_{\epsilon,1}(\mathbb{F})\}$$

via a fast Fourier transform in time $O(s \log(|\mathbb{F}|)|\mathbb{F}|^s)$. The procedure to recover e_2 is similar.

An appropriate value for parameter s is chosen so that the two stages take almost the same cost. The SD 2.0 algorithm [CDAMHT22] achieves better advantage than that achieved by the traditional SD attack.

For the GV bound decoding over \mathbb{F}_2 , SD 2.0 [CDAMHT22] uses the same collision technique invoked by the advanced ISD algorithms [Ste88, Dum91, MMT11, BJMM12] to find a set \mathcal{V} with a much smaller w . However, SD 2.0 [CDAMHT22] does not perform better on low-noise LPN problems, as the collision technique takes exponential time $|\mathbb{F}|^{\theta(k)}$ that is already larger than the subexponential time $2^{o(k)}$ of ISD to solve the low-noise LPN problem. We adapt the SD 2.0 algorithm [CDAMHT22] to solve low-noise LPN problems by replacing the collision technique with other collision techniques (e.g., [BCGI18]). The traditional SD attack is a special case of the adapted SD 2.0 attack by using the collision technique that sets s as a constant. Below, we analyze the SD 2.0 cost of solving low-noise LPN by considering two cases: binary field \mathbb{F}_2 and larger fields.

Binary field \mathbb{F}_2 . Given a (HW, \mathbb{F}) -LPN(k, N, t) instance, the adapted SD 2.0 algorithm finds a parity-check vector set with Hamming weight $w = w(s)$. According to [Lyu05, Lemma 3], we have that

$$\epsilon_s = \Pr[\langle \mathbf{v}_2, \mathbf{e}_2 \rangle = 0] - 1/|\mathbb{F}| = \frac{1}{2} \left(\frac{N - 2w - t + 1}{N - t + 1} \right)^t,$$

and the adapted algorithm solves the problem in time $T = \min_s (T_1 \cdot (1/\epsilon_s)^2 + s \cdot 2^s)$, where $w = (k+1-s)/2$ and $T_1 = k + 1$ following the work [BCGI18].

Larger fields. For fields of size $|\mathbb{F}| \geq 4t$, we adapt the SD 2.0 algorithm [CDAMHT22] to analyze the cost of solving the (HW, \mathbb{F}) -LPN(k, N, t) problem. In particular, we have the following theorem.

Theorem 10. *For $w = w(s) \in \mathbb{N}$ and a finite field \mathbb{F} with $|\mathbb{F}| \geq 4t$, the adapted SD 2.0 algorithm solves the decisional (HW, \mathbb{F}) -LPN(k, N, t) problem in time*

$$T = \min_s \left(T_1 \cdot \left(\frac{\binom{N}{t}}{\binom{N-w}{t}} \cdot \frac{2|\mathbb{F}|}{|\mathbb{F}| - 1} \right)^2 + s \cdot \log(|\mathbb{F}|) \cdot |\mathbb{F}|^s \right),$$

with constant advantage, where T_1 is the time of finding one parity-check vector.

The proof of Theorem 10 is postponed to Appendix E. We set $w = k - s + 1$ and assume $T_1 = k + 1$ following the work [BCGI18]. For the case of a field \mathbb{F} with $2 < |\mathbb{F}| < 4t$, we can still use above formula to estimate the cost of the adapted SD 2.0 algorithm, which is smaller than the actual cost and makes our analysis more conservative.

Below, we show that for solving LPN problems over large fields, both the optimal SD attack and the SD 2.0 attack (adapted to the low-noise setting) take more cost than the Prange's ISD algorithm [Pra62]. In the following theorem, we assume that LPN problems adopt random linear codes as in [DT17].

Theorem 11. *For any (HW, \mathbb{F}) -LPN(k, N, t) problem with $|\mathbb{F}| = k^{\omega(1)}$, $(1 + \beta)k \leq N = \text{poly}(k)$ for a constant $\beta > 0$ and $t = o(N)$, both the traditional SD attack and our adapted SD 2.0 attack (that distinguishes with constant advantage) require more cost than the Prange's ISD algorithm (that recovers the secret with constant probability).*

Proof. Since the traditional SD attack has more cost than the adapted SD 2.0 attack, it suffices to prove that SD 2.0 attack (adapted to the low-noise setting) is less efficient than Prange’s ISD algorithm. For $s = \Omega(k)$ the statement is trivially true since the SD 2.0 algorithm takes at least $|\mathbb{F}|^s$ to solve the (HW, \mathbb{F}) -LPN(k, N, t) problem, which is much greater than the time complexity $2^{o(k)}$ of the Prange’s ISD algorithm. Hence, we only need to consider $s = o(k)$ in the SD 2.0 framework.

The SD 2.0 framework needs to find parity-check vectors from

$$\mathcal{V} \subset \left\{ \mathbf{v} \stackrel{\text{def}}{=} [\mathbf{v}_1 \quad \mathbf{v}_2] \mid \mathbf{v}_1 \in \mathbb{F}^s, \mathbf{v} \cdot \mathbf{A} = \mathbf{0} \text{ and } |\mathbf{v}_2| = w \right\},$$

with a sufficiently small $w > 0$. The expected number of parameter- w parity-check vectors is $\mathbb{E}[|\mathcal{V}|] \leq \binom{N-s}{w} |\mathbb{F}|^{w-k+s}$. We need $\mathbb{E}[|\mathcal{V}|] \geq 1$ to guarantee existence. In particular, we have the following:

$$2^{w \cdot \log N + (w-k+s) \cdot \log |\mathbb{F}|} = N^w \cdot |\mathbb{F}|^{w-k+s} \geq \binom{N-s}{w} \cdot |\mathbb{F}|^{w-k+s} \geq 1.$$

From the above relation, $|\mathbb{F}| = k^{\omega(1)}$ and $N = \text{poly}(k)$, we have that $w \geq (1 - o(1))k$. Since the cost of SD 2.0 increases as w increases, we set $w = (1 - o(1))k$ for an optimal attack in the SD 2.0 framework and we conservatively assume that the time of finding a parity-check vector with minimal weight w is at least the time using the Gaussian elimination method (denoted by T_{Gauss}).¹² According to Theorem 10, we have that the SD 2.0 solves the LPN problem with constant probability in time

$$T = \min_s \left(T_{Gauss} \cdot \left(\frac{\binom{N}{t}}{\binom{N-w}{t}} \cdot \frac{2|\mathbb{F}|}{|\mathbb{F}| - 1} \right)^2 + s \log (|\mathbb{F}|) |\mathbb{F}|^s \right).$$

The cost of the Prange’s ISD algorithm is upper bounded by $T_{Gauss} \cdot \frac{\binom{N}{t}}{\binom{N-k}{t}}$ [BCGI18, HOSS18]. The cost-ratio between the optimal attack in the SD 2.0 framework and the Prange’s ISD algorithm is

$$\frac{\left(\frac{\binom{N}{t}}{\binom{N-w}{t}} \right)^2}{\frac{\binom{N}{t}}{\binom{N-k}{t}}} = \frac{\binom{N}{t} \cdot \binom{N-k}{t}}{\binom{N-w}{t}^2} = \prod_{i=1}^t \frac{(N-t+i) \cdot (N-k-t+i)}{(N-w-t+i)^2}. \quad (2)$$

For $N \geq (1 + \beta)k$, $t = o(N)$ and $w = (1 - o(1))k$, we have that $w \cdot (2N - 2t - w) \geq k(N - t)$. It follows that the above value in (2) is at least 1. \square

5.4 The Hardness of LPN against Algebraic Attack

Very recently, Briaud and Øygarden [BØ23] proposed a new algebraic attack that is able to take advantage of regular noise distributions. In particular, their attack solves a polynomial system in the coordinates of a regular noise vector e by utilizing the quadratic system that encodes the regular structure but does not depend on the particular LPN instance.

Compared to linear attacks including Pooled Gauss, ISD and SD attacks, their algebraic attack has lower cost for solving LPN problems with regular noises and small code rates. Since the dual-LPN problems used in PCG-like protocols have constant code rate (often 1/2 or 3/4), their algebraic attack does not perform better than ISD attacks. When the LPN problems under the parameter sets given in [BCGI18] have small

¹²If there exist more efficient algorithms to find such a parity-check vector, then they can also be used to improve the ISD algorithms.

code rates, Briaud and Øyngarden [BØ23] shown that their algebraic attack outperforms ISD attacks for the following parameter sets:

$$\begin{aligned} (N, k, t) &= (2^{10}, 652, 106), & (N, k, t) &= (2^{18}, 15336, 1312), \\ (N, k, t) &= (2^{20}, 32771, 2467), & (N, k, t) &= (2^{22}, 64770, 4788), \\ (N, k, t) &= (2^{12}, 1589, 172) \text{ for finite field } \mathbb{F}_{2^{128}}. \end{aligned}$$

Compared to ISD attacks, the increase in the cost of their algebraic attack when going from \mathbb{F}_2 to $\mathbb{F}_{2^{128}}$ is smaller for most parameter sets. Compared to the cost of LPN under the above parameter sets using our script, their algebraic attack allows one to reduce the cost by $3 \sim 45$ bits. For the LPN problems with larger code rates used in the PCG-like protocols [YWL⁺20, WYKW21], the new algebraic attack [BØ23] does not perform better than ISD attacks. In the one-time setup phase, these protocols adopt the LPN problems with small code rates, where the algebraic attack will perform best. Particularly, for $(N, k, t) = (642048, 19870, 2508)$, the algebraic attack reduces the cost by 28 bits compared to ISD.

Given the number of samples, we can increase the dimension k to add more cost into the algebraic attack [BØ23]. In the PCG-like protocols, the overhead of increasing k is negligibly small for the extension phase as the technique of bootstrapping iterations is used. But larger dimension k requires more cost for the one-time setup phase. In the future, we will integrate the algebraic attack into our script, and make the script be publicly available.

Acknowledgements

Work of Xiao Wang is supported in part by DARPA under Contract No. HR001120C0087, NSF award #2016240, research awards from Facebook and Google, and a Gift from PlatON. The views, opinions, and/or findings expressed are those of the author(s) and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government. Work of Kang Yang is supported by the National Natural Science Foundation of China (Grant Nos. 62102037, 61932019, 62022018). Work of Yu Yu is supported by the National Key Research and Development Program of China (Grant Nos. 2020YFA0309705 and 2018YFA0704701) and the National Natural Science Foundation of China (Grant Nos. 62125204 and 61872236). Yu Yu also acknowledges the support from the XPLOER PRIZE.

References

- [AIK07] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. Cryptography with constant input locality. In Alfred Menezes, editor, *CRYPTO 2007*, volume 4622 of *LNCS*, pages 92–110. Springer, Heidelberg, August 2007.
- [Aka08] Adi Akavia. *Learning Noisy Characters, Multiplication Codes, and Cryptographic Hardcore Predicates*. PhD thesis, Massachusetts Institute of Technology, 2008.
- [Al 01] A. Kh. Al Jabri. A statistical decoding algorithm for general linear block codes. In Bahram Honary, editor, *8th IMA International Conference on Cryptography and Coding*, volume 2260 of *LNCS*, pages 1–8. Springer, Heidelberg, December 2001.
- [Ale03] Michael Alekhnovich. More on average case vs approximation complexity. In *44th FOCS*, pages 298–307. IEEE Computer Society Press, October 2003.
- [AMBD⁺18] Carlos Aguilar-Melchor, Olivier Blazy, Jean-Christophe Deneuville, Philippe Gaborit, and Gilles Zémor. Efficient encryption from random quasi-cyclic codes. *IEEE Transactions on Information Theory*, 64, 2018.

- [App16] Benny Applebaum. Garbling XOR gates “for free” in the standard model. *Journal of Cryptology*, 29(3):552–576, July 2016.
- [BA21] Gregory Bard and Martin Albrecht. M4ri(e)- linear algebra over \mathbf{F}_2 (and \mathbf{F}_{2^e}). In *Free Open Source Software*, 2021.
- [BBC⁺19] Marco Baldi, Alessandro Barenghi, Franco Chiaraluce, Gerardo Pelosi, and Paolo Santini. A finite regime analysis of information set decoding algorithms. *Algorithms*, 12(10), 2019.
- [BBMH⁺21] Carsten Baum, Lennart Braun, Alexander Munch-Hansen, Benoît Razet, and Peter Scholl. Appenzeller to brie: Efficient zero-knowledge proofs for mixed-mode arithmetic and \mathbb{Z}_2^k . In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 192–211. ACM Press, November 2021.
- [BBMHS22] Carsten Baum, Lennart Braun, Alexander Munch-Hansen, and Peter Scholl. Moz \mathbb{Z}_{2^k} arella: Efficient vector-OLE and zero-knowledge proofs over \mathbb{Z}_{2^k} . In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part IV*, volume 13510 of *LNCS*, pages 329–358. Springer, Heidelberg, August 2022.
- [BC23] Dung Bui and Geoffroy Couteau. Private set intersection from pseudorandom correlation generators. In *PKC 2023, Part I*, *LNCS*. Springer, Heidelberg, May 2023.
- [BCG⁺19a] Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, Peter Rindal, and Peter Scholl. Efficient two-round OT extension and silent non-interactive secure computation. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 291–308. ACM Press, November 2019.
- [BCG⁺19b] Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, and Peter Scholl. Efficient pseudorandom correlation generators: Silent OT extension and more. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 489–518. Springer, Heidelberg, August 2019.
- [BCG⁺22] Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, Nicolas Resch, and Peter Scholl. Correlated pseudorandomness from expand-accumulate codes. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 603–633. Springer, Heidelberg, August 2022.
- [BCGI18] Elette Boyle, Geoffroy Couteau, Niv Gilboa, and Yuval Ishai. Compressing vector OLE. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018*, pages 896–912. ACM Press, October 2018.
- [BFKL94] Avrim Blum, Merrick L. Furst, Michael J. Kearns, and Richard J. Lipton. Cryptographic primitives based on hard learning problems. In Douglas R. Stinson, editor, *CRYPTO’93*, volume 773 of *LNCS*, pages 278–291. Springer, Heidelberg, August 1994.
- [BGI15] Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 337–367. Springer, Heidelberg, April 2015.
- [BJMM12] 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 David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 520–536. Springer, Heidelberg, April 2012.

- [BKW00] Avrim Blum, Adam Kalai, and Hal Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. In *32nd ACM STOC*, pages 435–440. ACM Press, May 2000.
- [BLVW19] Zvika Brakerski, Vadim Lyubashevsky, Vinod Vaikuntanathan, and Daniel Wichs. Worst-case hardness for LPN and cryptographic hashing via code smoothing. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part III*, volume 11478 of *LNCS*, pages 619–635. Springer, Heidelberg, May 2019.
- [BM18] Leif Both and Alexander May. Decoding linear codes with high error rate and its impact for LPN security. In Tanja Lange and Rainer Steinwandt, editors, *Post-Quantum Cryptography - 9th International Conference, PQCrypto 2018*, pages 25–46. Springer, Heidelberg, 2018.
- [BMRS21] Carsten Baum, Alex J. Malozemoff, Marc B. Rosen, and Peter Scholl. Mac’n’cheese: Zero-knowledge proofs for boolean and arithmetic circuits with nested disjunctions. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part IV*, volume 12828 of *LNCS*, pages 92–122, Virtual Event, August 2021. Springer, Heidelberg.
- [BØ23] Pierre Briaud and Morten Øygarden. A new algebraic approach to the regular syndrome decoding problem and implications for pcg constructions. In *EUROCRYPT 2023*, *LNCS*. Springer, Heidelberg, June 2023.
- [CDAMHT22] Kevin Carrier, Thomas Debris-Alazard, Charles Meyer-Hilfiger, and Jean-Pierre Tillich. Statistical decoding 2.0: Reducing decoding to lpn. In *ASIACRYPT 2022*. Springer, 2022.
- [CDE⁺18] Ronald Cramer, Ivan Damgård, Daniel Escudero, Peter Scholl, and Chaoping Xing. SPD \mathbb{Z}_{2^k} : Efficient MPC mod 2^k for dishonest majority. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 769–798. Springer, Heidelberg, August 2018.
- [CG90] John T. Coffey and Rodney M. Goodman. The complexity of information set decoding. *IEEE Transactions on Information Theory*, 36(5), 1990.
- [CRR21] Geoffroy Couteau, Peter Rindal, and Srinivasan Raghuraman. Silver: Silent VOLE and oblivious transfer from hardness of decoding structured LDPC codes. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part III*, volume 12827 of *LNCS*, pages 502–534, Virtual Event, August 2021. Springer, Heidelberg.
- [DDN14] Bernardo David, Rafael Dowsley, and Anderson C. A. Nascimento. Universally composable oblivious transfer based on a variant of LPN. In Dimitris Gritzalis, Aggelos Kiayias, and Ioannis G. Askoxylakis, editors, *CANS 14*, volume 8813 of *LNCS*, pages 143–158. Springer, Heidelberg, October 2014.
- [DEM19] Claire Delaplace, Andre Esser, and Alexander May. Improved low-memory subset sum and LPN algorithms via multiple collisions. In Martin Albrecht, editor, *17th IMA International Conference on Cryptography and Coding*, volume 11929 of *LNCS*, pages 178–199. Springer, Heidelberg, December 2019.
- [DI14] Erez Druk and Yuval Ishai. Linear-time encodable codes meeting the gilbert-varshamov bound and their cryptographic applications. In Moni Naor, editor, *ITCS 2014*, pages 169–182. ACM, January 2014.

- [DILO22a] Samuel Dittmer, Yuval Ishai, Steve Lu, and Rafail Ostrovsky. Authenticated garbling from simple correlations. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part IV*, volume 13510 of *LNCS*, pages 57–87. Springer, Heidelberg, August 2022.
- [DILO22b] Samuel Dittmer, Yuval Ishai, Steve Lu, and Rafail Ostrovsky. Improving line-point zero knowledge: Two multiplications for the price of one. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 829–841. ACM Press, November 2022.
- [DIO21] Samuel Dittmer, Yuval Ishai, and Rafail Ostrovsky. Line-point zero knowledge and its applications. In *2nd Conference on Information-Theoretic Cryptography*, 2021.
- [DPSZ12] Ivan Damgård, Valerio Pastro, Nigel P. Smart, and Sarah Zakarias. Multiparty computation from somewhat homomorphic encryption. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 643–662. Springer, Heidelberg, August 2012.
- [DT17] Thomas Debris-Alazard and Jean-Pierre Tillich. Statistical decoding. In *ISIT 2017*, 2017.
- [Dum91] Ilya Dumer. On minimum distance decoding of linear codes. In *Proc. 5th Joint Soviet-Swedish Int. Workshop Inform. Theory*, 1991.
- [EB22] Andre Esser and Emanuele Bellini. Syndrome decoding estimator. In *PKC 2022*, volume 13177, 2022.
- [EHK⁺18] Andre Esser, Felix Heuer, Robert Kübler, Alexander May, and Christian Sohler. Dissection-BKW. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 638–666. Springer, Heidelberg, August 2018.
- [EKM17] Andre Esser, Robert Kübler, and Alexander May. LPN decoded. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part II*, volume 10402 of *LNCS*, pages 486–514. Springer, Heidelberg, August 2017.
- [EKZ21] Andre Esser, Robert Kübler, and Floyd Zweyding. A faster algorithm for finding closest pairs in hamming metric. In *FSTTCS 2021*, volume 213, 2021.
- [EMZ22] Andre Esser, Alexander May, and Floyd Zweyding. McEliece needs a break - solving McEliece-1284 and quasi-cyclic-2918 with modern ISD. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part III*, volume 13277 of *LNCS*, pages 433–457. Springer, Heidelberg, May / June 2022.
- [FJR22] Thibault Feneuil, Antoine Joux, and Matthieu Rivain. Syndrome decoding in the head: Shorter signatures from zero-knowledge proofs. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 541–572. Springer, Heidelberg, August 2022.
- [FKI07] Marc P. C. Fossorier, Kazukuni Kobara, and Hideki Imai. Modeling bit flipping decoding based on nonorthogonal check sums with application to iterative decoding attack of mceliece cryptosystem. *IEEE Trans. Inf. Theory*, 53(1), 2007.
- [FKL⁺21] Nicholas Franzese, Jonathan Katz, Steve Lu, Rafail Ostrovsky, Xiao Wang, and Chenkai Weng. Constant-overhead zero-knowledge for RAM programs. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 178–191. ACM Press, November 2021.

- [FS09] Matthieu Finiasz and Nicolas Sendrier. Security bounds for the design of code-based cryptosystems. In Mitsuru Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 88–105. Springer, Heidelberg, December 2009.
- [GKH17] Cheikh Thiécoumba Gueye, Jean Belo Klamti, and Shoichi Hirose. Generalization of BJMM-ISD using may-ozeroov nearest neighbor algorithm over an arbitrary finite field \mathbb{F}_q . In *International Conference on Codes, Cryptology, and Information Security*, volume 10194, 2017.
- [HB01] Nicholas J. Hopper and Manuel Blum. Secure human identification protocols. In Colin Boyd, editor, *ASIACRYPT 2001*, volume 2248 of *LNCS*, pages 52–66. Springer, Heidelberg, December 2001.
- [Hir16] Shoichi Hirose. May-ozeroov algorithm for nearest-neighbor problem over $\mathbb{F}(q)$ and its application to information set decoding. In *SECITC 2016*, volume 10006, 2016.
- [HJ10] Nick Howgrave-Graham and Antoine Joux. New generic algorithms for hard knapsacks. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 235–256. Springer, Heidelberg, May / June 2010.
- [HOSS18] Carmit Hazay, Emmanuela Orsini, Peter Scholl, and Eduardo Soria-Vazquez. TinyKeys: A new approach to efficient multi-party computation. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 3–33. Springer, Heidelberg, August 2018.
- [HS13] Yann Hamdaoui and Nicolas Sendrier. A non asymptotic analysis of information set decoding. Cryptology ePrint Archive, Report 2013/162, 2013. <https://eprint.iacr.org/2013/162>.
- [HSS20] Carmit Hazay, Peter Scholl, and Eduardo Soria-Vazquez. Low cost constant round MPC combining BMR and oblivious transfer. *Journal of Cryptology*, 33(4):1732–1786, October 2020.
- [IKOS08] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Cryptography with constant computational overhead. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 433–442. ACM Press, May 2008.
- [JKPT12] Abhishek Jain, Stephan Krenn, Krzysztof Pietrzak, and Aris Tentes. Commitments and efficient zero-knowledge proofs from learning parity with noise. In Xiaoyun Wang and Kazue Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 663–680. Springer, Heidelberg, December 2012.
- [KOS16] Marcel Keller, Emmanuela Orsini, and Peter Scholl. MASCOT: Faster malicious arithmetic secure computation with oblivious transfer. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 830–842. ACM Press, October 2016.
- [KSS10] Jonathan Katz, Ji Sun Shin, and Adam Smith. Parallel and concurrent security of the HB and HB+ protocols. *Journal of Cryptology*, 23(3):402–421, July 2010.
- [LY22] Hanlin Liu and Yu Yu. A non-heuristic approach to time-space tradeoffs and optimizations for bkw. In *ASIACRYPT 2022, Part III*, LNCS. Springer, Heidelberg, December 2022.

- [Lyu05] Vadim Lyubashevsky. The parity problem in the presence of noise, decoding random linear codes, and the subset sum problem. In *RANDOM 2005*, volume 3624, 2005.
- [Meu12] Alexander Meurer. A coding-theoretic approach to cryptanalysis. Universität Bochum Ruhr, November 2012. Dissertation thesis.
- [MM11] Daniele Micciancio and Petros Mol. Pseudorandom knapsacks and the sample complexity of LWE search-to-decision reductions. In Phillip Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 465–484. Springer, Heidelberg, August 2011.
- [MMT11] Alexander May, Alexander Meurer, and Enrico Thomae. Decoding random linear codes in $\tilde{O}(2^{0.054n})$. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 107–124. Springer, Heidelberg, December 2011.
- [MO15] Alexander May and Ilya Ozerov. On computing nearest neighbors with applications to decoding of binary linear codes. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 203–228. Springer, Heidelberg, April 2015.
- [MTSB13] Rafael Misoczki, Jean-Pierre Tillich, Nicolas Sendrier, and Paulo S. L. M. Barreto. Mdpcc-eceliece: New mceliece variants from moderate density parity-check codes. In *Proceedings of the 2013 IEEE International Symposium on Information Theory, 2013*, pages 2069–2073. IEEE, 2013.
- [NNOB12] Jesper Buus Nielsen, Peter Sebastian Nordholt, Claudio Orlandi, and Sai Sheshank Burra. A new approach to practical active-secure two-party computation. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 681–700. Springer, Heidelberg, August 2012.
- [Ove06] Raphael Overbeck. Statistical decoding revisited. In Lynn Margaret Batten and Reihaneh Safavi-Naini, editors, *ACISP 06*, volume 4058 of *LNCS*, pages 283–294. Springer, Heidelberg, July 2006.
- [Pet10] Christiane Peters. Information-set decoding for linear codes over F_q . In Nicolas Sendrier, editor, *The Third International Workshop on Post-Quantum Cryptography, PQCRYPTO 2010*, pages 81–94. Springer, Heidelberg, May 2010.
- [Pra62] Eugene Prange. The use of information sets in decoding cyclic codes. *IRE Trans. Inf. Theory*, 8, 1962.
- [RR22] Srinivasan Raghuraman and Peter Rindal. Blazing fast PSI from improved OKVS and subfield VOLE. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 2505–2517. ACM Press, November 2022.
- [RS21] Peter Rindal and Phillipp Schoppmann. VOLE-PSI: Fast OPRF and circuit-PSI from vector-OLE. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part II*, volume 12697 of *LNCS*, pages 901–930. Springer, Heidelberg, October 2021.
- [Sen11] Nicolas Sendrier. Decoding one out of many. In Bo-Yin Yang, editor, *Post-Quantum Cryptography - 4th International Workshop, PQCrypto 2011*, pages 51–67. Springer, Heidelberg, November / December 2011.

- [SGRR19] Philipp Schoppmann, Adrià Gascón, Leonie Reichert, and Mariana Raykova. Distributed vector-OLE: Improved constructions and implementation. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 1055–1072. ACM Press, November 2019.
- [Ste88] Jacques Stern. A method for finding codewords of small weight. In *Coding Theory and Applications*, volume 388, 1988.
- [TS16] Rodolfo Canto Torres and Nicolas Sendrier. Analysis of information set decoding for a sub-linear error weight. In Tsuyoshi Takagi, editor, *Post-Quantum Cryptography - 7th International Workshop, PQCrypto 2016*, pages 144–161. Springer, Heidelberg, 2016.
- [Wag02] David Wagner. A generalized birthday problem. In Moti Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 288–303. Springer, Heidelberg, August 2002.
- [WRK17a] Xiao Wang, Samuel Ranellucci, and Jonathan Katz. Authenticated garbling and efficient maliciously secure two-party computation. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 21–37. ACM Press, October / November 2017.
- [WRK17b] Xiao Wang, Samuel Ranellucci, and Jonathan Katz. Global-scale secure multiparty computation. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 39–56. ACM Press, October / November 2017.
- [WYKW21] Chenkai Weng, Kang Yang, Jonathan Katz, and Xiao Wang. Wolverine: Fast, scalable, and communication-efficient zero-knowledge proofs for boolean and arithmetic circuits. In *2021 IEEE Symposium on Security and Privacy*, pages 1074–1091. IEEE Computer Society Press, May 2021.
- [WYX⁺21] Chenkai Weng, Kang Yang, Xiang Xie, Jonathan Katz, and Xiao Wang. Mystique: Efficient conversions for zero-knowledge proofs with applications to machine learning. In Michael Bailey and Rachel Greenstadt, editors, *USENIX Security 2021*, pages 501–518. USENIX Association, August 2021.
- [WYY⁺22] Chenkai Weng, Kang Yang, Zhaomin Yang, Xiang Xie, and Xiao Wang. AntMan: Interactive zero-knowledge proofs with sublinear communication. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 2901–2914. ACM Press, November 2022.
- [YS16] Yu Yu and John P. Steinberger. Pseudorandom functions in almost constant depth from low-noise LPN. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 154–183. Springer, Heidelberg, May 2016.
- [YSWW21] Kang Yang, Pratik Sarkar, Chenkai Weng, and Xiao Wang. QuickSilver: Efficient and affordable zero-knowledge proofs for circuits and polynomials over any field. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 2986–3001. ACM Press, November 2021.
- [YWL⁺20] Kang Yang, Chenkai Weng, Xiao Lan, Jiang Zhang, and Xiao Wang. Ferret: Fast extension for correlated OT with small communication. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020*, pages 1607–1626. ACM Press, November 2020.

- [YWZ20] Kang Yang, Xiao Wang, and Jiang Zhang. More efficient MPC from improved triple generation and authenticated garbling. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020*, pages 1627–1646. ACM Press, November 2020.
- [YZW⁺19] Yu Yu, Jiang Zhang, Jian Weng, Chun Guo, and Xiangxue Li. Collision resistant hashing from sub-exponential learning parity with noise. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part II*, volume 11922 of *LNCS*, pages 3–24. Springer, Heidelberg, December 2019.

A Proofs of Theorems and Lemmas

Lemma 4. Let $\mathbf{A}_i, \mathbf{P}_i, \mathbf{r}_i^\top$ for $i \in [1, \alpha t]$, \mathbf{B} and \mathbf{M} be as defined in the proof of Theorem 2. Then,

$$\text{SD} \left(\left(\left[\begin{array}{c} [\mathbf{A}_1 \parallel \mathbf{P}_1] \cdot \mathbf{M} \\ \mathbf{r}_1^\top - \mathbf{1}^\top (\mathbf{A}_1 \mathbf{B}) \end{array} \right], \dots, \left[\begin{array}{c} [\mathbf{A}_{\alpha t} \parallel \mathbf{P}_{\alpha t}] \cdot \mathbf{M} \\ \mathbf{r}_{\alpha t}^\top - \mathbf{1}^\top (\mathbf{A}_{\alpha t} \mathbf{B}) \end{array} \right] \right), (U_q^{(m+1) \times (k+\alpha t)})^{\alpha t} \right) \leq 2^{-k+2},$$

where $\text{SD}(\cdot, \cdot)$ denotes the statistical distance between two distributions, and $U_q^{m \times n}$ denotes the uniform distribution over $\mathbb{F}_q^{m \times n}$ independent of any other random variable in consideration.

Proof. Denote the event that $\mathbf{r}_1^\top, \dots, \mathbf{r}_{\alpha t}^\top$ are linearly independent by \mathcal{E}_{ind} , where $\Pr[\mathcal{E}_{ind}] > 1 - 2^{-k+1}$. Conditioned on \mathcal{E}_{ind} , square matrix \mathbf{M} is invertible, and thus $[\mathbf{A}_1 \parallel \mathbf{P}_1] \cdot \mathbf{M}, \dots, [\mathbf{A}_{\alpha t} \parallel \mathbf{P}_{\alpha t}] \cdot \mathbf{M}$ are identically distributed to $(U_q^{m \times (k+\alpha t)})^{\alpha t}$ and are independent of $\mathbf{r}_1^\top, \dots, \mathbf{r}_{\alpha t}^\top$ (despite that \mathbf{r}_i^\top is part of \mathbf{M}). Overall,

$$\text{SD} \left(\left(\left[\begin{array}{c} [\mathbf{A}_i \parallel \mathbf{P}_i] \cdot \mathbf{M} \\ \mathbf{r}_i^\top \end{array} \right], \mathbf{A}_i \mathbf{B} \right)_{i \in [1, \alpha t]}, \left(\left[\begin{array}{c} [\mathbf{A}_i \parallel \mathbf{P}_i] \cdot \mathbf{M} \\ U_q^{1 \times (k+\alpha t)} \end{array} \right], \mathbf{A}_i \mathbf{B} \right)_{i \in [1, \alpha t]} \right) \leq 2^{-k+1},$$

which implies

$$\text{SD} \left(\left(\left[\begin{array}{c} [\mathbf{A}_i \parallel \mathbf{P}_i] \cdot \mathbf{M} \\ \mathbf{r}_i^\top - \mathbf{1}^\top (\mathbf{A}_i \mathbf{B}) \end{array} \right] \right)_{i \in [1, \alpha t]}, \left(\left[\begin{array}{c} [\mathbf{A}_i \parallel \mathbf{P}_i] \cdot \mathbf{M} \\ U_q^{1 \times (k+\alpha t)} \end{array} \right] \right)_{i \in [1, \alpha t]} \right) \leq 2^{-k+1}.$$

Combining with the fact that $([\mathbf{A}_i \parallel \mathbf{P}_i] \cdot \mathbf{M})_{i \in [1, \alpha t]}$ is 2^{-k+1} -close to $(U_q^{m \times (k+\alpha t)})^{\alpha t}$, we complete the proof by a triangle inequality. \square

Corollary 3 (Corollary 2, restated). *If decisional $(\text{Ber}, \mathbb{F}_2)$ -LPN($k, N, \mu/2^\lambda$) is hard, then computational $(\text{HW}, \mathbb{Z}_{2^\lambda})$ -LPN($k, N, t = (1 - 2^{-\lambda})\mu N$) is hard.*

Proof. By Theorem 4 it suffices to show that decisional $(\text{Ber}, \mathbb{Z}_{2^\lambda})$ -LPN implies its computational analogue, which in turns implies computational $(\text{HW}, \mathbb{Z}_{2^\lambda})$ -LPN. The former is trivial and thus it left out to show the latter. The difference is that for $e \leftarrow \text{Ber}_{\mu, N}(\mathbb{Z}_{2^\lambda})$ e has expected (instead of exact) Hamming weight $\mu N \frac{2^\lambda - 1}{2^\lambda}$. The difference is not substantial for computational problems as conditioned on $|e| = (1 - 2^{-\lambda})\mu N$, which has probability $\Omega(1/\sqrt{N})$, e follows the exact-weight distribution $\text{HW}_{(1-2^{-\lambda})\mu N, N}(\mathbb{Z}_{2^\lambda})$, which completes the proof. \square

Theorem 12 (Theorem 7, restated). *If computational $(\text{HW}, \mathbb{Z}_{2^\lambda})$ -LPN(k, N, t) is $(\lambda T + \text{poly}(k, N), \epsilon^{\lambda+1})$ -hard, then computational $(\text{HW}, \mathbb{F}_2)$ -LPN(k, N, t') is $(T, \frac{2\epsilon}{1 - \exp(-\delta^2 t/6)})$ -hard, where $t' = t \frac{2^{\lambda-1}}{2^\lambda - 1} (1 + \delta)$ for any constant $\delta > 0$.*

Proof. Similar to the proof of Theorem 6, we can show if there exists $\mathcal{A}_{\text{LPN}_2}$ that breaks LPN over \mathbb{F}_2 and noise $\widetilde{\text{HW}}_{t, N}(\mathbb{F}_2)$ with probability more than 2ϵ within time T , then it can be used to breaks LPN over \mathbb{Z}_{2^λ} and noise $\text{HW}_{t, N}(\mathbb{Z}_{2^\lambda})$ with probability $\epsilon^{\lambda+1}$, where $\widetilde{\text{HW}}_{t, N}(\mathbb{F}_2)$ is the distribution of e^i when $e \leftarrow \text{HW}_{t, N}(\mathbb{Z}_{2^\lambda})$. The expected weight of e^i is $t \frac{2^{\lambda-1}}{2^\lambda - 1}$, and thus by a Chernoff bound e^i is a convex combination of distributions $\text{HW}_{1, N}(\mathbb{F}_2), \dots, \text{HW}_{t', N}(\mathbb{F}_2)$ with $t' = t \frac{2^{\lambda-1}}{2^\lambda - 1} (1 + \delta)$ and $\delta > 0$ except for an error probability bounded by $\exp(-\delta^2 t/6)$. Since $\mathcal{A}_{\text{LPN}_2}$ works on LPN over \mathbb{F}_2 with noise $\text{HW}_{t', N}(\mathbb{F}_2)$, it should work on that with noise $\text{HW}_{i, N}(\mathbb{F}_2)$ of weight up to $i = t'$ (and any their convex combination) as well.¹³ Therefore, $\mathcal{A}_{\text{LPN}_2}$ that breaks $(\text{HW}, \mathbb{F}_2)$ -LPN(k, N, t') with probability $2\epsilon/(1 - \exp(-\delta^2 t/6))$ is the hypothetical algorithm desired to complete the proof. \square

¹³Strictly speaking, $\mathcal{A}_{\text{LPN}_2}$ implies such an algorithm with roughly the same complexity and success probability, which can be seen by a simple reduction.

Theorem 13 (Theorem 9, restated). *If the computational $(\text{HW}, \mathbb{F}_2)$ -LPN(k, N, t') problem can be broken (say by $\mathcal{A}_{\text{LPN}_2}$) in time T with success probability at least $(1 - \epsilon/2)$, then the computational $(\text{HW}, \mathbb{Z}_{2^\lambda})$ -LPN(k, N, t) problem can be broken by $\mathcal{A}_{\text{LPN}_{2^\lambda}}$ (see Algorithm 1) in time $\lambda T + \text{poly}(k, N)$ with success probability at least $1 - (\lambda + 1)\sqrt{\epsilon}$, where $t' = t \frac{2^{\lambda-1}}{2^\lambda - 1} (1 + \delta)$ for any δ and ϵ satisfying $\delta^2 t \geq 6 \ln(2/\epsilon)$.*

Proof. Similar to the proof of Theorem 8, as long as there is $\mathcal{A}_{\text{LPN}_{2^\lambda}}$ that succeeds in breaking the LPN over \mathbb{F}_2 and noise e^i with probability at least $(1 - \epsilon)$, then the rest follows from Markov inequality and a union bound. As analyzed in the proof of Theorem 7, e^i is $\exp(-\delta^2 t/6)$ -close to a convex combination of distributions $\text{HW}_{1,N}(\mathbb{F}_2), \dots, \text{HW}_{t',N}(\mathbb{F}_2)$ with $t' = t \frac{2^{\lambda-1}}{2^\lambda - 1} (1 + \delta)$ and $\delta > 0$. Therefore, we need $\mathcal{A}_{\text{LPN}_{2^\lambda}}$ to be successful on over \mathbb{F}_2 and noise $\text{HW}_{t',N}(\mathbb{F}_2)$ with probability at least

$$\frac{1 - \epsilon}{1 - \exp(-\delta^2 t/6)} \leq 1 - \left(\epsilon - \exp(-\delta^2 t/6) \right) \leq 1 - \frac{\epsilon}{2} .$$

□

B The Hardness of LPN over More General Rings

We generalize the reductions between LPN over $\mathbb{Z}_{p^{\lambda_1} q^{\lambda_2}}$ and LPN over $\mathbb{F}_p/\mathbb{F}_q$ for distinct primes p, q . Our techniques extend to an arbitrary integer ring with more than two prime factors as well.

We recall that every number $\mathbf{a} \in [p^{\lambda_1} q^{\lambda_2}]$ can be uniquely represented using the multi-base $(1, p, \dots, p^{\lambda_1}, p^{\lambda_1} q, \dots, p^{\lambda_1} q^{\lambda_2 - 1})$ for distinct primes p, q . We define function `DigitDecomp` that decomposes a number/vector/matrix $\mathbf{a} \in [p^{\lambda_1} q^{\lambda_2}]^{\dim}$ (where $\dim = 1, n, n_1 \times n_2$ for number, vector, or matrix respectively) by applying the operation component-wise) into the above multi-base representation, i.e.,

$$\text{DigitDecomp}(\mathbf{a}) = (\mathbf{a}^0, \mathbf{a}^1, \dots, \mathbf{a}^{\lambda_1 + \lambda_2 - 1})$$

such that $\mathbf{a}^i \in [p]^{\dim}$ and $\mathbf{a}^j \in [q]^{\dim}$ for every $i \in [\lambda_1]$ and $j \in [\lambda_1, \lambda_1 + \lambda_2]$, and $\mathbf{a} = \sum_{i=0}^{\lambda_1 - 1} p^i \mathbf{a}^i + \sum_{j=\lambda_1}^{\lambda_1 + \lambda_2 - 1} p^{\lambda_1} q^{j - \lambda_1} \mathbf{a}^j$, and its inverse `DigitDecomp`⁻¹ such that $\mathbf{a} = \text{DigitDecomp}^{-1}(\mathbf{a}^0, \mathbf{a}^1, \dots, \mathbf{a}^{\lambda_1 + \lambda_2 - 1})$.

We show how to reduce the hardness of both decisional LPN over \mathbb{F}_p and decisional LPN over \mathbb{F}_q to that of decisional LPN over $\mathbb{Z}_{p^{\lambda_1} q^{\lambda_2}}$ with the noise distributions, $\text{IndBer}_{\mu, N}(\mathbb{Z}_{p^{\lambda_1} q^{\lambda_2}})$ and $\text{IndHW}_{t, N}(\mathbb{Z}_{p^{\lambda_1} q^{\lambda_2}})$, extending the noise distributions $\text{IndBer}_{\mu, N}$ and $\text{IndHW}_{t, N}$ from over \mathbb{Z}_{2^λ} to $\mathbb{Z}_{p^{\lambda_1} q^{\lambda_2}}$.

- $e \leftarrow \text{IndBer}_{\mu, N}(\mathbb{Z}_{p^{\lambda_1} q^{\lambda_2}})$ refers to $e := \text{DigitDecomp}^{-1}(e^0, e^1, \dots, e^{\lambda_1 + \lambda_2 - 1})$ with $e^i \leftarrow \text{Ber}_{\mu, N}(\mathbb{F}_p)$ and $e^j \leftarrow \text{Ber}_{\mu, N}(\mathbb{F}_q)$ for $i \in [\lambda_1]$ and $j \in [\lambda_1, \lambda_1 + \lambda_2]$.
- $e \leftarrow \text{IndHW}_{t, N}(\mathbb{Z}_{p^{\lambda_1} q^{\lambda_2}})$ means $e := \text{DigitDecomp}^{-1}(e^0, e^1, \dots, e^{\lambda_1 + \lambda_2 - 1})$ with $e^i \leftarrow \text{HW}_{t, N}(\mathbb{F}_p)$ and $e^j \leftarrow \text{HW}_{t, N}(\mathbb{F}_q)$ for $i \in [\lambda_1]$ and $j \in [\lambda_1, \lambda_1 + \lambda_2]$.

Theorem 14 (Equivalence of Decisional LPN over $\mathbb{Z}_{p^{\lambda_1} q^{\lambda_2}}$ and $\mathbb{F}_p/\mathbb{F}_q$).

1. Both decisional $(\mathcal{D}_1, \mathbb{F}_p)$ -LPN(k, N, w) and decisional $(\mathcal{D}_1, \mathbb{F}_q)$ -LPN(k, N, w) are hard iff decisional $(\mathcal{D}_2, \mathbb{Z}_{p^{\lambda_1} q^{\lambda_2}})$ -LPN(k, N, w) is hard.
2. If decisional $(\mathcal{D}, \mathbb{Z}_{p^{\lambda_1} q^{\lambda_2}})$ -LPN(k, N, t) is hard, then both of decisional $(\mathcal{D}, \mathbb{F}_p)$ -LPN(k, N, w_p) and decisional $(\mathcal{D}, \mathbb{F}_q)$ -LPN(k, N, w_q) are hard.

where $(\mathcal{D}_1, \mathcal{D}_2, w) \in \{(\text{Ber}, \text{IndBer}, \mu), (\text{HW}, \text{IndHW}, t)\}$ and $(\mathcal{D}, w_1, w_p, w_q) \in \{(\text{HW}, t, \frac{(p-1)p^{\lambda_1-1}q^{\lambda_2}}{p^{\lambda_1}q^{\lambda_2}-1} \cdot t), (\text{Ber}, \mu, \mu, \mu)\}$.

Algorithm 2: $\mathcal{A}_{\text{LPN}_{p^{\lambda_1}q^{\lambda_2}}}$, the secret recovery algorithm on LPN over $\mathbb{Z}_{p^{\lambda_1}q^{\lambda_2}}$ ($\lambda_1 + \lambda_2 \geq 2$) with oracle access to $\mathcal{A}_{\text{LPN}_r}$ (the solver for LPN over \mathbb{F}_r) and $r \in \{p, q\}$.

Input: $(\mathcal{D}, \mathbb{Z}_{p^{\lambda_1}q^{\lambda_2}})$ -LPN(k, N, t) samples $(\mathbf{A}, \mathbf{b} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e} \pmod{p^{\lambda_1}q^{\lambda_2}})$

Output: $\mathbf{s} \in \mathbb{Z}_{p^{\lambda_1}q^{\lambda_2}}$

- 1 $(\mathbf{A}^0, \mathbf{A}^1, \dots, \mathbf{A}^{\lambda_1+\lambda_2-1}) := \text{DigitDecomp}(\mathbf{A});$
 - 2 $(\mathbf{b}^0, \mathbf{b}^1, \dots, \mathbf{b}^{\lambda_1+\lambda_2-1}) := \text{DigitDecomp}(\mathbf{b});$
 - 3 **if** $\lambda_1 \geq 1$ **then**
 - 4 $(\mathbf{s}^0, \mathbf{e}^0) \leftarrow \mathcal{A}_{\text{LPN}_p}(\mathbf{A}^0, \mathbf{b}^0);$
 - 5 $\mathbf{b}' := (\mathbf{b} - \mathbf{A} \cdot \mathbf{s}^0 - \mathbf{e}^0)/p;$
 - 6 **Return** $\mathbf{s} = \mathbf{s}^0 + p \cdot \mathcal{A}_{\text{LPN}_{p^{\lambda_1-1}q^{\lambda_2}}}(\mathbf{A}' := \mathbf{A} \pmod{p^{\lambda_1-1}q^{\lambda_2}}, \mathbf{b}')$;
 - 7 $(\mathbf{s}^0, \mathbf{e}^0) \leftarrow \mathcal{A}_{\text{LPN}_q}(\mathbf{A}^0, \mathbf{b}^0);$
 - 8 $\mathbf{b}' := (\mathbf{b} - \mathbf{A} \cdot \mathbf{s}^0 - \mathbf{e}^0)/q;$
 - 9 **Return** $\mathbf{s} = \mathbf{s}^0 + q \cdot \mathcal{A}_{\text{LPN}_{q^{\lambda_2-1}}}(\mathbf{A}' := \mathbf{A} \pmod{q^{\lambda_2-1}}, \mathbf{b}')$;
-

Proof. The proof of the first statement is essentially similar to that of Theorem 4 and Theorem 5. Let $(\mathbf{A}, \mathbf{b} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e} \pmod{p^{\lambda_1}q^{\lambda_2}})$ be the LPN over $\mathbb{Z}_{p^{\lambda_1}q^{\lambda_2}}$. Decompose the matrices and vectors into the corresponding size- $(\lambda_1 + \lambda_2)$ lists, $(\mathbf{s}^0, \mathbf{s}^1, \dots, \mathbf{s}^{\lambda_1+\lambda_2-1}) := \text{DigitDecomp}(\mathbf{s})$, $(\mathbf{e}^0, \mathbf{e}^1, \dots, \mathbf{e}^{\lambda_1+\lambda_2-1}) := \text{DigitDecomp}(\mathbf{e})$ and $(\mathbf{b}^0, \mathbf{b}^1, \dots, \mathbf{b}^{\lambda_1+\lambda_2-1}) := \text{DigitDecomp}(\mathbf{b})$. Therefore, for $i \in [\lambda_1]$ and $j \in [\lambda_1, \lambda_1 + \lambda_2]$, we can write

$$\begin{aligned} \mathbf{b}^i &= \mathbf{A}' \cdot \mathbf{s}^i + \mathbf{e}^i + f_i(\mathbf{A}, \mathbf{s}^0, \dots, \mathbf{s}^{i-1}, \mathbf{e}^0, \dots, \mathbf{e}^{i-1}) \pmod{p}, \\ \mathbf{b}^j &= \mathbf{A}'' \cdot \mathbf{s}^j + \mathbf{e}^j + f_j(\mathbf{A}, \mathbf{s}^0, \dots, \mathbf{s}^{j-1}, \mathbf{e}^0, \dots, \mathbf{e}^{j-1}) \pmod{q}, \end{aligned}$$

where $\mathbf{A}' = \mathbf{A} \pmod{p}$, $\mathbf{A}'' = \mathbf{A} \pmod{q}$ and f_i (resp., f_j) is the sum of all other terms involving the individual components of \mathbf{s} and \mathbf{e} with index up to $i - 1$ (resp., $j - 1$). Define the hybrid distributions $H_0, \dots, H_{\lambda_1+\lambda_2}$ as

$$\begin{aligned} H_0 &= (\mathbf{A}, \mathbf{u}_0, \dots, \mathbf{u}_{k-1}, \mathbf{u}_k \dots, \mathbf{u}_{\lambda_1+\lambda_2-1}) \\ &\quad \vdots \\ H_k &= (\mathbf{A}, \mathbf{b}^0, \dots, \mathbf{b}^{k-1}, \mathbf{u}_k \dots, \mathbf{u}_{\lambda_1+\lambda_2-1}) \\ &\quad \vdots \\ H_{\lambda_1+\lambda_2} &= (\mathbf{A}, \mathbf{b}^0, \dots, \mathbf{b}^{k-1}, \mathbf{b}^k \dots, \mathbf{b}^{\lambda_1+\lambda_2-1}) \end{aligned}$$

where every $\mathbf{u}_i \leftarrow \mathbb{F}_p^N$ and $\mathbf{u}_j \leftarrow \mathbb{F}_q^N$ is sampled independently for $i \in [\lambda_1]$ and $j \in [\lambda_1, \lambda_1 + \lambda_2]$. Note that all the \mathbf{s}^k 's are independent, and by the definition of $\mathbf{e} \leftarrow \text{IndBer}_{\mu, N}(\mathbb{Z}_{p^{\lambda_1}q^{\lambda_2}})$ (resp., $\mathbf{e} \leftarrow \text{IndHW}_{t, N}(\mathbb{Z}_{p^{\lambda_1}q^{\lambda_2}})$) we have that \mathbf{e}^i follows $\text{Ber}_{\mu, N}(\mathbb{F}_p)$ (resp., $\text{HW}_{t, N}(\mathbb{F}_p)$) and \mathbf{e}^j follows $\text{Ber}_{\mu, N}(\mathbb{F}_q)$ (resp., $\text{HW}_{t, N}(\mathbb{F}_q)$) given its (independent) prefix $\mathbf{e}^0, \dots, \mathbf{e}^{i-1}$ and prefix $\mathbf{e}^0, \dots, \mathbf{e}^{j-1}$ for $i \in [\lambda_1]$ and $j \in [\lambda_1, \lambda_1 + \lambda_2]$. Therefore, all the adjacent H_{i-1} and H_i are computationally indistinguishable and so are H_0 and H_λ by a hybrid argument.

The proof of the second statement is essentially similar to that of Theorem 3. Given the LPN over ring $\mathbb{Z}_{p^{\lambda_1}q^{\lambda_2}}$ samples $(\mathbf{A}, \mathbf{b} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e} \pmod{p^{\lambda_1}q^{\lambda_2}})$, we also observe that the samples $(\mathbf{A}^0 := \mathbf{A} \pmod{r}, \mathbf{b}^0 := \mathbf{b} \pmod{r})$ constitute exactly the LPN over \mathbb{F}_r samples for noise $\mathbf{e}^0 = \mathbf{e} \pmod{r}$, where $r \in \{p, q\}$. In case that $\mathbf{e} \leftarrow \text{HW}_{t, N}(\mathbb{Z}_{p^{\lambda_1}q^{\lambda_2}})$, the noise vector $\mathbf{e} \pmod{p}$ has expected weight $t_p = \frac{(p-1)p^{\lambda_1-1}q^{\lambda_2}}{p^{\lambda_1}q^{\lambda_2}-1} \cdot t$ and the noise vector $\mathbf{e} \pmod{q}$ has expected weight $t_q = \frac{(q-1)p^{\lambda_1}q^{\lambda_2-1}}{p^{\lambda_1}q^{\lambda_2}-1} \cdot t$. This implies

that with probability $\Omega(1/\sqrt{t})$ the noise vector $e \bmod p$ follows the exact-weight distribution $\text{HW}_{t_p, N}(\mathbb{F}_p)$ and the noise vector $e \bmod q$ follows the exact-weight distribution $\text{HW}_{t_q, N}(\mathbb{F}_q)$. The proof for the second statement is likewise, except when taking $e \leftarrow \text{Ber}_{\mu, N}(\mathbb{Z}_{p^{\lambda_1} q^{\lambda_2}})$ we get $e \bmod r \sim \text{Ber}_{\mu, N}(\mathbb{F}_r)$, where $r \in \{p, q\}$. \square

Similar with the analysis of Theorem 8, we give a reduction both computational LPN over \mathbb{F}_q and \mathbb{F}_p to that over a ring $\mathbb{Z}_{p^{\lambda_1} q^{\lambda_2}}$ (see Theorem 15), extending the corresponding reduction between computational LPN over \mathbb{F}_2 and \mathbb{Z}_{2^λ} .

Lemma 5. *Let $(\mathbf{A}, \mathbf{b} = \mathbf{A} \cdot \mathbf{s} + e \bmod p^{\lambda_1} q^{\lambda_2})$ be the LPN samples over $\mathbb{Z}_{p^{\lambda_1} q^{\lambda_2}}$, then $(\mathbf{A}', \mathbf{b}')$ as defined as in line 6 (resp., line 9) of Algorithm 2 constitute the LPN samples over $\mathbb{Z}_{p^{\lambda_1-1} q^{\lambda_2}}$ (resp., over $\mathbb{Z}_{q^{\lambda_2-1}}$).*

Proof. Let $(\mathbf{A}^0, \mathbf{A}^1, \dots, \mathbf{A}^{\lambda-1})$, $(\mathbf{b}^0, \mathbf{b}^1, \dots, \mathbf{b}^{\lambda-1})$ be as per Algorithm 1, and let $(\mathbf{s}^0, \mathbf{s}^1, \dots, \mathbf{s}^{\lambda-1}) := \text{DigitDecomp}(\mathbf{s})$, $(e^0, e^1, \dots, e^{\lambda-1}) := \text{DigitDecomp}(e)$. The proof of the statement about $(\mathbf{A}', \mathbf{b}')$ defined as in line 6 and line 9 are similar with that of Lemma 3. For the statement about $(\mathbf{A}', \mathbf{b}')$ defined as in line 6, it suffices to prove $\mathbf{b}' = \mathbf{A}' \cdot \mathbf{s}' + e' \bmod p^{\lambda_1-1} q^{\lambda_2}$, where $\mathbf{s}' = (\mathbf{s} - \mathbf{s}^0)/p$ and $e' = (e - e^0)/p$ are the secret and noise of the LPN over $\mathbb{Z}_{p^{\lambda_1-1} q^{\lambda_2}}$ respectively. That is,

$$\begin{aligned} p\mathbf{b}' - p(\mathbf{A}' \cdot \mathbf{s}' + e') &= \mathbf{b} - \mathbf{A} \cdot \mathbf{s}^0 - e^0 - p(\mathbf{A}' \cdot \mathbf{s}' + e') \\ &= \mathbf{b} - \mathbf{A} \cdot \mathbf{s}^0 - e^0 - \mathbf{A} \cdot (\mathbf{s} - \mathbf{s}^0) - e + e^0 \\ &= \mathbf{b} - \mathbf{A} \cdot \mathbf{s} - e = \mathbf{0} \bmod p^{\lambda_1} q^{\lambda_2}, \end{aligned}$$

where the first equality follows from $\mathbf{b}' = (\mathbf{b} - \mathbf{A} \cdot \mathbf{s}^0 - e^0)/p$, the second is due to $p\mathbf{A}' \cdot \mathbf{s}' = p\mathbf{A} \cdot \mathbf{s}' \bmod p^{\lambda_1} q^{\lambda_2}$ and $pe' = e - e^0$, and the last is by the assumption of LPN over $\mathbb{Z}_{p^{\lambda_1-1} q^{\lambda_2}}$. For the statement about $(\mathbf{A}', \mathbf{b}')$ defined as in line 9, it suffices to prove $\mathbf{b}' = \mathbf{A}' \cdot \mathbf{s}' + e' \bmod q^{\lambda_2-1}$, where $\mathbf{s}' = (\mathbf{s} - \mathbf{s}^0)/q$ and $e' = (e - e^0)/q$ are the secret and noise of the LPN over $\mathbb{Z}_{q^{\lambda_2-1}}$ respectively. That is,

$$\begin{aligned} q\mathbf{b}' - q(\mathbf{A}' \cdot \mathbf{s}' + e') &= \mathbf{b} - \mathbf{A} \cdot \mathbf{s}^0 - e^0 - q(\mathbf{A}' \cdot \mathbf{s}' + e') \\ &= \mathbf{b} - \mathbf{A} \cdot \mathbf{s}^0 - e^0 - \mathbf{A} \cdot (\mathbf{s} - \mathbf{s}^0) - e + e^0 \\ &= \mathbf{b} - \mathbf{A} \cdot \mathbf{s} - e = \mathbf{0} \bmod q^{\lambda_2}, \end{aligned}$$

where the first equality follows from $\mathbf{b}' = (\mathbf{b} - \mathbf{A} \cdot \mathbf{s}^0 - e^0)/q$, the second is due to $q\mathbf{A}' \cdot \mathbf{s}' = q\mathbf{A} \cdot \mathbf{s}' \bmod q^{\lambda_2}$ and $qe' = e - e^0$, and the last is by the assumption of LPN over $\mathbb{Z}_{q^{\lambda_2-1}}$. \square

Theorem 15. *If computational $(\text{Ber}, \mathbb{F}_p)$ -LPN(k, N, μ) and $(\text{Ber}, \mathbb{F}_q)$ -LPN(k, N, μ) problems can be broken with probability at least $(1-\epsilon)$ in time T respectively, then the computational $(\text{Ber}, \mathbb{Z}_{p^{\lambda_1} q^{\lambda_2}})$ -LPN(k, N, μ) problem can be broken in time $(\lambda_1 + \lambda_2)T + \text{poly}(k, N)$ with success probability at least $1 - (\lambda_1 + \lambda_2 + 2)\sqrt{\epsilon}$.*

Proof. Algorithm 2 translates an $(\text{Ber}, \mathbb{Z}_{p^{\lambda_1} q^{\lambda_2}})$ -LPN(k, N, μ) instance into λ_1 $(\text{Ber}, \mathbb{F}_p)$ -LPN(k, N, μ) instances and λ_2 $(\text{Ber}, \mathbb{F}_q)$ -LPN(k, N, μ) instances, which are independent except that all $(\text{Ber}, \mathbb{F}_p)$ -LPN(k, N, μ) instances share the same random matrix $\mathbf{A} \bmod p$, all $(\text{Ber}, \mathbb{F}_q)$ -LPN(k, N, μ) instances share the same random matrix $\mathbf{A} \bmod q$ and that the noise vectors of the $\lambda_1 + \lambda_2$ instances are somehow correlated.

For contradiction assume that there exists an algorithm $\mathcal{A}_{\text{LPN}_p}$ (resp., an algorithm $\mathcal{A}_{\text{LPN}_q}$) that recovers the secret of LPN over \mathbb{F}_p (resp., the secret of LPN over \mathbb{F}_q) with probability more than 2ϵ within time T . Similar to the proof of Theorem 6, we consider distribution $\text{Ber}_{\mu, N}(\mathbb{F}_r)$, for $r \in \{p, q\}$, being sampled in two steps: first pick each coordinate with probability with probability μ independently (and let the rest with 0's), and second assign the picked coordinates with uniform random field element.

We have by a Markov inequality that for at least a $(1 - \sqrt{\epsilon})$ fraction of $(\mathbf{A} \bmod p, \text{coin}_1(e^i)) \mathcal{A}_{\text{LPN}_p}$ recovers s^i (for $i \in [\lambda_1]$) from $(\mathbf{A} \bmod p, (\mathbf{A} \bmod p) \cdot s^i + e^i)$ with probability at least $1 - \sqrt{\epsilon}$. Thus, the secret of all $(\text{Ber}, \mathbb{F}_p)$ -LPN(k, N, μ) instances can be recovered with the following probability by a union bound

$$(1 - \sqrt{\epsilon})(1 - \lambda_1 \sqrt{\epsilon}) \geq 1 - (\lambda_1 + 1)\sqrt{\epsilon} .$$

The proof about the $(\text{Ber}, \mathbb{F}_q)$ -LPN(k, N, μ) instances is likewise, and we have that the secrets of all $(\text{Ber}, \mathbb{F}_q)$ -LPN(k, N, μ) instances can be recovered with the following probability by a union bound

$$(1 - \sqrt{\epsilon})(1 - \lambda_2 \sqrt{\epsilon}) \geq 1 - (\lambda_2 + 1)\sqrt{\epsilon} .$$

Therefore, overall $\mathcal{A}_{\text{LPN}_{p^{\lambda_1} q^{\lambda_2}}}$ succeeds with the following probability by a union bound

$$1 - (\lambda_1 + \lambda_2 + 2)\sqrt{\epsilon} .$$

□

C Attacks on LPN Problems

In the following, we outline the known attacks on LPN problems and refer the reader to the related works listed below for more details.

Gauss. Gauss attack [EKM17] is the most natural extension of Gaussian elimination to recover the secret vector from an LPN instance with a Bernoulli distribution. This attack guesses a fresh batch of k non-noisy LPN samples by picking them at random in each iteration, inverts the corresponding submatrix, computes a candidate secret s' , and then verifies whether s' is correct or not. For LPN with noise rate r , this attack recovers the secret in time $\tilde{\Theta}(1/(1-r)^k)$ using $\tilde{\Theta}(1/(1-r)^k)$ samples, where $r = \mu(|\mathcal{R}| - 1)/|\mathcal{R}|$ for a Bernoulli distribution $\text{Ber}_{\mu, N}(\mathcal{R})$. When considering the concrete LPN parameters, the number of samples required is *not* achieved.

To reduce the number of samples in Gauss attack, Esser, Kübler and May [EKM17] introduced *Pooled Gauss* attack, which guesses k non-noisy samples by picking them at random from a pool of the *fixed* $N = k^2 \log^2 k$ LPN samples in each iteration, and then inverts the corresponding subsystem to get a candidate vector s' and verifies if s' is correct. For LPN with noise rate r , this attack recovers the secret in time $\frac{k^3 \log^2 k}{(1-r)^k}$ using $k^2 \log^2 k$ samples. The Pooled Gauss attack [EKM17] only considers the LPN problem with a Bernoulli distribution. In Section 5.1, we extend it to solve LPN and dual-LPN with an exact noise distribution.

As pointed out in [EKM17], Pooled Gauss attack solves the LPN problem via finding a non-noisy index set, and such a non-noisy index set is called an information set in coding theory language. Therefore, we can view Pooled Gauss as a special case of the information set decoding (ISD) algorithm, particularly Pooled Gauss algorithm resembles the well-known algorithm of Prange [Pra62].

Information set decoding (ISD). As shown in the previous subsection, solving LPN is equivalent to solving its dual variant, which is able to be interpreted as the task of decoding a linear code from its syndrome. The best-known algorithms for this task are improvements of the Prange's ISD algorithm [Pra62], which aims to find a size- t subset of the rows of the parity-check matrix \mathbf{H} that spans $\mathbf{H} \cdot e$, where recall that t is the Hamming weight of noise vector e .

For solving dual-LPN over binary field \mathbb{F}_2 , the Prange's ISD algorithm [Pra62] was gradually improved in recent years. The improved ISD algorithm [Ste88, Dum91] is called the Stern-Dumer variant (SD-ISD) in [HS13]. Later, the May-Meurer-Thomae variant (MMT-ISD) [MMT11] and the Becker-Joux-May-Meurer variant (BJMM-ISD) [BJMM12] improved the SD-ISD by using the generalized birthday

algorithm [Wag02]. We give an overview of the three ISD variants in Appendix D. Recently, several works [BBC⁺19, EB22, EMZ22] reduced the significant space consumption of the MMT-ISD and BJMM-ISD algorithms.

Compared to the case of \mathbb{F}_2 , the ISD algorithms to solve dual-LPN over larger fields are less studied. An initial study of ISD over a field \mathbb{F} with $|\mathbb{F}| > 2$ was given by Coffey and Goodman [CG90], who provided an asymptotic analysis of the Prange’s ISD algorithm over \mathbb{F} . Peters [Pet10] generalized the more efficient SD-ISD algorithm from binary field \mathbb{F}_2 to any finite field \mathbb{F} . Later, Meurer [Meu12] proposed a new generalization of the SD-ISD algorithm over any finite field.

These NN-ISD [MO15, BM18, EKZ21, Hir16, GKH17] applied “nearest neighbors” search to the SD-ISD or BJMM-ISD algorithms, and obtain better asymptotic complexities, where the works [MO15, BM18, EKZ21] focus on the case of binary field \mathbb{F}_2 , and the works [Hir16, GKH17] allow an arbitrary finite field \mathbb{F} . However, these NN-ISD introduce quite large polynomial overheads.

Statistical decoding. While Gauss and ISD attacks recover the secret vector \mathbf{s} , statistical decoding attack [AI 01, Ove06, FKI07, DT17] (a.k.a., low-weight parity-check attack in [BCGI18, BCG⁺19a, BCG⁺19b]) can directly distinguish LPN samples $(\mathbf{A}, \mathbf{b} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e})$ over a finite field \mathbb{F} from random samples (\mathbf{A}, \mathbf{u}) . The core of this attack is to find a set of parity-check vectors

$$\mathcal{V} \subset \{\mathbf{v} \mid \mathbf{v} \cdot \mathbf{A} = \mathbf{0} \text{ and } |\mathbf{v}| = w \text{ with a sufficiently small } w > 0\}.$$

While $\Pr[\langle \mathbf{v}, \mathbf{u} \rangle = 0] = 1/|\mathbb{F}|$ for random samples, $\Pr[\langle \mathbf{v}, \mathbf{b} \rangle = \langle \mathbf{v}, \mathbf{e} \rangle = 0] = 1/|\mathbb{F}| + \epsilon$ with some $\epsilon > 0$ for LPN samples, since both Hamming weights of vectors \mathbf{e} and \mathbf{v} are small. For each $\mathbf{v} \in \mathcal{V}$, this attack can compute a vote $\langle \mathbf{v}, \mathbf{y} \rangle$ with either $\mathbf{y} = \mathbf{b}$ or $\mathbf{y} = \mathbf{u}$. By repeating the process $|\mathcal{V}| = 1/\epsilon^2$ times, this attack outputs a majority of votes indicating whether $\mathbf{y} = \mathbf{b}$ or $\mathbf{y} = \mathbf{u}$.

In terms of asymptotic costs, Debris-Alazard and Tillich [DT17] show that an optimal statistical decoding algorithm requires more cost than the Prange’s ISD algorithm [Pra62] for the GV bound decoding over \mathbb{F}_2 .

Recently, Carrier, Debris-Alazard, Meyer-Hilfiger and Tillich [CDAMHT22] tweaked the framework of statistical decoding and beat the above bound by introducing the fast Fourier transform. Furthermore, they [CDAMHT22] proposed an advanced attack, called the SD 2.0 algorithm, which improved statistical decoding considerably and even outperforms the ISD algorithm for the GV bound decoding over \mathbb{F}_2 and $k < 0.3N$. The SD 2.0 algorithm [CDAMHT22] invokes two techniques to reduce the time complexity, as follows:

1. generalize the majority voting to the fast Fourier transform, leading to a smaller ϵ .
2. use the same collision technique invoked by the ISD algorithms [Ste88, Dum91, MMT11, BJMM12] to find set \mathcal{V} with a smaller w .

However, the running time of the second technology is $|\mathbb{F}|^{\theta(k)}$, which limits its advantage in analyzing the $(\text{HW}, \mathbb{F})\text{-LPN}(k, N, t)$ problem of security $|\mathbb{F}|^{\theta(k)}$. Therefore, the low-noise LPN assumption is beyond its strengths, because of the subexponential security $2^{O(k\mu)}$ of the low-noise LPN assumption, where dimension k and low noise rate $\mu = k^{-c}$ for constant $0 < c < 1$.

Our adapted SD 2.0 algorithm adjusts the SD 2.0 algorithm [CDAMHT22] for analyzing the low-noise LPN assumption by replacing the second technology with other collision techniques, e.g., [BCGI18]. However, this adapted SD algorithm still performs poorly on all the LPN parameters given in [BCGI18] (see Table 3 and Table 4 of Section 5). Moreover, in Section 5.3, we will prove that both an optimal statistical decoding attack and SD 2.0 attack (adapted to the low-noise setting) require more cost than the Prange’s ISD algorithm [Pra62] for analyzing the LPN assumptions over a large field, while allowing any weight $t = o(N)$ satisfied by the low-noise LPN assumption.

BKW. Blum, Kalai and Wassermann [BKW00] introduced the first sub-exponential algorithm (referred to as the BKW algorithm) that solves LPN problems using a sub-exponential number of samples. In particular, the BKW algorithm recovers the secret of a $(\mathcal{D}, \mathbb{F}_2)$ -LPN (k, N, t) instance with noise rate r in time $2^{O(k/\log(k/r))}$ using $2^{O(k/\log(k/r))}$ samples. Later, Lyubashevsky [Lyu05] reduced the sample complexity of the BKW algorithm to $k^{1+\epsilon}$ (for any constant $\epsilon > 0$) at the price of increasing the time complexity to $2^{O(k/\log \log(k/r))}$. The BKW algorithms have worse performance in time and the number of samples for solving low-noise LPN over larger fields. More recent works [EHK⁺18, DEM19, LY22] focused on reducing the space complexity of the BKW algorithm. Esser, Kübler and May [EKM17] also described the combinations and refinements of Pooled Gauss, ISD and BKW attacks, which are called Well-Pooled Gauss, Hybrid and Well-Pooled MMT attacks. All these attacks either require a sub-exponential number of samples that is *not* satisfied in the LPN-based protocols under the PCG framework, or take significantly more time than the previous three attacks for solving LPN with low noise rate. Thus, our analysis does not consider the BKW algorithms.

D Variants of Information Set Decoding

Following the analysis [FS09, Sen11, HS13, TS16], we summarize the ISD variants by Stern-Dumer [Ste88, Dum91], May et al. [MMT11] and Becker et al. [BJMM12]. We refer the reader to the corresponding papers and the surveys [FS09, Sen11, HS13, TS16] for a more detailed description.

D.1 Stern-Dumer Variant

The SD-ISD attack introduces two additional parameters p and ℓ , and adjusts both parameters to minimize the whole running time. Specifically, given an instance of the $(\text{HW}, \mathbb{F}_2)$ -dual-LPN $(N, N - k, t)$ problem $(\mathbf{H} \in \mathbb{F}_2^{(N-k) \times N}, \mathbf{y} = \mathbf{H} \cdot \mathbf{e} \in \mathbb{F}_2^{(N-k)})$, the SD-ISD attack first transforms the instance to the following equation (3) via partial Gaussian elimination, where \mathbf{U} is a non-singular $(N - k) \times (N - k)$ matrix and \mathbf{P} is a random $N \times N$ permutation matrix.

$$\mathbf{U} \cdot \mathbf{H} \cdot \mathbf{P} = \begin{array}{c} \begin{array}{|c|c|} \hline \begin{array}{c} N - k - \ell \\ 1 \\ \vdots \\ 1 \end{array} & \begin{array}{c} k + \ell \\ \mathbf{R}_0 \end{array} \\ \hline \begin{array}{c} \ell \\ \mathbf{0} \end{array} & \begin{array}{c} \mathbf{R}_1 \end{array} \\ \hline \end{array} \end{array}, \quad \mathbf{U} \cdot \mathbf{y} = \begin{array}{|c|} \hline \mathbf{y}_0 \\ \hline \mathbf{y}_1 \\ \hline \end{array} \quad (3)$$

Then, this attack finds $\mathbf{e}_1 \in \mathbb{F}_2^{(k+\ell)}$ such that $\mathbf{R}_1 \cdot \mathbf{e}_1 = \mathbf{y}_1$ and $|\mathbf{e}_1| \leq p$ via the following meet-in-the-middle attack:

1. For each $\mathbf{e}_{1,0} \in \mathbb{F}_2^{(k+\ell)/2}$ with $|\mathbf{e}_{1,0}| \leq p/2$, add the vector $\mathbf{R}_1 \cdot [\mathbf{e}_{1,0} \mid \mathbf{0}]$ into a sorted set \mathcal{S}_0 .
2. For each $\mathbf{e}_{1,1} \in \mathbb{F}_2^{(k+\ell)/2}$ with $|\mathbf{e}_{1,1}| \leq p/2$, add the vector $\mathbf{y}_1 - \mathbf{R}_1 \cdot [\mathbf{0} \mid \mathbf{e}_{1,1}]$ into a sorted set \mathcal{S}_1 .
3. Search for identical elements in sets \mathcal{S}_0 and \mathcal{S}_1 , and then add the corresponding vectors $(\mathbf{e}_{1,0} + \mathbf{e}_{1,1})$ into a set $\mathcal{S} \subseteq \{\mathbf{e}_1 \mid \mathbf{R}_1 \cdot \mathbf{e}_1 = \mathbf{y}_1\}$.

This attack repeats the above steps until $|\mathbf{e}_0| + |\mathbf{e}_1| \leq t$ where $\mathbf{e}_0 = \mathbf{R}_0 \cdot \mathbf{e}_1 + \mathbf{y}_0$, and then outputs a noise vector $\mathbf{e} = \mathbf{P} \cdot (\mathbf{e}_0, \mathbf{e}_1)^\top$.

Many variants [MMT11, BJMM12, MO15, BM18] improved the above step 3 of SD-ISD attack (finding candidate e_1) via the generalized birthday algorithm [Wag02], the representation technique [HJ10] and the “Nearest Neighbours” search.

D.2 May-Meurer-Thomae Variant

The May-Meurer-Thomae variant (MMT-ISD) [MMT11] replaced the birthday algorithm of Stern-Dumer variant [Ste88, Dum91] with an order-2 generalized birthday algorithm [Wag02]. This variant applied this and the representation technique [HJ10] to improve ISD asymptotically and does as following,

1. For all $e_{1,0} \in \mathbb{F}_2^{(k+\ell)/2}$ with $|e_{1,0}| \leq p/4$, store $\mathbf{R}_1 \cdot [\mathbf{0} \mid e_{1,0}]$ in sorted \mathcal{S}_1 and \mathcal{S}_3 with $\mathcal{S}_1 = \mathcal{S}_3$.
2. For all $e_{1,1} \in \mathbb{F}_2^{(k+\ell)/2}$ with $|e_{1,1}| \leq p/4$, store $\mathbf{R}_1 \cdot [e_{1,1} \mid \mathbf{0}]$ in sorted \mathcal{S}_2 and store $\mathbf{y}_1 - \mathbf{R}_1 \cdot [e_{1,1} \mid \mathbf{0}]$ in sorted \mathcal{S}_4 .
3. Search from \mathcal{S}_1 to \mathcal{S}_4 and generate $\mathcal{S} \subseteq \{(e_1 \mid \mathbf{R}_1 \cdot e_1 = \mathbf{y}_1)\}$ by an order-2 generalized birthday algorithm [Wag02].

Note that the set \mathcal{C} is a singleton in the original algorithm. By allowing several $c \in \mathcal{C}$ we allow larger values of r_1 and give more flexibility in the search for optimal parameters. A larger r_1 also allows smaller memory requirements with the same algorithmic complexity.

D.3 Becker-Joux-May-Meurer Variant

The Becker-Joux-May-Meurer variant (BJMM-ISD) [BJMM12] further applied an order 3 generalized birthday algorithm [Wag02] and does as following,

1. For all $e_{1,0} \in \mathbb{F}_2^{(k+\ell)/2}$ with $|e_{1,0}| \leq p_2/2$, store $\mathbf{R}_1 \cdot [\mathbf{0} \mid e_{1,0}]$ in sorted $\mathcal{S}_1, \mathcal{S}_3, \mathcal{S}_5$ and \mathcal{S}_7 with $\mathcal{S}_1 = \mathcal{S}_3 = \mathcal{S}_5 = \mathcal{S}_7$.
2. For all $e_{1,1} \in \mathbb{F}_2^{(k+\ell)/2}$ with $|e_{1,1}| \leq p_2/2$, store $\mathbf{R}_1 \cdot [e_{1,1} \mid \mathbf{0}]$ in sorted $\mathcal{S}_2, \mathcal{S}_4$ and \mathcal{S}_6 with $\mathcal{S}_2 = \mathcal{S}_4 = \mathcal{S}_6$ and store $\mathbf{y}_1 - \mathbf{R}_1 \cdot [e_{1,1} \mid \mathbf{0}]$ in sorted \mathcal{S}_8 .
3. Search from \mathcal{S}_1 to \mathcal{S}_8 and generate generate $\mathcal{S} \subseteq \{(e_1 \mid \mathbf{R}_1 \cdot e_1 = \mathbf{0})\}$ by an order-3 generalized birthday algorithm [Wag02], where $p_2 = O(p)$ is positive additional parameters.

D.4 Cost of ISD variants

We essentially follow and simplify the analysis done in [FS09, Sen11, HS13, TS16] and count complexity by the number of field operations. The expected run-time of ISD attack consists of the below parts.

1. We denote T_{Gauss} as the cost of the partial Gaussian elimination. A naive implementation leads to $T_{Gauss} = (N - k - \ell)N^2$ field operation. Fast linear algebra [BA21] leads to $T_{Gauss} = (N - k - \ell)N^2 / \log(N - k - \ell)$.
2. We estimate the success probability of one iteration. It is common in existing literature [Sen11] that each individual e_1 leads independently to success with the probability

$$\varepsilon(p, \ell)2^\ell \approx \binom{N - k - \ell}{t - p} 2^\ell / \binom{N}{t}.$$

It follows that the probability of success of one iteration is equal to

$$\mathcal{P}(p, \ell) \approx \varepsilon(p, \ell) 2^\ell |\mathcal{S}|$$

The expected value of the set \mathcal{S} will depend on various birthday_decoding.

3. Complexity of various birthday_decoding.
4. The final test cost $2|\mathcal{S}|N$ field operation.

Theorem 16 (SD-ISD [HS13, Ste88, Dum91]). *The $(\text{HW}, \mathbb{F}_2)$ -LPN(k, N, t) problem can be solved by the SD-ISD variant in expected time*

$$T_{SD}(N, k, t) = \min_{p, \ell} \frac{1}{\mathcal{P}(p, \ell)} \left(T_{Gauss} + 2L_0 \cdot N + 2\mathbb{E}[|\mathcal{S}|] \cdot N \right),$$

where $L_0 = |\mathcal{S}_1| = |\mathcal{S}_2| = \binom{(k+\ell)/2}{p/2}$ and $\mathbb{E}[|\mathcal{S}|] = \frac{L_0^2}{2^\ell}$.

Theorem 17 (MMT-ISD variant [HS13, MMT11]). *the $(\text{HW}, \mathbb{F}_2)$ -LPN(k, N, t) problem can be solved in expected time $T_{MMT}(N, k, t)$ by the MMT-ISD variant as below*

$$T_{MMT}(N, k, t) = \min_{\ell, r_1, p, |\mathcal{C}|} \frac{1}{\mathcal{P}(p, \ell)} \left(T_{Gauss} + |\mathcal{C}| \cdot N \cdot \left(4L_0 + \frac{2L_0^2}{2^{r_1}} + \frac{2L_0^4}{2^{\ell+r_1}} \right) \right),$$

where $L_0 = \binom{(k+\ell)/2}{p/4}$ and $|\mathcal{S}| = \frac{|\mathcal{C}|L_0^4}{2^{\ell+r_1}}$.

Theorem 18 (BJMM-ISD variant [HS13, BJMM12]). *The $(\text{HW}, \mathbb{F}_2)$ -LPN(k, N, t) problem can be solved in expected time $T_{BJMM}(N, k, t)$ by the BJMM-ISD variant as below*

$$T_{BJMM}(N, k, t) = \min_{p, \ell, r_1, r_2, e_1, e_2} \frac{1}{\mathcal{P}(p, \ell)} \left(T_{Gauss} + (8S_3 + 4C_3 + 2C_2 + 2C_1) \cdot N \right),$$

where $S_3 = \binom{(k+\ell)/2}{p_2}$, $C_3 = \frac{S_3^2}{2^{r_2}}$, $C_2 = \frac{C_3^2}{2^{r_1}}$, $C_1 = \frac{S_1^2}{2^{\ell-r_1-r_2}}$, $S_1 = \min\{\mu_2 C_2, \binom{(k+\ell)}{2r_1+r_2}\}$, $|\mathcal{S}| = \min\{\mu_1 C_1, \frac{(k+\ell)}{2^\ell}\}$,
 $\mu_1 = \frac{\binom{p_1}{e_1} \binom{(k+\ell-p_1)}{p_1-e_1}}{\binom{(k+\ell)}{p_1}}$, $\mu_2 = \frac{\binom{p_2}{e_2} \binom{(k+\ell-p_2)}{p_2-e_2}}{\binom{(k+\ell)}{p_2}}$, $p_2 = p_1/2 + e_2$ and $p_1 = p/2 + e_1$.

E Proofs of Theorem 10

Theorem 19 (Theorem 10, restated). *For $w = w(s) \in \mathbb{N}$ and a finite field \mathbb{F} with $|\mathbb{F}| \geq 4t$, the adapted SD 2.0 algorithm solves the decisional (HW, \mathbb{F}) -LPN(k, N, t) problem in time*

$$T = \min_s \left(T_1 \cdot \left(\frac{\binom{N}{t}}{\binom{N-w}{t}} \cdot \frac{2|\mathbb{F}|}{|\mathbb{F}| - 1} \right)^2 + s \log(|\mathbb{F}|) |\mathbb{F}|^s \right),$$

with constant advantage, where T_1 is the time of finding one parity check vector.

Proof. We first analyze the time complexity with a slightly different noise distribution $\text{HW}'_{t,N}(\mathbb{F})$ (to be defined below) and then show when switching from $\text{HW}'_{t,N}(\mathbb{F})$ back to $\text{HW}_{t,N}(\mathbb{F})$, the advantage doesn't drop much (remains constant) while preserving the same time complexity. We define $\text{HW}'_{t,N}(\mathbb{F})$ as: 1) sample t out of N positions uniformly at random (and set the rest $N - t$ positions to 0); 2) for each selected position, use a random element in \mathbb{F} . Using the modified noise distribution $\text{HW}'_{t,N}(\mathbb{F})$, we analyze the

cost of statistical decoding attack to solve the decisional $(\text{HW}', \mathbb{F})\text{-LPN}(k, N, t)$ problem by extending the analysis approach for binary field \mathbb{F}_2 [CDAMHT22], where the noise follows distribution $\text{HW}'_{t,N}(\mathbb{F})$. This attack can obtain a distinguishing advantage at least $\frac{1}{2}$ in time T . Then, we show that the advantage remains above $\frac{1}{4}$, when the noise distribution switches from $\text{HW}'_{t,N}(\mathbb{F})$ to $\text{HW}_{t,N}(\mathbb{F})$ by showing their statistical distance is upper bounded by $1/4$.

Consider $\mathbf{b} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e}'$ with $\mathbf{e}' \leftarrow \text{HW}'_{t,N}(\mathbb{F})$, we denote by E the event that there is no intersection between the set of non-zero positions of vector $\mathbf{v}' = [\mathbf{0}^s \quad \mathbf{v}_2]$ and that of the noisy positions of \mathbf{e}' . Note that $\langle \mathbf{v}', \mathbf{e}' \rangle = \langle \mathbf{v}_2, \mathbf{e}'_2 \rangle$, where $\mathbf{e}' \stackrel{\text{def}}{=} [\mathbf{e}'_1 \quad \mathbf{e}'_2]$ with $\mathbf{e}'_2 \in \mathbb{F}^{N-s}$. Then, we have the following:

$$\begin{aligned} \Pr[\langle \mathbf{v}', \mathbf{b} \rangle = \langle \mathbf{v}', \mathbf{e}' \rangle = 0] &= \Pr[\langle \mathbf{v}', \mathbf{e}' \rangle = 0 \mid E] \cdot \Pr[E] + \Pr[\langle \mathbf{v}', \mathbf{e}' \rangle = 0 \mid \neg E] \cdot \Pr[\neg E] \\ &= \Pr[E] + \frac{1}{|\mathbb{F}|} \cdot \Pr[\neg E] = \frac{\binom{N-w}{t}}{\binom{N}{t}} + \frac{1}{|\mathbb{F}|} \cdot \left(1 - \frac{\binom{N-w}{t}}{\binom{N}{t}}\right). \end{aligned}$$

The SD 2.0 algorithm [CDAMHT22] performs maximum likelihood decoding on $1/\epsilon^2$ votes, which takes time $O(s \log(|\mathbb{F}|)|\mathbb{F}|^s)$ by a fast Fourier transform. With bias $\epsilon = (\Pr[\langle \mathbf{v}', \mathbf{b} \rangle = 0] - 1/|\mathbb{F}|)/2 = \frac{\binom{N-w}{t}}{\binom{N}{t}} \cdot \frac{|\mathbb{F}|-1}{2|\mathbb{F}|}$, our adapted SD 2.0 algorithm solves the decisional $(\text{HW}', \mathbb{F})\text{-LPN}(k, N, t)$ with advantage at least $\frac{1}{2}$ in time

$$T = \min_s \left(T_1 \cdot \left(\frac{\binom{N}{t}}{\binom{N-w}{t}} \cdot \frac{2|\mathbb{F}|}{|\mathbb{F}|-1} \right)^2 + s \log(|\mathbb{F}|)|\mathbb{F}|^s \right).$$

We consider the statistical distance between $\mathbf{e} \leftarrow \text{HW}_{t,N}(\mathbb{F})$ and $\mathbf{e}' \leftarrow \text{HW}'_{t,N}(\mathbb{F})$

$$\text{SD}(\mathbf{e}, \mathbf{e}') \leq \text{SD}(\mathcal{U}_{(\mathbb{F} \setminus \{0\})^t}, \mathcal{U}_{\mathbb{F}^t}) = 1 - \left(1 - \frac{1}{|\mathbb{F}|}\right)^t \leq \frac{t}{|\mathbb{F}|} \leq \frac{1}{4},$$

where the first inequality is because $\text{HW}'_{t,N}(\mathbb{F})$ shares the same first-step sampling procedure as $\text{HW}_{t,N}(\mathbb{F})$ ($\mathcal{U}_{\mathcal{R}}$ denotes a uniform distribution over \mathcal{R}), and the first equality is due to that for $\mathbf{b}' \leftarrow \mathcal{U}_{\mathbb{F}^t}$, conditioned on $\mathbf{b}' \in (\mathbb{F} \setminus \{0\})^t$, which has probability $(1 - 1/|\mathbb{F}|)^t$, \mathbf{b}' follows distribution $\mathcal{U}_{(\mathbb{F} \setminus \{0\})^t}$. Therefore, taking into account the statistical difference between \mathbf{e} and \mathbf{e}' , our adapted SD 2.0 algorithm solves the decisional $(\text{HW}, \mathbb{F})\text{-LPN}(k, N, t)$ problem with advantage at least $1/4$. \square