

On Public Key Encryption from Noisy Codewords

Eli Ben-Sasson* Iddo Ben-Tov† Ivan Damgård‡ Yuval Ishai§ Noga Ron-Zewi¶

June 9, 2015

Abstract

Several well-known public key encryption schemes, including those of Alekhnovich (FOCS 2003), Regev (STOC 2005), and Gentry, Peikert and Vaikuntanathan (STOC 2008), rely on the conjectured intractability of inverting noisy linear encodings. These schemes are limited in that they either require the underlying field to grow with the security parameter, or alternatively they can work over the binary field but have a low noise entropy that gives rise to sub-exponential attacks.

Motivated by the goal of efficient public key cryptography, we study the possibility of obtaining improved security over the binary field by using different noise distributions. Inspired by an abstract encryption scheme of Micciancio (PKC 2010), we consider an abstract encryption scheme that unifies all the three schemes mentioned above and allows for arbitrary choices of the underlying field and noise distributions.

Our main result establishes an unexpected connection between the power of such encryption schemes and additive combinatorics. Concretely, we show that under the “approximate duality conjecture” from additive combinatorics (Ben-Sasson and Zewi, STOC 2011), every instance of the abstract encryption scheme over the binary field can be attacked in time $2^{O(\sqrt{n})}$, where n is the maximum of the ciphertext size and the public key size (and where the latter excludes public randomness used for specifying the code). On the flip side, counter examples to the above conjecture (if false) may lead to candidate public key encryption schemes with improved security guarantees.

We also show, using a simple argument that relies on agnostic learning of parities (Kalai, Mansour and Verbin, STOC 2008), that any such encryption scheme can be *unconditionally* attacked in time $2^{O(n/\log n)}$, where n is the ciphertext size. Combining this attack with the security proof of Regev’s cryptosystem, we immediately obtain an algorithm that solves the *learning parity with noise (LPN)* problem in time $2^{O(n/\log \log n)}$ using only $n^{1+\epsilon}$ samples, reproducing the result of Lyubashevsky (Random 2005) in a conceptually different way.

Finally, we study the possibility of instantiating the abstract encryption scheme over constant-size rings to yield encryption schemes with no decryption error. We show that over the binary field decryption errors are inherent. On the positive side, building on the construction of matching vector families (Grolmusz, Combinatorica 2000; Efremenko, STOC 2009; Dvir, Gopalan and Yekhanin, FOCS 2010), we suggest plausible candidates for secure instances of the framework over constant-size rings that can offer perfectly correct decryption.

*Department of Computer Science, Technion, Haifa, Israel, eli@cs.technion.ac.il

†Department of Computer Science, Technion, Haifa, Israel, iddo@cs.technion.ac.il

‡Department of Computer Science, Aarhus University, Aarhus, Denmark, ivan@cs.au.dk

§Department of Computer Science, Technion, Haifa, Israel, yuvali@cs.technion.ac.il. Supported by ERC grant no. 259426 CaC, ISF grant 1709/14, and BSF grant 2012378.

¶School of Mathematics, Institute for Advanced Study, Princeton, NJ, USA, nogazewi@ias.edu. Supported by the Rothschild fellowship and NSF grant number CCF-1412958.

1 Introduction

Public key encryption is one of the most intriguing concepts of modern cryptography. Decades after the introduction of the first public key encryption schemes [12, 41, 37, 30, 16], there are still only a handful of candidate constructions. While public key encryption schemes such as RSA are widely deployed in practice, their concrete efficiency, including the size of keys and ciphertexts, leaves much to be desired. In particular, there is still a considerable efficiency gap between the best known public key encryption schemes and their private key counterparts.

Motivated by the goal of finding new public key encryption schemes with attractive efficiency features, we consider an abstract encryption scheme which captures a class of known schemes that rely on the hardness of inverting a *noisy linear encoding*. This class includes the public key encryption scheme of Alekhnovich [3], whose security is based on the conjectured intractability of the “learning parity with noise” (LPN) problem, and the schemes of Regev [39] and of Gentry, Peikert and Vaikuntanathan (GPV) [17], whose security is based on the conjectured intractability of the “learning with errors” (LWE) problem.

In all of the above schemes, there is a publicly known linear code which is typically chosen at random, and the public keys and ciphertexts are generated by picking a secret uniform random codeword and adding a secret random noise vector, or alternatively by computing the syndrome of such a noisy codeword. Among other differences, the schemes differ in the choice of the underlying field and the distribution from which the noise is picked. In the schemes proposed by Regev and GPV, the field size grows polynomially with the security parameter and the noise distribution is a discrete Gaussian. The scheme of Alekhnovich has the advantage of working over the *binary* field, but its noise distribution is restricted to noise patterns whose Hamming weight is smaller than the square root of the ciphertext size and public key size.¹

The choice of binary field made by Alekhnovich [3] is attractive because of the potential for better concrete efficiency, especially on light-weight devices [22, 11, 36]. However, the choice of noise distribution made in [3] has a negative impact on efficiency since the low-weight noise makes a brute-force guessing attack possible. In particular, if we require the scheme to resist 2^t time attacks then this requires the public keys as well as the ciphertexts to be of size at least $\Omega(t^2)$, even when encrypting a single bit. In contrast, the known attacks on the schemes of Regev and GPV, using lattice algorithms, only require the public keys and ciphertexts to be of size $\Theta(t \log t)$. The main question we study is whether it is possible to obtain a similar or better level of succinctness by using linear codes over the binary field, thus obtaining a cryptosystem that enjoys the best of both worlds.

1.1 Overview of contribution

Towards a systematic study of the above question, we start by formalizing our abstract encryption scheme which unifies the schemes of Regev, GPV, and Alekhnovich, and allows for arbitrary choices of the underlying field and noise distributions. This scheme is inspired by an abstract encryption scheme of Micciancio described in the online talk [32], which unifies the encryption schemes of Regev and GPV.

Next, we unconditionally rule out the possibility of instantiating the abstract encryption scheme over the binary field to yield an *optimally succinct* cryptosystem, in the sense that the ciphertexts and public keys are only $O(t)$ bits long.² This result is obtained using a simple argument that relies on a previous result of

¹We view the code specification as a global public parameter and do not count it towards the public key size. This is justified by the possibility of picking the code pseudorandomly or using special classes of codes that can be succinctly described (cf. [29, 11]).

²Recall that we do not include global public parameters, such as the specification of a random linear code, in the public key size. Currently, the only plausible candidates for public key encryption schemes that are optimally succinct in the above sense are based on special families of elliptic curves. Unlike typical code-based constructions, these schemes are inherently susceptible to

Kalai et al. on agnostic learning of parities [23]. Combining this attack with the security proof of Regev’s cryptosystem [39] immediately yields an algorithm that solves the learning parity with noise (LPN) problem in time $2^{O(n/\log \log n)}$ using only $n^{1+\epsilon}$ samples, providing a conceptually different proof for the main result of Lyubashevsky [28].

Our main result establishes an unexpected connection between the power of such encryption schemes and additive combinatorics. We show that under a conjecture from additive combinatorics it is also impossible to obtain near-optimal succinctness over the binary field in the case in which the decryption error of a single encryption is a sufficiently small constant. More concretely, every instance of the abstract encryption scheme over the binary field can be attacked in time $2^{O(\sqrt{n})}$, where n is the maximum of the ciphertext size and the public key size. This suggests that the parameters of Alekhnovich’s original construction cannot be significantly improved by choosing different noise distributions.

The high level idea behind this result is as follows. An encryption scheme in the unified framework is parameterized by three independent noise distributions: a distribution μ_{sk} , applied during the key generation, and distributions μ_0 and μ_1 that are used for encrypting the messages 0 and 1 respectively. To enable correct decryption with high probability, it must be the case that the distributions $\langle \mu_{sk}, \mu_0 \rangle$ and $\langle \mu_{sk}, \mu_1 \rangle$ are statistically far (where $\langle \cdot, \cdot \rangle$ denotes the inner product of independent random samples). On the other hand, the security of the scheme implies that noisy linear encoding with respect to these noise distributions must be one-way, and in particular these distributions should *not* satisfy certain combinatorial properties that enable an adversary to guess the noise and solve the resulting system of linear equations. Our conditional negative results are obtained by applying the *approximate duality conjecture* from [7] to establish limits on the existence of distributions which satisfy the above. On the flip side, counter examples to the approximate duality conjecture (if false) would give distributions μ_{sk}, μ_0, μ_1 that can potentially serve as a basis for cryptosystems (over the binary field) that resist exponential time attacks.

As a secondary contribution of this work, we study the possibility of instantiating the unified framework over constant-size rings to yield encryption schemes with no decryption error. We show that over the binary field, a small decryption error probability is inherent. On the positive side, building on the construction of matching vector families from [13], which builds in turn on the constructions of [20, 15], we suggest plausible candidates for secure instances of the framework over constant-size rings that can offer perfectly correct decryption.

Before providing a more detailed account of our results, we provide some background on the problem of noisy linear decoding and public key encryption schemes based on its conjectured hardness.

1.2 Learning parity with noise

The *learning parity with noise* (LPN) problem is the problem of solving random linear equations over \mathbb{F}_2 which are corrupted by some noise. More specifically, in this problem there is an unknown vector $s \in \mathbb{F}_2^n$, and one is given independent random samples of the form (a_i, b_i) , where a_i is a uniform random vector in \mathbb{F}_2^n , $b_i = \langle a_i, s \rangle + e_i$, and each noise bit $e_i \in \{0, 1\}$ is 1 with probability $\eta < \frac{1}{2}$ and 0 otherwise independently of a_i (all operations are performed over \mathbb{F}_2). The goal is to recover the unknown vector s from these samples. If the *noise rate* η equals 0 then this can simply be done using Gaussian elimination. When $\eta > 0$ the problem is conjectured to be intractable. Indeed, solving LPN given m samples can

quantum attacks. The work of Sahai and Waters [42] shows that public key encryption with optimally succinct ciphertexts can be based on indistinguishability obfuscation and an exponentially strong one-way function. However, obfuscation-based constructions have large public keys and their known instances are currently quite far from being practical.

be viewed as the problem of decoding a noisy codeword in a random linear code of block length m and dimension n , a longstanding problem in coding theory.

It is known that the hardness of solving the above *search* version of LPN with a uniform random unknown vector s implies the hardness of the *decision* version of LPN, namely distinguishing between samples of the form (a_i, b_i) as above and uniformly random and independent vectors in \mathbb{F}_2^{n+1} [9, 5]. From a coding theory perspective, this means that if it is hard to decode noisy random codewords in a random linear code, then the joint distribution (G, b) is pseudorandom, where G is a random generator matrix of a random linear code and b is a noisy random codeword in the code.

A naive approach for solving LPN is to search among all vectors in \mathbb{F}_2^n to find a vector s' which satisfies the largest number of equations. This algorithm takes $2^{O(n)}$ time and one can show, using the Chernoff bound, that $O(n)$ independent random samples suffice to ensure that s' will be the correct solution with high probability. In [10], Blum et al. showed that, quite surprisingly, one can solve the LPN problem in time $2^{O(n/\log n)}$. However, a drawback of this algorithm is that it requires $2^{O(n/\log n)}$ independent random samples. In [28] (see also [24]) it was shown that the number of samples could be reduced to $n^{1+\epsilon}$ at the price of increasing the running time to $2^{O(n/\log \log n)}$. More specifically, they showed that using only $n^{1+\epsilon}$ initial independent random samples one can generate additional “almost fresh” random samples by XORing sufficiently large random subsets of the initial samples. These new samples can be used in turn as an input to the algorithm of [10].

1.3 Alekhnovich’s public key encryption scheme

In 2003, Alekhnovich [3] proposed a public key encryption scheme whose security was based on the intractability of the LPN problem. Roughly speaking, this scheme can be used to encrypt a bit $\sigma \in \{0, 1\}$ as follows. Let n be a security parameter, $m = 2n$, and $k = n^{1/2-\epsilon}$. The key generation proceeds by choosing a random noise vector $e \in \mathbb{F}_2^m$ in which each entry is set to 1 with independent probability $\eta = k/m$, a uniform random $m \times n$ matrix G over the binary field, and a uniform random $w \in \text{Image}(G)$ (that is, w is uniform in the column span of G). The private key is the noise vector e and the public key is the $m \times (n+1)$ matrix $\tilde{G} = (G \mid b)$ obtained from G by appending the noisy codeword $b = w + e$ to the right of the matrix G . (As discussed above, we do not count G towards the size of the public key.)

The encryption of $\sigma = 0$ is a random vector $c \in \mathbb{F}_2^m$ of the form $c = \tilde{w} + \tilde{e}$, where \tilde{w} is a uniform random vector in $\ker(\tilde{G}^T)$ and $\tilde{e} \in \mathbb{F}_2^m$ is a random noise vector distributed identically to (but independently of) the private key e . The encryption of $\sigma = 1$ is a uniform random vector in \mathbb{F}_2^m . In order to decrypt a ciphertext $c \in \mathbb{F}_2^m$, one simply outputs the inner product $\langle c, e \rangle$. It can be easily seen that this inner product is a nearly uniform random bit when c is an encryption of 1, and is equal to the inner product $\langle e, \tilde{e} \rangle$ when c is an encryption of 0. By the birthday paradox, the inner product $\langle e, \tilde{e} \rangle$ is 0 with probability $1 - o(1)$ and consequently, by repeating the encryption process $\text{polylog}(n)$ times, one can distinguish between encryptions of 0 and 1 with negligible error probability.

The security of the above scheme can be based on the intractability of the LPN problem with noise rate η . Indeed, since the matrix \tilde{G} is indistinguishable from a uniform random matrix, the code from which \tilde{w} is picked is indistinguishable from a random linear code, implying that the noisy codeword c is also pseudorandom. However, by the choice of the noise rate η , the Hamming weight of the private key e is bounded by $n^{1/2-\epsilon/2}$ with overwhelming probability. By trying all different possibilities for such a private key, the scheme can be attacked in time $2^{O(\sqrt{n})}$.

It is instructive to consider the abstract requirements from the noise distributions e and \tilde{e} that are necessary for the correctness and security of the above scheme. To enable correct decryption with high probability, the inner product of e and \tilde{e} (where the two noise vectors are independently sampled) should be statistically

far from uniform, i.e., significantly biased towards either 0 or 1. On the other hand, a sufficient condition for security is that the LPN decision problem be intractable with respect to both of the noise distributions e and \tilde{e} . The main question that motivates this work is whether there can be other choices of noise distributions that satisfy the above correctness requirement and may provide substantially better security than the original choice of Alekhnovich.

1.4 Learning with errors

The *learning with errors* (LWE) problem, introduced by Regev for the construction of his public key encryption scheme [39], is a generalization of the LPN problem to arbitrary rings \mathbb{Z}_q . More specifically, in this problem one is given independent random samples of the form (a_i, b_i) where now a_i is a uniform random vector in \mathbb{Z}_q^n , $b_i = \langle a_i, s \rangle + e_i$ for a fixed unknown vector $s \in \mathbb{Z}_q^n$ and e_i is distributed according to some fixed distribution χ on \mathbb{Z}_q independently of a_i (all operations are performed over \mathbb{Z}_q). Concretely, the distribution χ is usually chosen to be some small *discrete Gaussian*. The goal is again to recover the unknown vector s .

As was the case with LPN, assuming that the distribution χ is sufficiently far from uniform, one can solve LWE naively in time $q^{O(n)}$ using $O(n \log q)$ samples, and the algorithm of Blum et al. [10] can be adapted to solve this problem in time $q^{O(n/\log n)}$ using $q^{O(n/\log n)}$ samples. However, what is remarkable about LWE is that its hardness can be based on the *worst-case hardness* of well-studied lattice problems. This makes all cryptographic constructions based on the hardness of LWE secure under assumptions on the worst-case hardness of these lattice problems. See the survey [40] for more details.

1.5 Public key encryption based on learning with errors

As mentioned above, Regev introduced the LWE problem as a basis for the construction of his public key encryption scheme [39] which can be used to encrypt a bit $\sigma \in \{0, 1\}$ as follows. Let n be a security parameter, $m = (1 + \epsilon)n \log q$ and $q = \text{poly}(n)$. The key generation proceeds by choosing a random noise vector $e \in \mathbb{F}_q^m$ in which each coordinate is distributed independently according to a small discrete Gaussian, a uniform random $m \times n$ matrix G over \mathbb{F}_q , and a uniform random $w \in \text{Image}(G)$. The private key is the noise vector e and the public key is the $m \times (n + 1)$ matrix $\tilde{G} = (G \mid b)$ obtained from G by appending the noisy codeword $b = w + e$ to the right of the matrix G .

The encryption of a bit $\sigma \in \{0, 1\}$ is a random vector $c \in \mathbb{F}_q^{n+1}$ of the form $c = \tilde{G}^T \cdot \tilde{e} + v_\sigma$ where \tilde{e} is a uniform random vector in $\{0, 1\}^m$ and $v_\sigma \in \mathbb{F}_q^{n+1}$ is the vector all of whose coordinates equal 0 except for the $(n + 1)$ -th coordinate which equals $\sigma \cdot \lfloor \frac{q}{2} \rfloor$. In order to decrypt a ciphertext $c \in \mathbb{F}_q^{n+1}$ one computes $\sigma' =: c_{n+1} - \langle x, P_n(c) \rangle$, where $P_n : \mathbb{F}_q^{n+1} \rightarrow \mathbb{F}_q^n$ denotes the projection on the first n coordinates and x is such that $w = Gx$, and outputs 0 if σ' is closer to 0 than to $\lfloor \frac{q}{2} \rfloor$ and 1 otherwise. Finally, it can be verified that $\sigma' = \sigma \cdot \lfloor \frac{q}{2} \rfloor + \langle e, \tilde{e} \rangle$. Consequently, if one chooses the Gaussian distribution of the coordinates of e to be small enough then $\langle e, \tilde{e} \rangle$, which is the sum of at most m such independent Gaussians, would be smaller than $\lfloor \frac{q}{4} \rfloor$ in absolute value with high probability and therefore would enable one to distinguish between encryptions of 0 and 1 with small error probability. In fact, the error here can be completely eliminated by truncating the tail of the Gaussian noise distribution.

Inspection reveals that Regev's encryption scheme described above has various similarities with the Alekhnovich encryption scheme (one of the contributions of this work is establishing a formal connection between these encryption schemes, see Section 2.1 below for more details). However, a main difference is that, as mentioned above, Alekhnovich's encryption scheme can be attacked in time $2^{O(\sqrt{n})}$ by enumerating over all possible private keys while the best known attacks on Regev's encryption scheme, using lattice

algorithms, run in time $2^{O(n)}$. This advantage of Regev’s scheme stems from the possibility to exploit the large modulus q for picking noise distributions e and \tilde{e} whose inner product is statistically far from uniform and yet the noisy decoding problem corresponding to these distributions can be conjectured to have nearly exponential hardness. Note, however, that since q is polynomial in n , the ciphertext is of size $\Omega(n \log n)$ and therefore falls slightly short of being optimally succinct.

Another related public key encryption scheme, based on the hardness of LWE, is the public key encryption scheme proposed by Gentry, Peikert and Vaikuntanathan (GPV) [17] which is described by the authors as a “dual of Regev’s scheme in which the key generation and the encryption algorithms are swapped” (again, one of the contributions of this work is establishing a formal connection between these schemes, see Section 2.1 below for more details). A useful property of the encryption scheme of [17] is that it allows an *identity-based encryption* in which arbitrary strings are allowed to serve as public keys.

1.6 Related work

Originating from the seminal work of Ajtai [1], there has been a large body of research on basing lattice-based cryptosystems on the minimal possible assumptions and improving the efficiency of such provably secure constructions. In particular, the work of Micciancio and Mol [33] considers the possibility of replacing the standard Gaussian noise by other noise distributions, which may admit a more efficient sampling algorithm, while maintaining provable security under standard assumptions. In contrast, the goal of the present work is to explore the space of constructions that *might* be secure, in the sense that they resist known attacks, regardless of the underlying intractability assumption or the way security may be argued. Moreover, unlike the work on lattice-based cryptography, our main focus is on constructions that use linear codes over the *binary* field.

As noted above, our unified framework is inspired by the abstract encryption scheme described in Micciancio’s online talk [32] which generalizes the encryption schemes of Regev and GPV. In particular, as is the case in [32], our unifying approach relies on duality between noisy codeword encoding and syndrome encoding. This duality has also been noticed and used in other settings in the context of lattice-based public key encryption, for example in [33, 43]. However, we are not aware of any formal statement in the literature relating the encryption schemes of Regev and GPV to that of Alekhnovich via duality. Such a formal relation is given by our unified framework.³

Finally, one should note that our unified framework does not capture all of the code-based and lattice-based public key encryption schemes from the literature. For instance, it does *not* capture the code-based McEliece cryptosystem and its variants [30, 35], as well as lattice- and LWE-based cryptosystems such as [2, 18, 21, 38, 29, 4, 34, 31]. However, these alternative constructions do not seem well suited to the goal of obtaining near-optimal succinctness over binary fields. The former code-based schemes require the public key size to grow quadratically with the security parameter, whereas the latter lattice-based schemes do not admit a “native” implementation over binary fields.

³A different unified view of the schemes of Regev and Alekhnovich was previously given by Lindner and Peikert [25] who suggested to add an additional noise vector in the encryption process of Regev’s scheme. This allowed them to argue about the security of Regev’s scheme using Alekhnovich-style security proof and consequently reduce key sizes in Regev’s scheme.

2 Our results in more detail

2.1 Unified framework

We start by defining a unified framework that captures the public key encryption schemes of Alekhnovich [3], Regev [39] and Gentry, Peikert and Vaikuntanathan (GPV) [17]. More specifically, for each of the schemes [3, 39, 17] we define an abstract version that we call $\Pi_{\text{Alek}}, \Pi_{\text{Reg}}, \Pi_{\text{GPV}}$, respectively, in which the field size as well as the noise distributions used in the key generation and encryption processes are allowed to be arbitrary.

We further show that for an identical choice of parameters all the abstract schemes are equivalent to each other in terms of security: Given a pair of schemes $E, E' \in \{\Pi_{\text{Alek}}, \Pi_{\text{Reg}}, \Pi_{\text{GPV}}\}$, there exists an efficiently computable randomized mapping which for every bit $\sigma \in \{0, 1\}$ maps the joint distribution of the public key pk and the encryption of σ using pk in E to the joint distribution of the public key pk' and the encryption of σ using pk' in E' .⁴ We note that the above equivalence was previously observed by Micciancio in the online talk [32] for the special case of Regev and GPV. Since we could not find a formal statement and proof of this argument in the literature we added such a formal treatment in this paper. We further observe that this equivalence holds also when considering the abstract versions Π_{Reg} of [39] and Π_{GPV} of [17], as well as the abstract version Π_{Alek} of [3].

At a high level, all the abstract schemes work as follows (see Table 1 for more details). Each of the schemes is parametrized by integers $n < m$, a field \mathbb{F}_q (whose size may depend on n), a distribution μ_{sk} over \mathbb{F}_q^m and a pair of distributions μ_0, μ_1 over \mathbb{F}_q^{m+1} . In all three schemes the private key is a random noise vector $e \sim \mu_{\text{sk}}$. The public key consists of two parts: A random linear code $C : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^m$, specified by either a uniform random generator matrix $G^T \in \mathbb{F}_q^{n \times m}$ (in Π_{Alek} and Π_{Reg}) or a uniform random parity-check matrix $H^T \in \mathbb{F}_q^{(m-n) \times m}$ (in Π_{GPV}), together with either a noisy codeword $b = w + e$ where w is a random codeword in C (in Π_{Alek} and Π_{Reg}) or its syndrome $u = H^T \cdot e$ (in Π_{GPV}).

The encryption process is similar: Let $\tilde{C} : \mathbb{F}_q^{m-n} \rightarrow \mathbb{F}_q^{m+1}$ be the code specified by the parity-check matrix

$$\tilde{G}^T = \begin{pmatrix} G & b \\ 0_n^T & -1 \end{pmatrix}^T,$$

where \tilde{G} is the $(m+1) \times (n+1)$ matrix obtained by appending the column b to the right of the matrix G and adding below a row whose first n entries equal zero and whose last entry equals -1 . Let

$$\tilde{H} = \begin{pmatrix} H \\ u^T \end{pmatrix}$$

be the $(m+1) \times (m-n)$ matrix obtained by adding the row u^T below the matrix H , and note that \tilde{H}^T is a generator matrix for the code \tilde{C} . In order to encrypt a bit $\sigma \in \{0, 1\}$ one chooses a random noise vector $\tilde{e} \sim \mu_\sigma$. The encryption of σ is either a noisy codeword $b = \tilde{w} + \tilde{e}$ where \tilde{w} is a uniform random codeword in \tilde{C} (in Π_{Alek} and Π_{GPV}) or its syndrome $\tilde{G}^T \cdot \tilde{e}$ (in Π_{Reg}).

Finally, in all the three schemes using the private key e one can obtain the inner product $\langle e \circ (-1), \tilde{e} \rangle$, where $e \circ (-1)$ denotes the vector obtained from e by adding -1 below the vector e . To enable decryption one has to choose noise distributions μ_{sk}, μ_0 and μ_1 such that it is possible to distinguish between the distributions $\langle \mu_{\text{sk}} \circ (-1), \mu_0 \rangle$ and $\langle \mu_{\text{sk}} \circ (-1), \mu_1 \rangle$ efficiently.

⁴Note that we do not claim that the original encryption schemes of Alekhnovich, Regev and GPV are equivalent to each other in terms of security but rather that for each pair of schemes $E, E' \in \{\text{Alekhnovich, Regev, GPV}\}$ one can change the field size and noise distributions in E (but not the syntactics of E') to obtain an encryption scheme that is equivalent to E' in terms of security.

2.2 Unconditional negative result

Our first negative result shows a simple unconditional attack running in time $2^{O(n/\log n)}$ on any instance of the abstract encryption scheme over the binary field. The attack uses a simple argument based on the algorithm for agnostic learning of parities of Kalai et al. [23], a powerful algorithm that learns parities with noise from *arbitrary* distributions. More specifically, this algorithm is given independent random samples of the form (a_i, b_i) , where $b_i = \langle a_i, s \rangle + e_i$ for a fixed unknown vector $s \in \mathbb{F}_2^n$ and (a_i, e_i) are distributed according to an arbitrary distribution over $\mathbb{F}_2^n \times \mathbb{F}_2$ (In particular, the e_i 's may depend on the a_i 's). Assuming that the noise bit e_i is non-zero with probability at most η (or alternatively, $b_i \neq \langle a_i, s \rangle$ with probability at most η), the algorithm returns a circuit $h : \mathbb{F}_2^n \rightarrow \mathbb{F}_2$ that errs with probability at most η on future examples, that is $\Pr_{(a_i, b_i)}[h(a_i) \neq b_i] \leq \eta$. The running time and number of samples used by this algorithm is $2^{O(n/\log n)}$ which matches the performance of the original LPN algorithm of [10]. Note that though quite powerful, this algorithm is not a proper learner since it returns an arbitrary circuit which is not necessarily a parity function. For simplicity, assume for now that the algorithm returns the original vector s .

By the equivalence of the abstract encryption schemes Π_{Alek} , Π_{Reg} and Π_{GPV} it suffices to show an attack on the encryption scheme Π_{Reg} . The property of this scheme that we shall use for the attack is that the decryption of a ciphertext $c \in \mathbb{F}_2^{n+1}$ is $c_{n+1} - \langle s, P_n(c) \rangle$ where $P_n : \mathbb{F}_2^{n+1} \rightarrow \mathbb{F}_2^n$ denotes the projection on the first n bits and $s \in \mathbb{F}_2^n$ is such that $w = Gs$. Using the public key we generate $2^{O(n/\log n)}$ samples of the form $(P_n(c'), c'_{n+1} - \xi)$ where $\xi \in \{0, 1\}$ is a uniform random bit and c' is a random encryption of ξ and feed them to the algorithm for agnostic learning of parities described above. Assuming that the decryption algorithm has low error probability we have that $c'_{n+1} - \langle s, P_n(c') \rangle = \xi$ with probability at least $1 - \eta$, or alternatively, $\langle s, P_n(c') \rangle \neq c'_{n+1} - \xi$ with probability at most η . Hence the algorithm of [23] will recover the vector s and consequently we can recover the private key $e = b - Gs$.

The attack described above has also some *positive* consequences to learning, where it can be used for learning parities corrupted by arbitrary noise distributions in sub-exponential time using a relatively small number of samples. More specifically, we observe that Regev's security proof [39], which shows that his original encryption scheme is secure assuming the hardness of LWE, can be generalized to show the security of the abstract encryption scheme under similar assumptions. In more detail, one can show that any instance of the abstract encryption scheme over an arbitrary field \mathbb{F}_q , using an arbitrary noise distribution μ_{sk} and noise distributions μ_0, μ_1 of sufficiently high min-entropy, is secure assuming the hardness of learning linear functions over \mathbb{F}_q corrupted by noise coming from the distribution μ_{sk} . We further observe that this security guarantee holds even assuming the hardness of learning such functions using a relatively small number of samples.

Stated positively, the above says that any attack on an instance of the abstract encryption scheme as above can be turned into an algorithm that learns linear functions over \mathbb{F}_q corrupted by noise coming from the distribution μ_{sk} using a relatively small number of samples. In particular, the attack described above can be turned into such an algorithm. We further observe that an instance of this latter algorithm solves the LPN problem in time $2^{O(n/\log \log n)}$ using $n^{1+\epsilon}$ samples, reproducing the result of [28] (see also [24]) in a conceptually different way.

2.3 Conditional negative results

Our main result is a (non-uniform) attack running in time $2^{O(\sqrt{m})}$ on any instance of the abstract encryption scheme over the binary field in the case in which the decryption error of a single encryption is a sufficiently small constant, assuming the 'approximate duality conjecture' of [7]. For the attacks we first formulate combinatorial properties of the distributions μ_{sk}, μ_0 and μ_1 that imply an attack on the abstract encryption

scheme and then show that these combinatorial properties are satisfied assuming the approximate duality conjecture or its variant. We elaborate on these two parts below.

Attacks based on combinatorial properties of $\mu_{\text{sk}}, \mu_0, \mu_1$. The main combinatorial property we shall use for the attacks is *sparsity*. More precisely, we say that a distribution μ over \mathbb{F}_2^m for $m \geq n$ is (n, k, ρ) -sparse if there exist k subsets $A_1, \dots, A_k \subseteq \mathbb{F}_2^m$ (not necessarily distinct) and k full rank linear transformations $L_1, \dots, L_k : \mathbb{F}_2^m \rightarrow \mathbb{F}_2^n$ (not necessarily distinct) such that $\Pr_\mu \left(\bigcup_{i=1}^k A_i \right) \geq \rho$ and $L_i(A_i)$ is constant for every $i \in [k]$. In other words, this means that there exist k affine subspaces $V_1, \dots, V_k \subseteq \mathbb{F}_2^m$, of co-dimension n each, such that with probability at least ρ a random vector sampled from μ falls into the union of these subspaces.

We show that if an instance of the abstract encryption scheme over the binary field satisfies that the distribution μ_{sk} is (n, k, ρ) -sparse or one of the distributions μ_0 or μ_1 is $(m+1-n+\log k+\log(1/\rho), k, \rho)$ -sparse, and the decryption error of a single encryption is relatively small compared to ρ , then one can attack this instance in time $O(k)$. To illustrate the idea behind our attacks assume that we are attacking $\Pi_{\text{Ale}k}$ and that the distribution μ_{sk} is (n, k, ρ) -sparse. In this case one can search for a 'good' private key e' by enumerating over all $i \in [k]$ and solving a corresponding system of linear equations to find a vector $e' \in \bigcup_{i=1}^k V_i$ and a vector $x' \in \mathbb{F}_2^n$ such that $b = Gx' + e'$. We can further test whether e' is a 'good' private key by generating random encryptions of 0 and 1 using the public key and computing the success probability of e' in decrypting these encryptions. Since the distribution μ_{sk} is (n, k, ρ) -sparse, with probability at least ρ we will succeed in finding a 'good' private key e' which can be used in turn in order to decrypt the ciphertext.

The case in which one of the distributions μ_0 or μ_1 is $(m+1-n+\log k+\log(1/\rho), k, \rho)$ -sparse is a bit more tricky. In this case it will be convenient to attack the scheme Π_{GPV} and by symmetry it suffices to show such an attack in the case in which μ_0 is $(m+1-n+\log k+\log(1/\rho), k, \rho)$ -sparse. As in the μ_{sk} case, we can still search in time $O(k)$ for $e' \in \bigcup_{i=1}^k V_i$ and a vector $x' \in \mathbb{F}_2^{m-n}$ such that $c = \tilde{H} \cdot x' + e'$. Our main observation is that since $\bigcup_{i=1}^k V_i$ is not too large, with high probability over the choice of the matrix \tilde{H} , there is no $e' \neq \tilde{e}$ such that $e' \in \bigcup_{i=1}^k V_i$ and $c = \tilde{H}x' + e'$ for some $x' \in \mathbb{F}_2^{m-n}$. This implies in turn that by enumerating over all $i \in [k]$ and solving a corresponding system of linear equations, with high probability one can verify whether $\tilde{e} \in \bigcup_{i=1}^k V_i$ and if this is the case one can also find \tilde{e} . It thus suffices to be able to distinguish between $\tilde{e} \sim \mu_0$ and $\tilde{e} \sim \mu_1$, conditioned on the event that $\tilde{e} \in \bigcup_{i=1}^k V_i$. Assuming that the decryption error is sufficiently small compared to ρ , this can be done by computing the inner product $\langle e^{(\text{sk})} \circ (-1), \tilde{e} \rangle$ with a random $e^{(\text{sk})} \sim \mu_{\text{sk}}$.

Attacks based on the approximate duality conjecture. For a pair of subsets $A, B \subseteq \mathbb{F}_2^m$ their *duality measure* is given by

$$D(A, B) = \mathbb{E}_{a \in A, b \in B} \left[(-1)^{\langle a, b \rangle} \right]. \quad (1)$$

Note that $D(A, B) = 1$ implies that $\langle a, b \rangle$ is constant. The question is what can be said about the structure of A, B when $D(A, B)$ is sufficiently large but strictly smaller than 1. The approximate duality conjecture of [7] (cf., also Conjecture 1.7.2 in [26]) postulates that in this case there exist large subsets $A' \subseteq A, B' \subseteq B$, of density at least $2^{-O(\sqrt{m})}$ inside A, B respectively, with $D(A', B') = 1$.

We note that the bound of $2^{-O(\sqrt{m})}$ in the approximate duality conjecture is tight, and to see this take $A = B = \left(\frac{m}{\sqrt{m}} \right)$ to be the set of vectors that have \sqrt{m} ones. The birthday paradox shows that $D(A, B)$ is a fixed positive constant, independent of m (in fact, taking vectors of weight $\alpha\sqrt{m}$ for α approaching 0 makes $D(A, B)$ approach 1). But it can be verified that for any pair $A' \subset A$ and $B' \subset B$ satisfying $D(A', B') = 1$,

the size of one of the sets A' or B' is a $2^{-\sqrt{m}}$ fraction of $|A|$. Such a pair is obtained by taking A' (B' respectively) to contain all vectors supported on the first (last, respectively) $m/2$ coordinates.

In [6] it was shown that assuming the well-known polynomial Freiman-Ruzsa conjecture from additive combinatorics (cf., [19]), the approximate duality conjecture holds when replacing the lower bound $2^{-O(\sqrt{m})}$ on the ratios $|A'|/|A|$ and $|B'|/|B|$ with the weaker bound of $2^{-O(m/\log m)}$. Furthermore, in [27] a version of the approximate duality conjecture over the reals was shown to hold (unconditionally) with the stronger bound of $2^{-O(\sqrt{m})}$. The approximate duality conjecture has found so far various applications in complexity theory: To the construction of two-source extractors [7], to relating rank to communication complexity [6] and to lower bounds on matching vector codes [8].

We show that the approximate duality conjecture implies that in any instance of the abstract encryption scheme over the binary field one of the distributions μ_{sk}, μ_0 or μ_1 is sparse which by the above implies an attack on this instance. To see this suppose that Π is an instance of the abstract encryption scheme over the binary field in which $\mu_{\text{sk}} \circ (-1), \mu_0, \mu_1$ are distributed uniformly over subsets $A, B_0, B_1 \subseteq \mathbb{F}_2^{m+1}$ respectively. Then by correctness of the decryption algorithm we have that either $D(A, B_0) \geq 1 - \epsilon$ or $D(A, B_1) \leq -(1 - \epsilon)$ for some constant $\epsilon < 1$. Without loss of generality assume that $D(A, B_0) \geq 1 - \epsilon$ and note that in this case the approximate duality conjecture implies that there exist subsets $A' \subseteq A, B' \subseteq B_0$, of density at least $2^{-c\sqrt{m}}$ inside A, B_0 respectively, with $D(A', B') = 1$. The latter implies in turn that $\dim(\text{span}(A')) + \dim(\text{span}(B')) \leq m + 2$. Consequently, we have that either $\dim(\text{span}(A')) \leq m + 2 - n + 2c\sqrt{m}$ in which case A' is contained in the union of $2^{2c\sqrt{m}+1}$ affine subspaces of co-dimension n and so $\mu_{\text{sk}} \circ (-1)$ is $(n, 2^{2c\sqrt{m}+1}, 2^{-c\sqrt{m}})$ -sparse or that $\dim(\text{span}(B')) \leq n - 2c\sqrt{m}$ in which case μ_0 is $(m - n + 2c\sqrt{m}, 1, 2^{-c\sqrt{m}})$ -sparse. This implies in turn an attack running in time $2^{O(\sqrt{m})}$ in the case in which the decryption error is $2^{-\Omega(\sqrt{m})}$. Note that the attack is non-uniform since the attacker needs to know the subsets A' and B' .

In order to show such an attack in the case in which $\mu_{\text{sk}}, \mu_0, \mu_1$ are general distributions, not necessarily uniform over a subset, we prove that the standard formulation of the approximate duality conjecture implies a generalized version of it that holds also when the expectation in (1) is taken over arbitrary distributions. In order to handle larger decryption errors we apply the approximate duality conjecture iteratively to obtain $t = 2^{O(\sqrt{m})}$ pairs of subsets $A_1, B_1, \dots, A_t, B_t$ such that $D(A_i, B_i) = 1$ for all $1 \leq i \leq t$ and such that the probability of being contained in the union of $\Omega(t)$ such subsets is $\Omega(1 - \epsilon)$. This implies that either $\mu_{\text{sk}} \circ (-1)$ is $(n, 2^{3c\sqrt{m}+1}, \Omega(1 - \epsilon))$ -sparse or μ_0 is $(m - n + 2c\sqrt{m}, 2^{c\sqrt{m}}, \Omega(1 - \epsilon))$ -sparse which implies in turn an attack that runs in time $2^{O(\sqrt{m})}$ in the case in which the decryption error is a sufficiently small constant.

Finally, we note that if the approximate duality conjecture is false, then a counter example to this conjecture would be a pair of sets $A, B \subseteq \mathbb{F}_2^n$ such that $D(A, B)$ is high but no large pair of subsets A', B' of A, B respectively are dual. In this case, if we let Π be a (possibly non-uniform) instance of the unified scheme in which $\mu_{\text{sk}}, \mu_0, \mu_1$ are distributed uniformly over the sets $A, B \circ 0, B \circ 1$ respectively, then the fact that $D(A, B)$ is high implies that the advantage of the decryption algorithm in Π is high. On the other hand, the lack of linear structure in the above distributions makes them secure against our brute-force linear algebra attacks which could potentially make Π secure against sub-exponential time attacks.

2.4 Perfectly correct decryption

Our last collection of results is concerned with the possibility of achieving perfectly correct decryption in the abstract encryption scheme over constant-size rings. As mentioned above, when the field size is polynomial in n , one can truncate the tail of the Gaussian noise distribution used in Regev's original encryption scheme [39] to achieve a perfectly correct decryption. We investigate whether one can achieve perfect decryption

also over constant-size rings.

Our first result in this regard is negative, showing that over the binary field any instance of the abstract encryption scheme with perfectly correct decryption can be attacked in time $\text{poly}(m)$. On the positive side, we propose to use the construction of matching vector families from [13], which builds on the constructions of [20, 15], to obtain candidates for instances of the abstract encryption scheme over constant-size rings that achieve perfectly correct decryption but resist $\text{poly}(m)$ -time attacks.

It should be noted that Dwork et al. [14] provide a general method for eliminating decryption errors in public key encryption schemes. However, applying their method has a high toll on efficiency and it only guarantees perfectly correct decryption with high probability over the randomness of the key generation.

2.5 Open problems

We end this section by highlighting several open problems for future research.

The approximate duality conjecture and its implications to public key encryption. This work presents a new connection between additive combinatorics and public key encryption by showing non-trivial attacks on any binary instance of an abstract public key encryption scheme that captures the schemes of Alekhnovich [3], Regev [39] and Gentry, Peikert and Vaikuntanathan [17], assuming the approximate duality conjecture from additive combinatorics. On the positive side, if the approximate duality conjecture is false then counter examples to this conjecture may lead to candidate binary instances of the abstract encryption scheme with improved security guarantees. This motivates further study of the connection between public key encryption from noisy codewords and additive combinatorics in general and the approximate duality conjecture in particular.

Handling larger decryption errors. Our conditional attacks only apply to encryption schemes whose error probability is bounded by a sufficiently small constant.⁵ While the error probability of any scheme can be decreased via repetition, the resulting scheme is no longer an instance of our general framework. We leave open the question of extending our negative results to the case of encryption schemes with an arbitrary (nontrivial) constant error.

Perfect decryption. We have shown that, over the binary field, our general framework *cannot* be instantiated to yield an encryption scheme with perfect decryption. We proposed a plausible approach for obtaining perfect decryption over constant-size rings by using matching vectors. The security of this construction, as well as the possibility of obtaining perfect security over constant-size *fields*, remain to be further studied.

2.6 Paper organization

The rest of this paper is organized as follows. In Section 3 we fix some notation and terminology. In Section 4 we present our unified view of the encryption schemes [3, 39, 17] by providing the formal definitions of the abstract encryption schemes Π_{Alek} , Π_{Reg} and Π_{GPV} and proving their equivalence in terms of security. In Section 5 we present our unconditional attack, running in time $2^{O(n/\log n)}$, on the abstract encryption scheme

⁵ Note that in the original variant of Alekhnovich’s scheme, the decryption algorithm has a one-sided error probability of $1/2$. However, there is a natural variant of this scheme in which the error probability is much smaller: in the distribution μ_{sk} each entry is set to 1 independently with probability $n^{-1/2-\epsilon}$ and the distributions μ_0, μ_1 are similar to μ_{sk} except that the last entry in μ_0 is always 0 and the last entry in μ_1 is always 1. It can be verified that a single encryption using this variant gives decryption error of $o(1)$.

over the binary field and consequences of this attack to learning. In Section 6 we present combinatorial properties of the distributions μ_{sk}, μ_0 and μ_1 that imply an attack on the abstract encryption scheme over the binary field. In Section 7 we show that these latter properties are satisfied assuming the approximate duality conjecture which implies an attack on the abstract encryption scheme over the binary field running in time $2^{O(\sqrt{m})}$. Finally, in Section 8 we discuss the possibility of achieving perfectly correct decryption together with super-polynomial security in the abstract encryption scheme over constant-size rings.

3 Preliminaries

We start with fixing some notation. For a prime power q let \mathbb{F}_q denote the finite field with q elements. All operations below are performed over \mathbb{F}_q and all vectors are assumed to be column vectors unless otherwise stated. For integers $m \geq n$ let $\mathcal{M}_{m \times n}^*(q)$ denote the set of all $m \times n$ full rank matrices over \mathbb{F}_q . For an integer m let $P_m : \mathbb{F}_q^{m+1} \rightarrow \mathbb{F}_q^m$ denote the projection on the first m coordinates. Let $0_m, 1_m$ denote the all-zeros and all-ones vectors of length m , respectively. For a pair of vectors u, v let $u \circ v$ denote their concatenation.

Let μ be a distribution over \mathbb{F}_q^m . For an element $a \in \mathbb{F}_q^m$ let $\Pr_\mu(a) = \Pr_{e \sim \mu}[e = a]$. The *support* $\text{supp}(\mu)$ of μ is the set containing all elements $a \in \mathbb{F}_q^m$ for which $\Pr_\mu(a) > 0$. For a subset $A \subseteq \mathbb{F}_q^m$ we let $\Pr_\mu(A) = \Pr_{e \sim \mu}[e \in A]$ and we denote by $\mu|A$ the distribution μ conditioned on the event that $e \in A$. For a pair of distributions μ, μ' over \mathbb{F}_q^m we denote by $\langle \mu, \mu' \rangle$ the distribution of $\langle e, e' \rangle$ where $e \sim \mu$ and $e' \sim \mu'$ independently. Finally, we write that $a \in_R A$ if a is chosen uniformly at random from the set A .

For the attacks we shall use the following well-known Chernoff-type tail bound for sampling, see e.g. [44].

Lemma 3.1 (Hoeffding bound for sampling). *If X is a random variable with $|X| \leq 1$ and $\hat{\mu}$ is the empirical average obtained from t samples, then*

$$\Pr[|\mathbb{E}[X] - \hat{\mu}| \geq \gamma] \leq 2 \exp(-2\gamma^2 t).$$

3.1 Public key encryption

A *public key encryption scheme* Π consists of three randomized polynomial time algorithms: the *key generation algorithm* Gen , the *encryption algorithm* Enc and the *decryption algorithm* Dec , which satisfy:

1. The key generation algorithm Gen takes as input the *security parameter* 1^n and outputs a pair of keys (sk, pk) where sk is the *private key* and pk is the *public key*. We write this as $(\text{sk}, \text{pk}) \leftarrow \text{Gen}(1^n)$.
2. The encryption algorithm Enc takes as input a public key pk and a *message bit* $\sigma \in \{0, 1\}$ and outputs a *ciphertext* c . We write this as $c \leftarrow \text{Enc}_{\text{pk}}(\sigma)$.
3. The decryption algorithm Dec takes as input a private key sk and a ciphertext c and outputs a bit $\sigma' \in \{0, 1\}$. We assume without loss of generality that Dec is deterministic and write this as $\sigma' := \text{Dec}_{\text{sk}}(c)$.

The *advantage of the decryption algorithm* is given by

$$\text{Adv}^{\text{Dec}}(n) = \Pr[\text{Dec}_{\text{sk}}(\text{Enc}_{\text{pk}}(1)) = 1] - \Pr[\text{Dec}_{\text{sk}}(\text{Enc}_{\text{pk}}(0)) = 1], \quad (2)$$

where the probabilities in (2) are taken over the internal coin tosses of the algorithms Gen and Enc . We say that the decryption algorithm is *perfectly correct* if $\text{Adv}^{\text{Dec}}(n) = 1$.

A typical choice of parameters in public key encryption schemes is that $\text{Adv}^{\text{Dec}}(n) = 1 - \eta(n)$ for $\eta(n)$ which is a negligible function in n . However, in the case where $\text{Adv}^{\text{Dec}}(n)$ is a fixed constant one can achieve $(1 - \eta(n))$ -advantage in the decryption process by repeating the key generation and encryption processes $\text{polylog}(n)$ times. In this work we are interested in negative results and our unconditional results hold even when $\text{Adv}^{\text{Dec}}(n)$ is negligible in n . Our conditional results, on the other hand, hold only if a single encryption (without repetitions) achieves advantage $\text{Adv}^{\text{Dec}}(n) = 1 - \epsilon$ where $\epsilon > 0$ is a sufficiently small constant.

A (*uniform*) attack \mathcal{A} on a public key encryption scheme Π is a randomized algorithm that takes as input a public key pk and a ciphertext c and outputs a bit $\sigma' \in \{0, 1\}$ and we write this as $\sigma' \leftarrow \mathcal{A}(\text{pk}, c)$. The *advantage of the attack* \mathcal{A} is given by

$$\text{Adv}^{\mathcal{A}}(n) = \Pr[\mathcal{A}(\text{pk}, \text{Enc}_{\text{pk}}(1)) = 1] - \Pr[\mathcal{A}(\text{pk}, \text{Enc}_{\text{pk}}(0)) = 1], \quad (3)$$

where the probabilities in (3) are taken over the internal coin tosses of the algorithms Gen and Enc as well as the attack \mathcal{A} . A *non-uniform attack* \mathcal{A} is defined similarly to the above except that it is modeled as a non-uniform Boolean circuit and we say that it has running time $t(n)$ if the associated circuit family has size $t(n)$.

4 Unified framework

In what follows we present the formal definitions of the abstract encryption schemes Π_{Alek} , Π_{Reg} and Π_{GPV} (see also summary in Table 1 below). For each of the abstract schemes Π_{Alek} , Π_{Reg} and Π_{GPV} we also discuss the choice of parameters made in the original corresponding scheme [3, 39, 17], respectively. For completeness, in Section 4.1 we provide a formal proof that these schemes are equivalent in terms of security.

General Parameters: Integers $m > n$, field \mathbb{F}_q (q may depend on n), efficiently samplable distribution μ_{sk} over \mathbb{F}_q^m , a pair of efficiently samplable distributions μ_0, μ_1 over \mathbb{F}_q^{m+1} , efficiently computable *decryption function* $g : \mathbb{F}_q \rightarrow \{0, 1\}$.

Π_{Alek} **scheme:**

- **Private key:** Choose a random vector $e \in \mathbb{F}_q^m$ according to the distribution μ_{sk} . The private key is e .
- **Public key:** Choose a uniform random matrix $G \in \mathcal{M}_{m \times n}^*(q)$ and a uniform random vector $w \in \text{Image}(G)$ and let $b = w + e$. The public key is $\tilde{G} = \begin{pmatrix} G & b \\ 0_n^T & -1 \end{pmatrix}$.
- **Encryption:** In order to encrypt a bit $\sigma \in \{0, 1\}$ choose a random vector $\tilde{e} \in \mathbb{F}_q^{m+1}$ according to the distribution μ_σ and a uniform random vector $\tilde{w} \in \ker(\tilde{G}^T)$. The encryption of σ is $\tilde{w} + \tilde{e}$.
- **Decryption:** The decryption of a vector $c \in \mathbb{F}_q^{m+1}$ is $g(\langle e \circ (-1), c \rangle)$.

Original setting of parameters in [3]: $m = 2n$, $q = 2$, μ_{sk} is the uniform distribution over all vectors in \mathbb{F}_2^m of hamming weight $n^{1/2-\epsilon}$. The distribution μ_0 is the uniform distribution over all vectors of the form $y \circ 0$ where $y \in \mathbb{F}_2^m$ has hamming weight $n^{1/2-\epsilon}$ and the distribution μ_1 is the uniform distribution over all vectors of the form $z \circ 0$ where $z \in \mathbb{F}_2^m$ is arbitrary. The decryption function g is the identity function over \mathbb{F}_2 .

Remark 4.1. In the original encryption scheme [3] the encryption algorithm was slightly different, namely the vector \tilde{w} was chosen to be a uniform random vector among all vectors in $\ker(\tilde{G}^T)$ which satisfy that $\tilde{w}_{m+1} = 0$. We note however that choosing \tilde{w} to be a uniform random vector in $\ker(\tilde{G}^T)$ does not weaken the security of the scheme. To see this it suffices to show an efficiently computable randomized mapping $\varphi : \mathbb{F}_q^{m+1} \rightarrow \mathbb{F}_q^{m+1}$ which for every vector \tilde{e} maps the distribution $w' + \tilde{e}$ where w' is a uniform random vector in $\ker(\tilde{G}^T)$ with $w'_{m+1} = 0$ to the distribution $w'' + \tilde{e}$ where w'' is a uniform random vector in $\ker(\tilde{G}^T)$. It can be verified that the function $\varphi(x) = x + v$ where $v \in_R \ker(\tilde{G}^T)$ satisfies the required properties.

Π_{Reg} **scheme:**

- **Private key:** Choose a random vector $e \in \mathbb{F}_q^m$ according to the distribution μ_{sk} . The private key is e .
- **Public key:** Choose a uniform random matrix $G \in \mathcal{M}_{m \times n}^*(q)$ and a uniform random $w \in \text{Image}(G)$ and let $b = w + e$. The public key is $\tilde{G} = \begin{pmatrix} G & b \\ 0_n^T & -1 \end{pmatrix}$.
- **Encryption:** In order to encrypt a bit $\sigma \in \{0, 1\}$ choose a random vector $\tilde{e} \in \mathbb{F}_q^{m+1}$ according to the distribution μ_σ . The encryption of σ is $\tilde{G}^T \cdot \tilde{e}$.
- **Decryption:** The decryption of a vector $c \in \mathbb{F}_q^{m+1}$ is $g(-\langle x \circ (-1), c \rangle)$ where $x \in \mathbb{F}_q^n$ is such that $b = Gx + e$.

Original setting of parameters in [39]: $m = (1 + \epsilon)(n + 1) \log q$, q is a prime number between n^2 and $2n^2$, the coordinates of μ_{sk} are i.i.d according to a discrete Gaussian $\overline{\Psi}_{\alpha(n)}$ for $\alpha(n) = o(1/(\sqrt{n} \log n))$. The distribution μ_0 is the uniform distribution over all vectors of the form $y \circ 0$ and μ_1 is the uniform distribution over all vectors of the form $y \circ (-\lfloor \frac{q}{2} \rfloor)$ where y is an arbitrary vector in $\{0, 1\}^m$. The decryption function g is given by

$$g(i) = \begin{cases} 1, & \lfloor \frac{q}{4} \rfloor \leq i \leq \lfloor \frac{3q}{4} \rfloor \\ 0, & \text{otherwise.} \end{cases}$$

Π_{GPV} **scheme:**

- **Private key:** Choose a random vector $e \in \mathbb{F}_q^m$ according to the distribution μ_{sk} . The private key is e .
- **Public key:** Choose a uniform random matrix $H \in \mathcal{M}_{m \times (m-n)}^*(q)$ and let $u = H^T \cdot e$. The public key is $\tilde{H} = \begin{pmatrix} H \\ u^T \end{pmatrix}$.
- **Encryption:** In order to encrypt a bit $\sigma \in \{0, 1\}$ choose a random vector $\tilde{e} \in \mathbb{F}_q^{m+1}$ according to the distribution μ_σ and a uniform random vector $\tilde{w} \in \text{Image}(\tilde{H})$. The encryption of σ is $\tilde{w} + \tilde{e}$.
- **Decryption:** The decryption of a vector $c \in \mathbb{F}_q^{m+1}$ is $g(\langle e \circ (-1), c \rangle)$.

Original setting of parameters in [17]: $m \geq 2(m - n) \log q$, $q \geq 5r(m + 1)$ for $r = \omega(\sqrt{\log m})$, the coordinates of μ_{sk} are i.i.d according to a discrete Gaussian $D_{\mathbb{Z}^m, r}$. The coordinates of μ_0 are i.i.d according

to a discrete Gaussian $\bar{\Psi}_{\alpha(m)}$ for $\alpha(m) = o(1/(r\sqrt{m\log m}))$ and μ_1 is distributed according to $y + 0_m \circ \lfloor \frac{q}{2} \rfloor$ where the coordinates of y are i.i.d according to $\bar{\Psi}_{\alpha(m)}$. The decryption function g is given by

$$g(i) = \begin{cases} 1, & \lfloor \frac{q}{4} \rfloor \leq i \leq \lfloor \frac{3q}{4} \rfloor \\ 0, & \text{otherwise.} \end{cases}$$

For convenience, we summarize the definitions above in the following table.

Table 1

	Π_{Alek}	Π_{Reg}	Π_{GPV}
sk	$e \sim \mu_{\text{sk}}$	$e \sim \mu_{\text{sk}}$	$e \sim \mu_{\text{sk}}$
	$sk = e$	$sk = e$	$sk = e$
pk	$G \in_R \mathcal{M}_{m \times n}^*(q)$ $w \in_R \text{Image}(G)$ $b = w + e$	$G \in_R \mathcal{M}_{m \times n}^*(q)$ $w \in_R \text{Image}(G)$ $b = w + e$	$H \in_R \mathcal{M}_{m \times (m-n)}^*(q)$ $u = H^T \cdot e$
	$\text{pk} = \tilde{G} = \begin{pmatrix} G & b \\ 0_n^T & -1 \end{pmatrix}$	$\text{pk} = \tilde{G} = \begin{pmatrix} G & b \\ 0_n^T & -1 \end{pmatrix}$	$\text{pk} = \tilde{H} = \begin{pmatrix} H \\ u^T \end{pmatrix}$
$\text{Enc}_{\text{pk}}(\sigma)$	$\tilde{e} \sim \mu_\sigma$ $\tilde{w} \in_R \ker(\tilde{G}^T)$	$\tilde{e} \sim \mu_\sigma$	$\tilde{e} \sim \mu_\sigma$ $\tilde{w} \in_R \text{Image}(\tilde{H})$
	$\text{Enc}_{\text{pk}}(\sigma) = \tilde{w} + \tilde{e}$	$\text{Enc}_{\text{pk}}(\sigma) = \tilde{G}^T \cdot \tilde{e}$	$\text{Enc}_{\text{pk}}(\sigma) = \tilde{w} + \tilde{e}$
$\text{Dec}_{\text{sk}}(c)$	$\text{Dec}_{\text{sk}}(c) = g(\langle e \circ (-1), c \rangle)$	$\text{Dec}_{\text{sk}}(c) = g(-\langle G^{-1}(b - e) \circ (-1), c \rangle)$	$\text{Dec}_{\text{sk}}(c) = g(\langle e \circ (-1), c \rangle)$

A straightforward computation gives the following.

Claim 4.2 (Advantage of decryption). *For every $\Pi \in \{\Pi_{\text{Alek}}, \Pi_{\text{Reg}}, \Pi_{\text{GPV}}\}$,*

$$\text{Adv}^{\text{Dec}}(n) = \Pr[g(\langle \mu_{\text{sk}} \circ (-1), \mu_1 \rangle) = 1] - \Pr[g(\langle \mu_{\text{sk}} \circ (-1), \mu_0 \rangle) = 1].$$

Corollary 4.3 (Necessary and sufficient conditions for decryption). *Let $\Pi \in \{\Pi_{\text{Alek}}, \Pi_{\text{Reg}}, \Pi_{\text{GPV}}\}$ and suppose that $\text{Adv}^{\text{Dec}}(n) \geq \epsilon$. Then the statistical distance between the distributions $\langle \mu_{\text{sk}} \circ (-1), \mu_0 \rangle$ and $\langle \mu_{\text{sk}} \circ (-1), \mu_1 \rangle$ is at least ϵ . Conversely, if q is constant and $\mu_{\text{sk}}, \mu_0, \mu_1$ are such that the statistical distance between the distributions $\langle \mu_{\text{sk}} \circ (-1), \mu_0 \rangle$ and $\langle \mu_{\text{sk}} \circ (-1), \mu_1 \rangle$ is at least ϵ then there exists an efficiently computable decryption function g for which $\text{Adv}^{\text{Dec}}(n) \geq \epsilon$.*

4.1 Equivalence of abstract encryption schemes

The following claim shows that for an identical setting of parameters all the abstract encryption schemes defined above are equivalent in terms of security. For completeness, we provide a proof of this claim below.

For an encryption scheme Π and a bit $\sigma \in \{0, 1\}$ let $(\text{pk}^\Pi, \text{Enc}_{\text{pk}}^\Pi(\sigma))$ denote the joint distribution of the public key and the encryption of the bit σ using this public key in Π .

Claim 4.4 (Equivalence of abstract encryption schemes). *For every $\Pi, \Pi' \in \{\Pi_{\text{Alek}}, \Pi_{\text{Reg}}, \Pi_{\text{GPV}}\}$ there exists a randomized mapping $\varphi_{\Pi \rightarrow \Pi'}$, computable in time $\text{poly}(m, q)$, such that for every bit $\sigma \in \{0, 1\}$ the distributions $\varphi_{\Pi \rightarrow \Pi'}(\text{pk}^\Pi, \text{Enc}_{\text{pk}}^\Pi(\sigma))$ and $(\text{pk}^{\Pi'}, \text{Enc}_{\text{pk}}^{\Pi'}(\sigma))$ are identical.*

Proof. It suffices to show the existence of three mappings $\varphi_1 = \varphi_{\Pi_{\text{Alek}} \rightarrow \Pi_{\text{Reg}}}$, $\varphi_2 = \varphi_{\Pi_{\text{Reg}} \rightarrow \Pi_{\text{GPV}}}$ and $\varphi_3 = \varphi_{\Pi_{\text{GPV}} \rightarrow \Pi_{\text{Alek}}}$.

The mapping φ_1 : Let $\varphi_1(\tilde{G}, c) = (\tilde{G}, \tilde{G}^T c)$ and note that $c = \tilde{w} + \tilde{e}$ for $\tilde{w} \in \ker(\tilde{G}^T)$ and so $\tilde{G}^T c = \tilde{G}^T(\tilde{w} + \tilde{e}) = \tilde{G}^T \tilde{e}$.

The mapping φ_2 : Let $\varphi_2(\tilde{G}, c) = (\tilde{H}, c')$, where $\tilde{H} = \begin{pmatrix} H \\ u^T \end{pmatrix}$, H is a uniform random matrix in $\mathcal{M}_{m \times (m-n)}^*(q)$ which satisfies $H^T \cdot G \equiv 0$, $u = H^T b$ and $c' \in \mathbb{F}_q^{m+1}$ is a uniform random solution for the system of linear equations $\tilde{G}^T \cdot c' = c$.

Since the matrix G is distributed uniformly over $\mathcal{M}_{m \times n}^*(q)$, by duality we have that the matrix H is distributed uniformly over $\mathcal{M}_{m \times (m-n)}^*(q)$. Furthermore, we have that $b = w + e$ for $w \in \text{Image}(G)$ and so $u = H^T b = H^T(w + e) = H^T e$. Finally, we also have that $c = \tilde{G}^T \cdot \tilde{e}$ and so c' is of the form $c' = \tilde{w} + \tilde{e}$ where \tilde{w} is a uniform random element in $\ker(\tilde{G}^T)$. Since $\tilde{G}^T \cdot \tilde{H} \equiv 0$ and both \tilde{H} and \tilde{G} are full rank, we have that $\text{Image}(\tilde{H}) = \ker(\tilde{G}^T)$ and so \tilde{w} is a uniform random element in $\text{Image}(\tilde{H})$.

The mapping φ_3 : Let $\varphi_3(\tilde{H}, c) = (\tilde{G}, c)$, where $\tilde{G} = \begin{pmatrix} G & b \\ 0_n^T & -1 \end{pmatrix}$, G is a uniform random matrix in $\mathcal{M}_{m \times n}^*(q)$ which satisfies $H^T \cdot G \equiv 0$ and $b \in \mathbb{F}_q^m$ is a uniform random solution for the system of linear equations $u = H^T \cdot b$.

Since H is distributed uniformly over $\mathcal{M}_{m \times (m-n)}^*(q)$, by duality we have that G is distributed uniformly over $\mathcal{M}_{m \times n}^*(q)$. Furthermore, we have that $u = H^T \cdot e$ and so b is of the form $b = w + e$ where w is a uniform random element in $\ker(H^T)$. Since $H^T \cdot G \equiv 0$ and both H and G are full rank we have that $\text{Image}(G) = \ker(H^T)$ and so w is a uniform random element in $\text{Image}(G)$. \square

5 Unconditional attack

In this section we show an unconditional simple attack, running in time $2^{O(n/\log n)}$, on any instance of the abstract encryption scheme over the binary field. The attack is based on the following algorithm for agnostic learning of parities (The theorem below is given as Theorem 2 in [23] for the special case in which $a = \log n/1000$, $b = n/a$, $\epsilon = 2^{-n^{0.99}}$ and the success probability is 0.99. The general parameters can be deduced from the proof of this theorem).

Theorem 5.1 (Agnostic learning of parities, [23]). *For any integers a, b such that $ab \geq n$ and for any $\epsilon > 0$ there exists a randomized algorithm running in time $\text{poly}(\epsilon^{-2^a}, 2^b)$ which satisfies the following guarantees for every distribution D over $(x, y) \in \mathbb{F}_2^n \times \mathbb{F}_2$. With probability at least $1 - \exp(-n)$, given $\text{poly}(\epsilon^{-2^a}, 2^b)$ independent random samples from D , the algorithm outputs a circuit computing $h : \mathbb{F}_2^n \rightarrow \mathbb{F}_2$ such that*

$$\Pr_{(x,y) \sim D}[h(x) \neq y] \leq \min_{s \in \mathbb{F}_2^n} \Pr_{(x,y) \sim D}[\langle x, s \rangle \neq y] + \epsilon.$$

Our main result in this section is the following.

Theorem 5.2. Let $\Pi \in \{\Pi_{\text{Alek}}, \Pi_{\text{Reg}}, \Pi_{\text{GPV}}\}$ be with $q = 2$ and $\text{Adv}^{\text{Dec}}(n) \geq \epsilon$. Then for any integers a, b such that $ab \geq n$ and for any $\gamma > 0$ there exists a (uniform) attack $\mathcal{A}_{\text{agnost}}$ on Π running in time $\text{poly}(\gamma^{-2^a}, 2^b, m)$ with $\text{Adv}^{\mathcal{A}_{\text{agnost}}}(n) \geq \epsilon - \gamma - \exp(-n)$.

By setting $a = \log n / 1000$, $b = n/a$ and $\gamma = 2^{-n^{0.99}}$ in the above theorem we obtain the following corollary.

Corollary 5.3. Let $\Pi \in \{\Pi_{\text{Alek}}, \Pi_{\text{Reg}}, \Pi_{\text{GPV}}\}$ be with $q = 2$ and $\text{Adv}^{\text{Dec}}(n) \geq \epsilon$. Then there exists a (uniform) attack $\mathcal{A}_{\text{agnost}}$ on Π running in time $\text{poly}(2^{n/\log n}, m)$ with $\text{Adv}^{\mathcal{A}_{\text{agnost}}}(n) \geq \epsilon - 2^{-n^{0.99}} - \exp(-n)$.

Proof of Theorem 5.2. By Claim 4.4 it suffices to prove the theorem for $\Pi = \Pi_{\text{Reg}}$ and without loss of generality we may assume that the decryption function g is the identity function over \mathbb{F}_2 . Let D be the distribution over $(x, y) \in \mathbb{F}_2^n \times \mathbb{F}_2$ where $x = G^T \cdot P_m(e')$ and $y = \langle b \circ (-1), e' \rangle - \xi$ for $\xi \in_R \{0, 1\}$ and $e' \sim \mu_\xi$. Note that D can be generated efficiently using the public key G . The attack $\mathcal{A}_{\text{agnost}}$ runs the algorithm guaranteed by Theorem 5.1 with the parameters a, b and $\gamma/2$ on the distribution D and outputs $c_{n+1} - h(P_n(c))$. By Theorem 5.1 we clearly have that the attack runs in time $\text{poly}(\gamma^{-2^a}, 2^b, m)$. It remains to analyze the advantage of the attack in guessing the message bit σ .

For a vector $y \in \mathbb{F}_2^m$ let

$$\epsilon(y) := \Pr[\langle y \circ (-1), \mu_1 \rangle = 1] - \Pr[\langle y \circ (-1), \mu_0 \rangle = 1],$$

and note that by Claim 4.2 we have that $\text{Adv}^{\text{Dec}}(n) = \mathbb{E}[\epsilon(\mu_{\text{sk}})]$. Let $s \in \mathbb{F}_2^n$ be such that $w = Gs$. Then we have that

$$\begin{aligned} & \Pr_{(x,y) \sim D}[\langle x, s \rangle \neq y] \\ &= \frac{1}{2} \cdot \Pr[\langle G^T \cdot P_m(\mu_1), s \rangle = \langle b \circ (-1), \mu_1 \rangle] + \frac{1}{2} \cdot \Pr[\langle G^T \cdot P_m(\mu_0), s \rangle = 1 + \langle b \circ (-1), \mu_0 \rangle] \\ &= \frac{1}{2} \cdot \Pr[\langle P_m(\mu_1), Gs \rangle = \langle b \circ (-1), \mu_1 \rangle] + \frac{1}{2} \cdot \Pr[\langle P_m(\mu_0), Gs \rangle = 1 + \langle b \circ (-1), \mu_0 \rangle] \\ &= \frac{1}{2} \cdot \Pr[\langle \mu_1, (b - w) \circ (-1) \rangle = 0] + \frac{1}{2} \cdot \Pr[\langle \mu_0, (b - w) \circ (-1) \rangle = 1] \\ &= \frac{1}{2} \cdot \Pr[\langle \mu_1, e \circ (-1) \rangle = 0] + \frac{1}{2} \cdot \Pr[\langle \mu_0, e \circ (-1) \rangle = 1] \\ &= \frac{1}{2} - \frac{1}{2} \cdot \left(\Pr[\langle \mu_1, e \circ (-1) \rangle = 1] - \Pr[\langle \mu_0, e \circ (-1) \rangle = 1] \right) \\ &= \frac{1 - \epsilon(e)}{2}. \end{aligned}$$

Consequently, with probability at least $1 - \exp(-n)$, the circuit h satisfies

$$\Pr_{(x,y) \sim D}[h(x) \neq y] \leq \frac{1 - \epsilon(e)}{2} + \frac{\gamma}{2} = \frac{1 - (\epsilon(e) - \gamma)}{2}.$$

Suppose that c is an encryption of a bit $\sigma \in \{0, 1\}$. Conditioned on the above, we have that

$$\begin{aligned} & \Pr[c_{n+1} - h(P_n(c)) = 1 \mid \sigma = 1] - \Pr[c_{n+1} - h(P_n(c)) = 1 \mid \sigma = 0] \\ &= \Pr[\langle b \circ (-1), \mu_1 \rangle - h(G^T \cdot P_m(\mu_1)) = 1] - \Pr[\langle b \circ (-1), \mu_0 \rangle - h(G^T \cdot P_m(\mu_0)) = 1] \\ &= 1 - \Pr[h(G^T \cdot P_m(\mu_1)) \neq 1 + \langle b \circ (-1), \mu_1 \rangle] - \Pr[h(G^T \cdot P_m(\mu_0)) \neq \langle b \circ (-1), \mu_0 \rangle] \\ &= 1 - 2\Pr_{(x,y) \sim D}[h(x) \neq y] \\ &\geq \epsilon(e) - \gamma. \end{aligned}$$

Averaging over all $e \sim \mu_{\text{sk}}$ we obtain that the advantage of the attack is at least

$$\mathbb{E}[\epsilon(\mu_{\text{sk}})] - \gamma - \exp(-n) = \text{Adv}^{\text{Dec}}(n) - \gamma - \exp(-n) \geq \epsilon - \gamma - \exp(-n).$$

□

5.1 Consequences to learning

In this section we use a generalized form of Regev's security proof (Lemma 5.4 below) to show that our unconditional attack Theorem 5.2 above can be turned into a sub-exponential algorithm that learns parities corrupted by arbitrary noise distributions using a relatively small number of samples. We then show that an instance of this latter algorithm solves the LPN problem in time $2^{O(n/\log \log n)}$ using $n^{1+\epsilon}$ samples, reproducing the result of [28] (see also [24]) in a conceptually different way.

To state Lemma 5.4 we start with a couple of definitions. First, since in the LPN problem one is interested in learning parities from uniform random linear equations which are not necessarily linearly independent, in this section we consider a variant of the generalized encryption scheme Π_{Reg} that we denote Π'_{Reg} . The encryption scheme Π'_{Reg} is defined similarly to Π_{Reg} except that the matrix G is chosen to be a uniform random $m \times n$ matrix over \mathbb{F}_q and we do not require it to be full rank. It can be verified that Theorem 5.2 holds also for Π'_{Reg} .

We shall also need slightly more refined definitions for the advantage of the decryption algorithm and the advantage of an attack. More specifically, we say that the decryption algorithm achieves an (ϵ, δ) -advantage, and denote this $\text{Adv}_{\delta}^{\text{Dec}}(n) \geq \epsilon$ if, with probability at least $1 - \delta$ over the internal coin tosses of Gen , it holds that

$$\Pr[\text{Dec}_{\text{sk}}(\text{Enc}_{\text{pk}}(1)) = 1] - \Pr[\text{Dec}_{\text{sk}}(\text{Enc}_{\text{pk}}(0)) = 1] \geq \epsilon, \quad (4)$$

where the probabilities in (4) are taken over the internal coin tosses of the algorithm Enc . Similarly, we say that an attack \mathcal{A} achieves an (ϵ, δ) -advantage, and denote this $\text{Adv}_{\delta}^{\mathcal{A}}(n) \geq \epsilon$ if, with probability at least $1 - \delta$ over the internal coin tosses of Gen , it holds that

$$\Pr[\mathcal{A}(\text{pk}, \text{Enc}_{\text{pk}}(1)) = 1] - \Pr[\mathcal{A}(\text{pk}, \text{Enc}_{\text{pk}}(0)) = 1] \geq \epsilon, \quad (5)$$

where the probabilities in (5) are taken over the internal coin tosses of the algorithm Enc as well as the attack \mathcal{A} . It can be verified that the proof of Theorem 5.2 gives $\text{Adv}_{\delta}^{\mathcal{A}_{\text{agnost}}}(n) \geq \epsilon - \gamma - \exp(-n)$ assuming that $\text{Adv}_{\delta}^{\text{Dec}}(n) \geq \epsilon$.

Finally, let μ be a distribution over \mathbb{F}_q^m . The *min-entropy* $H_{\infty}(\mu)$ of μ is the smallest integer k such that $\Pr_{\mu}(a) \leq 2^{-k}$ for every $a \in \mathbb{F}_q^m$. For integer $n \leq m$ and a vector $s \in \mathbb{F}_q^n$ let $D_{s,\mu}$ denote the distribution (G, b) where G is a uniform random $m \times n$ matrix over \mathbb{F}_q and $b = Gs + e$ for $e \sim \mu$.

Lemma 5.4 (Security of Π'_{Reg}). *Let Π be an instance of Π'_{Reg} in which q is some prime and the distributions μ_0 and μ_1 have min-entropy at least $(1+\alpha)(n+1) \log q$ each for some constant $\alpha > 0$. Suppose furthermore that there exists an attack \mathcal{A} on Π , running in time $t(n)$, with $\text{Adv}_{\delta}^{\mathcal{A}}(n) \geq \epsilon$ for $\epsilon \geq 4 \cdot 2^{-\alpha(n+1) \log q/4}$. Then for any integer ℓ there exists a randomized algorithm W , running in time $(t(n)\ell/\epsilon^2) \cdot \text{poly}(m, q)$, which for every $s \in \mathbb{F}_q^n$, given ℓ independent random samples from $D_{s,\mu_{\text{sk}}}$, outputs s with probability at least $1 - nq(\delta^{\ell} + \ell \exp(-n))$.*

The above lemma follows from Lemmas 5.4, 4.1 and 4.2 in [39] for the special case in which $m \geq (1 + \alpha)(n + 1) \log q$, μ_{sk} is i.i.d in each coordinate, μ_0 is the uniform distribution over all vectors of the form $y \circ 0$ and μ_1 is the uniform distribution over all vectors of the form $y \circ (-\lfloor \frac{q}{2} \rfloor)$ where y is an arbitrary

vector in $\{0, 1\}^m$, $t(n) = \text{poly}(n)$, $\epsilon \geq n^{-c_1}$ and $\delta \leq 1 - n^{-c_2}$. For completeness we provide a full proof for Lemma 5.4 in Appendix A. Combining this lemma with our unconditional attack Theorem 5.2 we obtain the following corollary.

Corollary 5.5. *Let Π be an instance of Π'_{Reg} in which $q = 2$ and the distributions μ_0 and μ_1 have min-entropy at least $(1 + \alpha)(n + 1)$ each for some constant $\alpha > 0$. Suppose furthermore that $\text{Adv}_\delta^{\text{Dec}}(n) \geq \epsilon$ for $\epsilon \geq 8 \cdot 2^{-\alpha(n+1)/4} + \exp(-n)$. Then for any integers a, b, ℓ such that $ab \geq n$ there exists a randomized algorithm W , running in time $\text{poly}(\epsilon^{-2^a}, 2^b, m) \cdot \ell$, which for every $s \in \mathbb{F}_2^n$, given ℓ independent random samples from $D_{s, \mu_{\text{sk}}}$, outputs s with probability at least $1 - 2n(\delta^\ell + \ell \exp(-n))$.*

An immediate consequence of the above corollary is the following result due to Lyubashevsky [28] (see also [24]).

Theorem 5.6. *For any constants $\epsilon > 0$ and $\delta < 1$ and for any $\eta < \frac{1}{2} - 2^{-(\log n)^\delta}$, there exists a randomized algorithm W , running in time $2^{O(n/\log \log n)}$, which satisfies the following guarantees for every $s \in \mathbb{F}_2^n$. Given $n^{1+\epsilon}$ independent random samples of the form (a_i, b_i) , where a_i is a uniform random vector in \mathbb{F}_2^n , $b_i = \langle a_i, s \rangle + e_i$ and $e_i \in \{0, 1\}$ is 1 with probability η and 0 otherwise independently of a_i , the algorithm W outputs s with probability at least $1 - \exp(-n)$.*

Proof. Let Π be an instance of Π'_{Reg} in which $q = 2$, $m = n^{1+\epsilon}$ and the coordinates of μ_{sk} are i.i.d where each coordinate equals 1 with probability η and 0 otherwise. Let μ_0 be the uniform distribution over all vectors of the form $y \circ 0$ and let μ_1 be the uniform distribution over all vectors of the form $y \circ 1$ where $y \in \mathbb{F}_2^m$ has hamming weight $\frac{2n}{\epsilon \log n}$. Finally, let the decryption function g be the identity function over \mathbb{F}_2 .

With the above setting of parameters, the algorithm W is given a single sample from $D_{s, \mu_{\text{sk}}}$ and its goal is to recover s . It therefore suffices to show that the conditions of Corollary 5.5 hold. For this note first that

$$H_\infty(\mu_0) = H_\infty(\mu_1) = \log \left(\frac{n^{1+\epsilon}}{\frac{2n}{\epsilon \log n}} \right) \geq \frac{2n}{\epsilon \log n} \cdot \log \left(\frac{n^{1+\epsilon}}{\frac{2n}{\epsilon \log n}} \right) > 2n. \quad (6)$$

Let $\xi = \frac{1}{2} - \eta$ and note that our assumption implies that $\xi > 2^{-(\log n)^\delta}$. Next we show that

$$\text{Adv}_{2^{-n}}^{\text{Dec}}(n) \geq \left(\frac{\xi}{2} \right)^{\frac{2n}{\epsilon \log n}}. \quad (7)$$

By Claim 4.2 we need to show that

$$\Pr_{e \sim \mu_{\text{sk}}} \left[\Pr[\langle e \circ (-1), \mu_1 \rangle = 1] - \Pr[\langle e \circ (-1), \mu_0 \rangle = 1] \geq \left(\frac{\xi}{2} \right)^{\frac{2n}{\epsilon \log n}} \right] \geq 1 - 2^{-n}.$$

By the Hoeffding bound for sampling (Lemma 3.1), a vector $e \sim \mu_{\text{sk}}$ will have at least $m(1 + \xi)/2$ zeros with probability at least

$$1 - 2 \exp(-2m(\xi/2)^2) \geq 1 - 2 \exp \left(-2n^{1+\epsilon} (2^{-(\log n)^\delta} / 2)^2 \right) \gg 1 - 2^{-n}.$$

Conditioned on the above, Claim 5.7 below implies that

$$\begin{aligned} & \Pr[\langle e \circ (-1), \mu_1 \rangle = 1] - \Pr[\langle e \circ (-1), \mu_0 \rangle = 1] \\ &= 2 \Pr[\langle e, P_m(\mu_0) \rangle = 0] - 1 \\ &\geq \left(\frac{n^{1+\epsilon} \xi - \frac{2n}{\epsilon \log n} + 1}{n^{1+\epsilon} - \frac{2n}{\epsilon \log n} + 1} \right)^{\frac{2n}{\epsilon \log n}} \\ &> \left(\frac{\xi}{2} \right)^{\frac{2n}{\epsilon \log n}} \end{aligned}$$

The conclusion follows from (6) and (7) by applying Corollary 5.5 with $a = \gamma \log \log n$ and $b = \frac{n}{\gamma \log \log n}$ for $\gamma < 1 - \delta$ and $\ell = 1$. \square

Claim 5.7 (Lemma 3 in [28]). *If a bucket contains m balls, $(\frac{1}{2} + p)m$ of which are colored white, and the rest colored black, and we select k balls at random without replacement, then the probability that we selected an even number of black balls is at least $\frac{1}{2} + \frac{1}{2} \left(\frac{2mp-k+1}{m-k+1} \right)^k$.*

6 Attacks based on combinatorial properties of $\mu_{\text{sk}}, \mu_0, \mu_1$

In Sections 6.1 and 6.2 below we present combinatorial properties of the distribution μ_{sk} and the pair of distributions μ_0, μ_1 , respectively, that imply an attack on the abstract encryption scheme over the binary field. In Section 7 we shall show that assuming the approximate duality conjecture at least one of the distributions μ_{sk}, μ_0 or μ_1 satisfies these combinatorial properties. This will imply in turn an attack on the abstract encryption scheme over the binary field assuming the approximate duality conjecture.

The main combinatorial property we shall utilize for the attacks is *sparsity*, defined as follows.

Definition 6.1 ((n, k, ρ) -sparse distribution). Suppose that μ is a distribution over \mathbb{F}_2^m for $m \geq n$. We say that μ is (n, k, ρ) -sparse if there exist k subsets $A_1, \dots, A_k \subseteq \mathbb{F}_2^m$ and k full rank linear transformations $L_1, \dots, L_k : \mathbb{F}_2^m \rightarrow \mathbb{F}_2^n$ such that $\Pr_{\mu} \left(\bigcup_{i=1}^k A_i \right) \geq \rho$ and $L_i(A_i)$ is constant for every $i \in [k]$.

Note that A_1, \dots, A_k and L_1, \dots, L_k in the definition above are not required to be distinct. At a high level, assuming that one of the noise distributions μ_{sk}, μ_0 or μ_1 is sparse one can 'guess' the noise vector used in the key generation process (in the case in which μ_{sk} is sparse) or in the encryption process (in the case in which μ_0 or μ_1 are sparse) by enumerating over all $i \in [k]$ and solving a corresponding system of linear equations.

6.1 Attack based on combinatorial properties of μ_{sk}

Lemma 6.2 (Attack based on combinatorial properties of μ_{sk}). *Let $\Pi \in \{\Pi_{\text{AleK}}, \Pi_{\text{Reg}}, \Pi_{\text{GPV}}\}$ be with $q = 2$ and $\text{Adv}^{\text{Dec}}(n) \geq 1 - \epsilon$ and suppose that the distribution μ_{sk} is (n, k, ρ) -sparse. Then there exists a non-uniform attack \mathcal{A}_{sk} on Π running in time $(k/\epsilon) \cdot \text{poly}(m)$ with $\text{Adv}^{\mathcal{A}_{\text{sk}}}(n) \geq (\rho - 4\sqrt{\epsilon})/10$.*

Proof. By Claim 4.4 it suffices to prove the lemma for $\Pi = \Pi_{\text{AleK}}$ and without loss of generality we may assume that g is the identity function over \mathbb{F}_2 . Since μ_{sk} is (n, k, ρ) -sparse there exist k subsets $A_1, \dots, A_k \subseteq \mathbb{F}_2^m$ and k full rank linear transformations $L_1, \dots, L_k : \mathbb{F}_2^m \rightarrow \mathbb{F}_2^n$ such that $\Pr_{\mu_{\text{sk}}} \left(\bigcup_{i=1}^k A_i \right) \geq \rho$ and $L_i(A_i)$ is constant for every $i \in [k]$.

Our main observation is that if e' satisfies that $b = w' + e'$ for some $w' \in \text{Image}(G)$ and in addition

$$\Pr[\langle e' \circ (-1), \mu_1 \rangle = 1] - \Pr[\langle e' \circ (-1), \mu_0 \rangle = 1] \geq 1 - \epsilon' \quad (8)$$

then decrypting the ciphertext using e' as the private key achieves advantage $1 - \epsilon'$. We search for e' that satisfies the above by enumerating over all $i \in [k]$ and solving a corresponding system of linear equations and we test whether e' satisfies (8) via sampling.

Fix $y \in \mathbb{F}_2^m$. By the Hoeffding bound for sampling (Lemma 3.1) if we draw $\ell = m/(\sqrt{\epsilon}/2)^2$ independent random samples $e_1^{(0)}, \dots, e_\ell^{(0)} \sim \mu_0$ and ℓ independent random samples $e_1^{(1)}, \dots, e_\ell^{(1)} \sim \mu_1$ then

$$\left| \left(\Pr[\langle y \circ (-1), \mu_1 \rangle = 1] - \Pr[\langle y \circ (-1), \mu_0 \rangle = 1] \right) - \left(\Pr_{j \in [\ell]} [\langle y \circ (-1), e_j^{(1)} \rangle = 1] - \Pr_{j \in [\ell]} [\langle y \circ (-1), e_j^{(0)} \rangle = 1] \right) \right| \leq \sqrt{\epsilon} \quad (9)$$

with probability at least $1 - 4 \cdot 2^{-2m}$. By union bound this implies in turn that (9) holds for every $y \in \mathbb{F}_2^m$ with probability at least $1 - 4 \cdot 2^{-m}$. In particular, there exist ℓ vectors $e_1^{(0)}, \dots, e_\ell^{(0)} \in \text{supp}(\mu_0)$ and ℓ vectors $e_1^{(1)}, \dots, e_\ell^{(1)} \in \text{supp}(\mu_1)$ for which (9) holds for every $y \in \mathbb{F}_2^m$.

\mathcal{A}_{sk}

- For every $i = 1, 2, \dots, k$:
 - Solve the system of linear equations

$$L_i b = L_i G x' + L_i(A_i)$$
 in the indeterminate x' .
 - If there is no solution continue to the next i .
 - Else let x' be an arbitrary solution and let $e' := b - Gx'$.
 - If e' satisfies that

$$\Pr_{j \in [\ell]} [\langle e' \circ (-1), e_j^{(1)} \rangle = 1] - \Pr_{j \in [\ell]} [\langle e' \circ (-1), e_j^{(0)} \rangle = 1] \geq 1 - 2\sqrt{\epsilon}, \quad (10)$$
 output $\langle e' \circ (-1), c \rangle$, else continue to the next i .
 - Else if no e' satisfies (10), output a random bit.

Inspection reveals that the attack above can be implemented using a non-uniform circuit of size $(k/\epsilon) \cdot \text{poly}(m)$. Next we analyze the advantage of the attack in guessing the message bit σ . We will show that with probability at least $(\rho - \sqrt{\epsilon})/10$ the attack finds e' which satisfies (10) and that in this case the advantage of guessing the correct message bit is at least $1 - 3\sqrt{\epsilon}$. This will imply in turn that $\text{Adv}^{\mathcal{A}_{\text{sk}}}(n) \geq (\rho - 4\sqrt{\epsilon})/10$.

We start by showing a lower bound on the probability that the attack finds e' which satisfies (10). Since $\text{Adv}^{\text{Dec}}(n) \geq 1 - \epsilon$ by Claim 4.2, together with a standard probabilistic argument, we have that e satisfies

$$\Pr[\langle e \circ (-1), \mu_1 \rangle = 1] - \Pr[\langle e \circ (-1), \mu_0 \rangle = 1] \geq 1 - \sqrt{\epsilon}$$

with probability at least $1 - \sqrt{\epsilon}$. By (9) this implies in turn that e satisfies (10) with probability at least $1 - \sqrt{\epsilon}$. Furthermore, since μ_{sk} is (n, k, ρ) -sparse with probability at least ρ we have that $e \in \bigcup_{i=1}^k A_i$. So with probability at least $\rho - \sqrt{\epsilon}$ we have that e satisfies (10) and in addition $e \in A_i$ for some $i \in [k]$. Finally, the matrix $L_i G$ is non-singular with probability at least $1/10$ (say), independently of the above.

Conditioned on all the above, in the i -th iteration we have that

$$e' = b - Gx' = b - G \cdot (L_i G)^{-1} (L_i b - L_i(A_i)) = b - G \cdot (L_i G)^{-1} \cdot (L_i b - L_i e) = e.$$

Consequently we have that the attack finds e' which satisfies (10) with probability at least $(\rho - \sqrt{\epsilon})/10$.

Next we show that if the attack finds e' which satisfies (10) then the advantage of guessing the correct message bit using e' is high. Since $e' = b - Gx'$ we have that

$$\begin{aligned} \langle e' \circ (-1), c \rangle &= \langle (b - Gx') \circ (-1), \tilde{w} \rangle + \langle (b - Gx') \circ (-1), \tilde{e} \rangle \\ &= \langle b \circ (-1), \tilde{w} \rangle - \langle x', G^T \cdot P_m(\tilde{w}) \rangle + \langle (b - Gx') \circ (-1), \tilde{e} \rangle \\ &= 0 - 0 + \langle e' \circ (-1), \tilde{e} \rangle \\ &= \langle e' \circ (-1), \tilde{e} \rangle. \end{aligned}$$

Furthermore, since e' satisfies (10), by (9) we have that

$$\Pr[\langle e' \circ (-1), \mu_1 \rangle = 1] - \Pr[\langle e' \circ (-1), \mu_0 \rangle = 1] \geq 1 - 3\sqrt{\epsilon},$$

so the advantage of the attack in this case is $1 - 3\sqrt{\epsilon}$. \square

6.2 Attack based on combinatorial properties of μ_0, μ_1

Lemma 6.3 (Attack based on combinatorial properties of μ_0, μ_1). *Let $\Pi \in \{\Pi_{\text{Alek}}, \Pi_{\text{Reg}}, \Pi_{\text{GPV}}\}$ be with $q = 2$ and $\text{Adv}^{\text{Dec}}(n) \geq 1 - \epsilon$ and suppose that there exists $\xi \in \{0, 1\}$ such that the distribution μ_ξ is $(m + 1 - n + r, k, \rho)$ -sparse. Then there exists a non-uniform attack \mathcal{A}_ξ on Π running in time $k \cdot \text{poly}(m)$ with $\text{Adv}^{\mathcal{A}_\xi}(n) \geq \rho/2 - \epsilon - 2k2^{-r}$.*

Proof. By Claim 4.4 it suffices to prove the lemma for $\Pi = \Pi_{\text{GPV}}$ and by symmetry we may further assume that $\xi = 0$. Without loss of generality we may assume that g is the identity function over \mathbb{F}_2 . Since μ_0 is $(m + 1 - n + r, k, \rho)$ -sparse there exist k subsets $A_1, \dots, A_k \subseteq \mathbb{F}_2^{m+1}$ and k full rank linear transformations $L_1, \dots, L_k : \mathbb{F}_2^{m+1} \rightarrow \mathbb{F}_2^{m+1-n+r}$ such that $\Pr_{\mu_0}(\bigcup_{i=1}^k A_i) \geq \rho$ and $L_i(A_i)$ is constant for every $i \in [k]$. For every $i \in [k]$ let $V_i = \{v \in \mathbb{F}_2^{m+1} \mid L_i(v) = L_i(A_i)\}$ and let $S = \bigcup_{i=1}^k V_i$. Since $\text{Adv}^{\text{Dec}}(n) \geq 1 - \epsilon$, by averaging there exists $e^{(\text{sk})} \in \text{supp}(\mu_{\text{sk}})$ such that

$$\Pr[\langle e^{(\text{sk})} \circ (-1), \mu_1 \rangle = 1] - \Pr[\langle e^{(\text{sk})} \circ (-1), \mu_0 \rangle = 1] \geq 1 - \epsilon.$$

Our main observation is that since S is not too large, with high probability over the choice of the matrix \tilde{H} , there is no $e' \in S \setminus \{\tilde{e}\}$ such that $c = \tilde{H}x' + e'$ for some $x' \in \mathbb{F}_2^{m-n}$. This implies in turn that by enumerating over all $i \in [k]$ and solving a corresponding system of linear equations, with high probability one can verify whether $\tilde{e} \in S$ and if this is the case one can also find \tilde{e} . It thus suffices to be able to distinguish between $\tilde{e} \sim \mu_0$ and $\tilde{e} \sim \mu_1$, conditioned on the event that $\tilde{e} \in S$. Assuming that ϵ is sufficiently small compared to ρ , this can be done by computing the inner product $\langle e^{(\text{sk})} \circ (-1), \tilde{e} \rangle$.

\mathcal{A}_0

- For every $i = 1, 2, \dots, k$:

- Solve the system of linear equations

$$L_i c = L_i \tilde{H} x' + L_i(A_i) \quad (11)$$

in the indeterminate x' .

- If there is no solution continue to the next i .

- Else let x' be an arbitrary solution and let $e' := c - \tilde{H} x'$.

- If e' satisfies that $\langle e^{(\text{sk})} \circ (-1), e' \rangle = 0$ output 0, else continue to the next i .

- Else if no e' satisfies the above, output a random bit.

Inspection reveals that the attack above can be implemented using a non-uniform circuit of size $k \cdot \text{poly}(m)$. Next we analyze the advantage of the attack in guessing the message bit σ .

We say that \tilde{H} is S -good for \tilde{e} if there is no $z \in S \setminus \{\tilde{e}\}$ such that $z - \tilde{e} \in \text{Image}(\tilde{H})$. We will show that for every \tilde{e} the probability that \tilde{H} is S -good for \tilde{e} is at least $1 - k \cdot 2^{-r}$. Consequently, for every \tilde{e} there exists a collection $\mathcal{H}_{\tilde{e}}$ of S -good matrices for \tilde{e} such that $\Pr[\tilde{H} \in \mathcal{H}_{\tilde{e}}] = 1 - k \cdot 2^{-r}$. We will then show that conditioned on the event that $\tilde{H} \in \mathcal{H}_{\tilde{e}}$ the attack outputs 0 with probability at least $(1 + \rho - \epsilon)/2$ when c is an encryption of 0 and it outputs 1 with probability at least $(1 - \epsilon)/2$ when c is an encryption of 1. This will imply in turn that the advantage of the attack is at least $(1 - k2^{-r})(\rho/2 - \epsilon) - k2^{-r} \geq \rho/2 - \epsilon - 2k2^{-r}$.

We start by showing that for every \tilde{e} the probability that \tilde{H} is S -good for \tilde{e} is at least $1 - k \cdot 2^{-r}$. For this note that for every $i \in [k]$ the subspace V_i has co-dimension $m + 1 - n + r$ and hence $|V_i| = 2^{n-r}$ and consequently $|S| \leq k2^{n-r}$. Thus by union bound it suffices to show that for every $z \in S \setminus \{\tilde{e}\}$ it holds that $z - \tilde{e} \in \text{Image}(\tilde{H})$ with probability at most 2^{-n} . To see this fix $z \in S \setminus \{\tilde{e}\}$ and suppose that $z - \tilde{e} \in \text{Image}(\tilde{H})$. Since $\tilde{H} = \begin{pmatrix} H \\ u^T \end{pmatrix}$ this implies in turn that $P_m(z - \tilde{e}) \in \text{Image}(H)$. Furthermore, since $z - \tilde{e} \neq 0$ and H is full rank we also have that $P_m(z - \tilde{e}) \neq 0$. So we obtained that $P_m(z - \tilde{e})$ is a non-zero point contained in $\text{Image}(H)$, a uniform random $(m - n)$ -dimensional space, which happens with probability at most 2^{-n} .

Next we show a lower bound on the probability that the attack outputs 0 when c is an encryption of 0, conditioned on the event that $\tilde{H} \in \mathcal{H}_{\tilde{e}}$. Since the event $\tilde{H} \in \mathcal{H}_{\tilde{e}}$ is independent of the choice of \tilde{e} , by union bound we have that the events $\tilde{e} \in \bigcup_{i=1}^k A_i$ and $\langle e^{(\text{sk})} \circ (-1), \tilde{e} \rangle = 0$ hold simultaneously with probability at least $\rho - \epsilon$. We will show that if these two events hold then the attack outputs 0. This will imply in turn that in the case in which c is an encryption of 0, conditioned on the event that $\tilde{H} \in \mathcal{H}_{\tilde{e}}$, the attack outputs 0 with probability at least $\rho - \epsilon$ and it outputs a random bit otherwise. So it outputs 0 in this case with probability at least $(1 + \rho - \epsilon)/2$.

Suppose that $\tilde{e} \in \bigcup_{i=1}^k A_i$ and $\langle e^{(\text{sk})} \circ (-1), \tilde{e} \rangle = 0$. Then in this case we have that

$$L_i c = L_i \tilde{w} + L_i \tilde{e} = L_i \tilde{H} \tilde{x} + L_i(A_i),$$

where $i \in [k]$ is such that $\tilde{e} \in A_i$ and \tilde{x} is such that $\tilde{w} = \tilde{H} \tilde{x}$. Consequently, the attack will find a solution for (11). Furthermore, we claim that if the attack finds a solution x' to (11) for some $j \in [k]$ then $e' = c - \tilde{H} x' = \tilde{e}$. To see this note that $L_j e' = L_j c - L_j \tilde{H} x' = L_j(A_j)$ and therefore $e' \in S$. Furthermore,

we have that $e' - \tilde{e} = (c - \tilde{H}x') - (c - \tilde{H}\tilde{x}) = \tilde{H}(\tilde{x} - x')$ and so $e' - \tilde{e} \in \text{Image}(\tilde{H})$. But due to our assumption that \tilde{H} is S -good for \tilde{e} this implies in turn that $\tilde{e} = e'$. So we have that $\tilde{e} = e'$ and due to our assumption that $\langle e^{(\text{sk})} \circ (-1), \tilde{e} \rangle = 0$ this implies in turn that the attack will output 0.

Finally, we show a lower bound on the probability that the attack outputs 1 when c is an encryption of 1, conditioned on the event that $\tilde{H} \in \mathcal{H}_{\tilde{e}}$. Since the event $\tilde{H} \in \mathcal{H}_{\tilde{e}}$ is independent of the choice of \tilde{e} , we have that $\langle e^{(\text{sk})} \circ (-1), \tilde{e} \rangle = 1$ with probability at least $1 - \epsilon$. Suppose that this latter event holds. If there is no solution for (11) for every $j \in [k]$ the attack outputs a random bit. Otherwise if the attack finds a solution x' for (11) for some $j \in [k]$ then similarly to the above the assumption that $\tilde{H} \in \mathcal{H}_{\tilde{e}}$ implies that $e' = c - \tilde{H}x' = \tilde{e}$. Due to our assumption that $\langle e^{(\text{sk})} \circ (-1), \tilde{e} \rangle = 1$ this implies in turn that the attack will output a random bit. Concluding, we obtained that in the case in which c is an encryption of 1, conditioned on the event that $\tilde{H} \in \mathcal{H}_{\tilde{e}}$, the attack outputs 1 with probability at least $(1 - \epsilon)/2$. \square

7 Attacks based on the approximate duality conjecture

Recall the definition of the duality measure given in (1). All results presented in this section assume that the following conjecture holds.

Conjecture 7.1 (Approximate duality conjecture [7]). *For every constant $\epsilon > 0$ there exists a constant c which depends only on ϵ such that the following holds. If $A, B \subseteq \mathbb{F}_2^m$ have $D(A, B) \geq \epsilon$ then there exist subsets $A' \subseteq A$ and $B' \subseteq B$ such that $|A'| \geq 2^{-c\sqrt{m}}|A|$, $|B'| \geq 2^{-c\sqrt{m}}|B|$ and $D(A', B') = 1$.*

Our main result in this section is the following.

Theorem 7.2. *Assuming the approximate duality conjecture (Conjecture 7.1) there exist constants $\epsilon, \gamma > 0$ such that the following holds. Let $\Pi \in \{\Pi_{\text{AleK}}, \Pi_{\text{Reg}}, \Pi_{\text{GPV}}\}$ be with $q = 2$ and $\text{Adv}^{\text{Dec}}(n) \geq 1 - \epsilon$. Then there exists a non-uniform attack \mathcal{A} on Π running in time $2^{O(\sqrt{m})}$ with $\text{Adv}^{\mathcal{A}}(n) \geq \gamma$.*

For the proof of the above theorem we first prove two consequences of Conjecture 7.1. The first consequence is a generalized form of this conjecture that applies to arbitrary distributions, not necessarily uniform over subsets A, B . For a pair of distributions μ_1, μ_2 over \mathbb{F}_2^m we define their duality measure as

$$D(\mu_1, \mu_2) = \mathbb{E} \left[(-1)^{\langle \mu_1, \mu_2 \rangle} \right].$$

Note that in the special case where μ_1, μ_2 are uniform distributions over subsets $A, B \subseteq \mathbb{F}_2^m$ respectively then $D(\mu_1, \mu_2) = D(A, B)$.

Lemma 7.3. *Assuming Conjecture 7.1, for every constant $\epsilon > 0$ there exists a constant c which depends only on ϵ such that the following holds. If a pair of distributions μ_1, μ_2 over \mathbb{F}_2^m have $D(\mu_1, \mu_2) \geq \epsilon$ then there exist subsets $A', B' \subseteq \mathbb{F}_2^m$ such that $\Pr_{\mu_1}(A') \geq 2^{-c\sqrt{m}}$, $\Pr_{\mu_2}(B') \geq 2^{-c\sqrt{m}}$ and $D(A', B') = 1$.*

The proof of the above lemma is given in Section 7.1. Note that the probability of being contained in the sets A' and B' in the above lemma is $2^{-c\sqrt{m}}$ and so using this lemma one can only obtain an attack on the abstract encryption scheme in the case in which the decryption error of a single encryption is $2^{-\Omega(\sqrt{m})}$. However, we are interested in an attack that works in the case in which the decryption error of a single encryption is a sufficiently small constant. For this we apply Lemma 7.3 iteratively to obtain $t \approx 2^{c\sqrt{m}}$ pairs of subsets A_i, B_i such that $D(A_i, B_i) = 1$ for all $1 \leq i \leq t$ and such that the probability of being contained in the union of $\Omega(t)$ of these subsets is $\Omega(\epsilon)$.

Lemma 7.4. *Assuming Conjecture 7.1, for every constant $\epsilon > 0$ there exists a constant c which depends only on ϵ such that the following holds for every integer $t \leq 2^{c\sqrt{m}}\epsilon/4$. If a pair of distributions μ_1, μ_2 over \mathbb{F}_2^m have $D(\mu_1, \mu_2) \geq \epsilon$, then there exist subsets $A_1, \dots, A_t \subseteq \mathbb{F}_2^m$ and $B_1, \dots, B_t \subseteq \mathbb{F}_2^m$ such that $D(A_i, B_i) = 1$ for all $i \in [t]$, and in addition for every $I \subseteq [t]$ it holds that $\Pr_{\mu_1}(\bigcup_{i \in I} A_i) \geq |I| \cdot 2^{-c\sqrt{m}}/4$ and $\Pr_{\mu_2}(\bigcup_{i \in I} B_i) \geq |I| \cdot 2^{-c\sqrt{m}}/4$.*

Note that the sets A_1, \dots, A_t and B_1, \dots, B_t in the above lemma may have non-empty intersections and in particular are not required to be distinct. The proof of the above lemma is given in Section 7.2.

In what follows we present the proof of our main Theorem 7.2 based on Lemma 7.4.

Proof of Theorem 7.2. We will show that assuming Conjecture 7.1 we have that the conditions of either Lemma 6.2 or Lemma 6.3 hold. Let c be the constant guaranteed by Lemma 7.4 for the constant $1 - 2\epsilon$. We shall show that the conclusion of the theorem holds for

$$\gamma = \min \left\{ (1 - 4\sqrt{\epsilon})/10, ((1 - 2\epsilon)/32 - 4\sqrt{\epsilon})/10, (1 - 2\epsilon)/64 - \epsilon - 2 \cdot 2^{-c\sqrt{m}} \right\}.$$

If $n \leq 2c\sqrt{m}$ we clearly have that the distribution μ_{sk} is $(n, 2^{2c\sqrt{m}}, 1)$ -sparse and consequently Lemma 6.2 implies an attack in time $2^{O(\sqrt{m})}$ with advantage $(1 - 4\sqrt{\epsilon})/10$. Hence from now on we shall assume that $n > 2c\sqrt{m}$.

Let $\xi \in \{0, 1\}$ be such that the decryption function g satisfies $g(0) = \xi$. Our main observation is that the assumption that $\text{Adv}^{\text{Dec}}(n) \geq 1 - \epsilon$ implies that $\Pr[\langle \mu_{\text{sk}} \circ (-1), \mu_\xi \rangle = 0] \geq 1 - \epsilon$ and consequently $D(\mu_{\text{sk}} \circ (-1), \mu_\xi) \geq 1 - 2\epsilon$. Thus we may apply Lemma 7.4 to the distributions $\mu_{\text{sk}} \circ (-1)$ and μ_ξ and conclude the existence of $t = 2^{c\sqrt{m}}(1 - 2\epsilon)/4$ subsets $A_1, \dots, A_t \subseteq \mathbb{F}_2^{m+1}$ and $B_1, \dots, B_t \subseteq \mathbb{F}_2^{m+1}$ such that $D(A_i, B_i) = 1$ for all $i \in [t]$, and in addition for every $I \subseteq [t]$ it holds that $\Pr_{\mu_{\text{sk}} \circ (-1)}(\bigcup_{i \in I} A_i) \geq |I| \cdot 2^{-c\sqrt{m}}/4$ and $\Pr_{\mu_\xi}(\bigcup_{i \in I} B_i) \geq |I| \cdot 2^{-c\sqrt{m}}/4$.

Fix $i \in [t]$. The fact that $D(A_i, B_i) = 1$ implies in turn that $\dim(\text{span}(A_i)) + \dim(\text{span}(B_i)) \leq m + 2$ and in particular we have that either $\dim(\text{span}(A_i)) \leq m + 2 - n + 2c\sqrt{m}$ or $\dim(\text{span}(B_i)) \leq n - 2c\sqrt{m}$. Let $I \subseteq [t]$ be the set of all indices i for which $\dim(\text{span}(A_i)) \leq m + 2 - n + 2c\sqrt{m}$. We shall show that if $|I| \geq t/2$ the conditions of Lemma 6.2 hold while if $|I| < t/2$ the conditions of Lemma 6.3 hold.

We start with the case in which $|I| \geq t/2$. Fix $i \in I$ and let v_1, \dots, v_{m+1} be a basis for \mathbb{F}_2^{m+1} such that the subspace spanned by $v_1, \dots, v_{m+2-n+2c\sqrt{m}}$ contains $\text{span}(A_i)$. Let $L_i : \mathbb{F}_2^{m+1} \rightarrow \mathbb{F}_2^n$ be the linear transformation which satisfies $L_i(\sum_{j=1}^{m+1} \alpha_j v_j) = (\alpha_{m-n+2}, \dots, \alpha_{m+1})$ for every $\alpha_1, \dots, \alpha_{m+1} \in \mathbb{F}_2$. Then $L_i(A_i)$ is supported only on the first $2c\sqrt{m} + 1$ bits and consequently $|L_i(A_i)| \leq 2^{2c\sqrt{m}+1}$. Furthermore, we have that $|I| \leq t = 2^{c\sqrt{m}}(1 - 2\epsilon)/4$ and $\Pr_{\mu_{\text{sk}} \circ (-1)}(\bigcup_{i \in I} A_i) \geq (t/2) \cdot 2^{-c\sqrt{m}}/4 = (1 - 2\epsilon)/32$. This implies in turn that the distribution $\mu_{\text{sk}} \circ (-1)$, and consequently also μ_{sk} , are $(n, 2^{3c\sqrt{m}+1}(1 - 2\epsilon)/4, (1 - 2\epsilon)/32)$ -sparse. Lemma 6.2 implies in turn that the encryption scheme can be attacked in time $2^{O(\sqrt{m})}$ with advantage $((1 - 2\epsilon)/32 - 4\sqrt{\epsilon})/10$.

Next we deal with the case in which $|I| < t/2$. Similarly to the previous case for every $i \notin I$ there exists a full rank linear transformation $L_i : \mathbb{F}_2^{m+1} \rightarrow \mathbb{F}_2^{m+1-n+2c\sqrt{m}}$ such that $L_i(B_i) \equiv 0$ and $\Pr_{\mu_\xi}(\bigcup_{i \notin I} B_i) \geq (1 - 2\epsilon)/32$. This implies in turn that μ_ξ is $(m + 1 - n + 2c\sqrt{m}, 2^{c\sqrt{m}}(1 - 2\epsilon)/4, (1 - 2\epsilon)/32)$ -sparse. So by Lemma 6.3 we have that the encryption scheme can be attacked in time $2^{O(\sqrt{m})}$ with advantage $(1 - 2\epsilon)/64 - \epsilon - 2 \cdot 2^{-c\sqrt{m}}$. \square

7.1 From uniform to general distributions – proof of Lemma 7.3

We start with the following lemma which says that every distribution can be approximated by a distribution which is a convex combination of not too many uniform distributions.

Lemma 7.5. Let μ be a distribution with support S , $|S| = N$, and let $t = \log(2N/\epsilon)/\log(1 + \epsilon/2)$. Then there exist a partition of S into at most $t + 2$ subsets S_0, \dots, S_{t+1} and a distribution χ which is a convex combination of uniform distributions on S_0, \dots, S_t such that μ is ϵ -close to χ .

Proof. Choose an arbitrary element $\beta \in S$. Let

$$S_0 = \left\{ \alpha \in S \setminus \{\beta\} \mid \Pr_\mu(\alpha) \leq \frac{\epsilon}{2N} \right\},$$

for all $1 \leq i \leq t$ let

$$S_i = \left\{ \alpha \in S \setminus \{\beta\} \mid \frac{\epsilon}{2N} \cdot (1 + \epsilon/2)^{i-1} < \Pr_\mu(\alpha) \leq \frac{\epsilon}{2N} \cdot (1 + \epsilon/2)^i \right\}$$

and let $S_{t+1} = \{\beta\}$.

Let χ be the distribution which satisfies

$$\Pr_\chi(\alpha) = \begin{cases} 0, & \alpha \in S_0 \\ \frac{\epsilon}{2N} \cdot (1 + \epsilon/2)^{i-1}, & \alpha \in S_i \text{ for } 1 \leq i \leq t \\ 1 - \sum_{\gamma \in S \setminus \{\beta\}} \Pr_\chi(\gamma), & \alpha = \beta. \end{cases}$$

We clearly have that S_0, \dots, S_{t+1} is a partition of S and that χ is a convex combination of uniform distributions on S_0, \dots, S_t .

It remains to show that μ is ϵ -close to the distribution χ . For this we compute

$$\begin{aligned} |\mu - \chi| &= \frac{1}{2} \sum_{\alpha \in S} |\Pr_\mu(\alpha) - \Pr_\chi(\alpha)| \\ &= \sum_{\alpha \in S \setminus \{\beta\}} (\Pr_\mu(\alpha) - \Pr_\chi(\alpha)) \\ &\leq \sum_{\alpha \in S_0} \frac{\epsilon}{2N} + \sum_{i=1}^t \sum_{\alpha \in S_i} \frac{\epsilon}{2} \Pr_\mu(\alpha) \\ &\leq \frac{\epsilon}{2N} \cdot N + \frac{\epsilon}{2} \sum_{\alpha \in S} \Pr_\mu(\alpha) \\ &= \epsilon. \end{aligned}$$

□

We shall also use the definition of the *spectrum* given below.

Definition 7.6 (Spectrum). For a distribution μ over \mathbb{F}_2^m and $\epsilon \in [0, 1]$ let the ϵ -spectrum of μ be the set

$$\text{Spec}_\epsilon(\mu) = \left\{ x \in \mathbb{F}_2^m \mid \mathbb{E}[(-1)^{\langle x, \mu \rangle}] \geq \epsilon \right\}. \quad (12)$$

Note that if $\text{supp}(\mu_1) \subseteq \text{Spec}_\epsilon(\mu_2)$ then $D(\mu_1, \mu_2) \geq \epsilon$. Conversely, a standard probabilistic argument shows that if $D(\mu_1, \mu_2) \geq \epsilon$ then $\Pr_{\mu_1}(\text{Spec}_{\epsilon/2}(\mu_2)) \geq \epsilon/2$.

Proof of Lemma 7.3. Let c' be the constant guaranteed by Conjecture 7.1 for the constant $\epsilon/4$.

Let $\mu'_1 = \mu_1|_{\text{Spec}_{\epsilon/2}(\mu_2)}$ and note that the fact that $D(\mu_1, \mu_2) \geq \epsilon$ implies that $\Pr_{\mu_1}(\text{Spec}_{\epsilon/2}(\mu_2)) \geq \epsilon/2$. By Lemma 7.5 there exists a partition of $\text{supp}(\mu'_1)$ into $t+2$ subsets $A_0, \dots, A_{t+1} \subseteq \mathbb{F}_2^m$ for $t = \log(2 \cdot 2^m/\delta)/\log(1 + \delta/2)$ such that μ'_1 is δ -close to a distribution χ_1 which is a convex combination of uniform distributions on A_0, \dots, A_{t+1} . Since $\text{supp}(\mu'_1) \subseteq \text{Spec}_{\epsilon/2}(\mu_2)$ we have that $A_i \subseteq \text{Spec}_{\epsilon/2}(\mu_2)$ for all $0 \leq i \leq t+1$ and so $D(A_i, \mu_2) \geq \epsilon/2$ for all $0 \leq i \leq t+1$.

Fix $0 \leq i \leq t+1$. Similarly to the above, let $\mu_2^{(i)} = \mu_2|_{\text{Spec}_{\epsilon/4}(A_i)}$ and note that the fact that $D(A_i, \mu_2) \geq \epsilon/2$ implies that $\Pr_{\mu_2}(\text{Spec}_{\epsilon/4}(A_i)) \geq \epsilon/4$. By Lemma 7.5 there exists a partition of $\text{supp}(\mu_2^{(i)})$ into $t+2$ subsets $B_0^{(i)}, \dots, B_{t+1}^{(i)} \subseteq \mathbb{F}_2^m$ for $t = \log(2 \cdot 2^m/\delta)/\log(1 + \delta/2)$ such that $\mu_2^{(i)}$ is δ -close to a distribution $\chi_2^{(i)}$ which is a convex combination of uniform distributions on $B_0^{(i)}, \dots, B_{t+1}^{(i)}$. Since $\text{supp}(\mu_2^{(i)}) \subseteq \text{Spec}_{\epsilon/4}(A_i)$ we have that $B_j^{(i)} \subseteq \text{Spec}_{\epsilon/4}(A_i)$ for all $0 \leq j \leq t+1$ and so $D(A_i, B_j^{(i)}) \geq \epsilon/4$ for all $0 \leq j \leq t+1$.

Summarizing, so far we found a collection of subsets $\{A_i\}_{0 \leq i \leq t+1}$ and a collection $\{B_j^{(i)}\}_{0 \leq i, j \leq t+1}$ such that:

- $\mu'_1 = \mu_1|_{\text{Spec}_{\epsilon/2}(\mu_2)}$ is close to a convex combination of uniform distributions on A_0, \dots, A_{t+1} .
- $\mu_2^{(i)} = \mu_2|_{\text{Spec}_{\epsilon/4}(A_i)}$ is close to a convex combination of uniform distributions on $B_0^{(i)}, \dots, B_{t+1}^{(i)}$ for all $0 \leq i \leq t+1$.
- $D(A_i, B_j^{(i)}) \geq \epsilon/4$ for all $0 \leq i, j \leq t+1$.

For every $0 \leq i, j \leq t+1$ we can apply Conjecture 7.1 to the sets $A_i, B_j^{(i)}$ and conclude the existence of subsets $\tilde{A}_j^{(i)} \subseteq A_i, \tilde{B}_j^{(i)} \subseteq B_j^{(i)}$ such that $D(\tilde{A}_j^{(i)}, \tilde{B}_j^{(i)}) = 1$ and $|\tilde{A}_j^{(i)}| \geq 2^{-c'\sqrt{m}}|A_i|, |\tilde{B}_j^{(i)}| \geq 2^{-c'\sqrt{m}}|B_j^{(i)}|$. So in order to prove the lemma it suffices to show the existence of a constant c and indices $0 \leq k, \ell \leq t+1$ for which $\Pr_{\mu_1}(\tilde{A}_\ell^{(k)}) \geq 2^{-c\sqrt{m}}$ and $\Pr_{\mu_2}(\tilde{B}_\ell^{(k)}) \geq 2^{-c\sqrt{m}}$.

By the pigeonhole principle, for every $0 \leq i \leq t+1$ there exists an index $0 \leq j_i \leq t+1$ such that

$$\Pr_{\mu_2}(\tilde{B}_{j_i}^{(i)}) \geq \frac{\Pr_{\mu_2}(\bigcup_{j=0}^{t+1} \tilde{B}_j^{(i)})}{t+2}.$$

Similarly, there exists $0 \leq k \leq t+1$ such that

$$\Pr_{\mu_1}(\tilde{A}_{j_k}^{(k)}) \geq \frac{\Pr_{\mu_1}(\bigcup_{i=0}^{t+1} \tilde{A}_{j_i}^{(i)})}{t+2}.$$

Let $A' = \tilde{A}_{j_k}^{(k)}$ and $B' = \tilde{B}_{j_k}^{(k)}$. Then we have that $D(A', B') = 1$ and in order to bound the probabilities $\Pr_{\mu_1}(A')$ and $\Pr_{\mu_2}(B')$ from below it suffices to bound the probabilities $\Pr_{\mu_2}(\bigcup_{j=0}^{t+1} \tilde{B}_j^{(k)})$ and

$\Pr_{\mu_1} \left(\bigcup_{i=0}^{t+1} \tilde{A}_{j_i}^{(i)} \right)$ from below. For this we compute

$$\begin{aligned} \Pr_{\mu_2} \left(\bigcup_{j=0}^{t+1} \tilde{B}_j^{(k)} \right) &\geq \frac{\epsilon}{4} \cdot \Pr_{\mu_2^{(k)}} \left(\bigcup_{j=0}^{t+1} \tilde{B}_j^{(k)} \right) \quad (\text{Since } \Pr_{\mu_2}(\text{Spec}_{\epsilon/4}(A_k)) \geq \epsilon/4) \\ &\geq \frac{\epsilon}{4} \cdot \left(\Pr_{\chi_2^{(k)}} \left(\bigcup_{j=0}^{t+1} \tilde{B}_j^{(k)} \right) - \delta \right) \quad (\text{Since } \mu_2^{(k)} \text{ and } \chi_2^{(k)} \text{ are } \delta\text{-close}) \\ &\geq \frac{\epsilon}{4} \cdot (2^{-c'\sqrt{m}} - \delta), \end{aligned}$$

where the last inequality follows since $\chi_2^{(k)}$ is a convex combination of uniform distributions on $B_0^{(k)}, \dots, B_{t+1}^{(k)}$ and $|\tilde{B}_j^{(k)}| \geq 2^{-c'\sqrt{m}} |B_j^{(k)}|$ for all $0 \leq j \leq t+1$.

Similarly, we have that

$$\begin{aligned} \Pr_{\mu_1} \left(\bigcup_{i=0}^{t+1} \tilde{A}_{j_i}^{(i)} \right) &\geq \frac{\epsilon}{2} \cdot \Pr_{\mu_1'} \left(\bigcup_{i=0}^{t+1} \tilde{A}_{j_i}^{(i)} \right) \\ &\geq \frac{\epsilon}{2} \cdot \left(\Pr_{\chi_1'} \left(\bigcup_{i=0}^{t+1} \tilde{A}_{j_i}^{(i)} \right) - \delta \right) \\ &\geq \frac{\epsilon}{2} \cdot (2^{-c'\sqrt{m}} - \delta). \end{aligned}$$

Concluding, we have found subsets A', B' such that $D(A', B') = 1$ and such that both $\Pr_{\mu_1}(A')$ and $\Pr_{\mu_2}(B')$ are bounded from below by $\frac{\epsilon}{4(t+2)} \cdot (2^{-c'\sqrt{m}} - \delta)$. The proof is completed by letting $\delta = 2^{-c'\sqrt{m}}/2$ and $t = \frac{\log(2 \cdot 2^m / \delta)}{\log(1 + \delta/2)}$ and noting that with this setting of parameters there exists a constant c which depends only on ϵ such that $\frac{\epsilon}{4(t+2)} \cdot (2^{-c'\sqrt{m}} - \delta) \geq 2^{-c\sqrt{m}}$ for a sufficiently large m . \square

7.2 Iterative application – proof of Lemma 7.4

In what follows we prove Lemma 7.4 based on Lemma 7.3.

Proof of Lemma 7.4. Let c be the constant guaranteed by Lemma 7.3 for the constant $\epsilon/2$. For ease of notations let $\gamma := 2^{-c\sqrt{m}}$.

We shall apply Lemma 7.3 iteratively to obtain at each step a new pair of sets A_i, B_i . Let $\mu_1^{(1)} := \mu_1$, $\mu_2^{(1)} := \mu_2$. For $i = 1, 2, \dots$, as long as $D(\mu_1^{(i)}, \mu_2^{(i)}) \geq \epsilon/2$ apply Lemma 7.3 to the distributions $\mu_1^{(i)}$ and $\mu_2^{(i)}$ to obtain minimal subsets A_i, B_i such that $D(A_i, B_i) = 1$ and $\Pr_{\mu_1^{(i)}}(A_i) \geq \gamma$, $\Pr_{\mu_2^{(i)}}(B_i) \geq \gamma$.

By minimality of A_i, B_i there exist elements $a_i \in A_i$ and $b_i \in B_i$ such that $\Pr_{\mu_1^{(i)}}(A_i \setminus \{a_i\}) < \gamma$ and $\Pr_{\mu_2^{(i)}}(B_i \setminus \{b_i\}) < \gamma$. Continue to the next step with the distributions $\mu_1^{(i+1)}, \mu_2^{(i+1)}$ given by

$$\Pr_{\mu_1^{(i+1)}}(a) = \begin{cases} 0, & a \in A_i \setminus \{a_i\} \\ \frac{\Pr_{\mu_1^{(i)}}(A_i) - \gamma}{1 - \gamma}, & a = a_i \\ \frac{\Pr_{\mu_1^{(i)}}(a)}{1 - \gamma}, & \text{Otherwise,} \end{cases}$$

and

$$\Pr_{\mu_2}^{(i+1)}(b) = \begin{cases} 0, & b \in B_i \setminus \{b_i\} \\ \frac{\Pr_{\mu_2}^{(i)}(B_i) - \gamma}{1 - \gamma}, & b = b_i \\ \frac{\Pr_{\mu_2}^{(i)}(b)}{1 - \gamma}, & \text{Otherwise.} \end{cases}$$

Note that at the i -th step we have

$$D(\mu_1^{(i+1)}, \mu_2^{(i+1)}) > D(\mu_1^{(i)}, \mu_2^{(i)}) - 2\gamma,$$

and so

$$D(\mu_1^{(i)}, \mu_2^{(i)}) > D(\mu_1^{(1)}, \mu_2^{(1)}) - 2(i-1)\gamma \geq \epsilon - 2(i-1)\gamma.$$

Thus the process continues as long as $i \leq \epsilon/(4\gamma)$.

It remains to bound the probabilities $\Pr_{\mu_1}(\bigcup_{i \in I} A_i)$ and $\Pr_{\mu_2}(\bigcup_{i \in I} B_i)$ from below for $I \subseteq [\epsilon/(4\gamma)]$, and by symmetry it suffices to bound $\Pr_{\mu_1}(\bigcup_{i \in I} A_i)$ from below. So we compute

$$\Pr_{\mu_1}\left(\bigcup_{i \in I} A_i\right) \geq \sum_{i \in I} (1 - \gamma)^{i-1} \cdot \gamma \geq |I|\gamma(1 - \gamma)^{1/\gamma} \geq |I|\gamma/4,$$

where the last inequality holds for a sufficiently large m recalling that $\gamma = 2^{-c\sqrt{m}}$. □

8 Perfectly correct decryption

We end this paper by discussing the possibility of achieving perfectly correct decryption in the abstract encryption scheme over constant-size rings. We first prove that when $q = 2$ any instance of the abstract encryption scheme with perfectly correct decryption can be attacked in time $\text{poly}(m)$. Then we propose to use the construction of matching vector families from [13], which builds on the constructions of [20, 15], to achieve an instance of the abstract encryption scheme with perfectly correct decryption over constant-size rings that cannot be attacked in time $\text{poly}(m)$.

8.1 Insecurity over the binary field

Claim 8.1. *Let $\Pi \in \{\Pi_{\text{Alek}}, \Pi_{\text{Reg}}, \Pi_{\text{GPV}}\}$ be with $q = 2$ and $\text{Adv}^{\text{Dec}}(n) = 1$. Then there exists a non-uniform attack \mathcal{A}_{per} on Π running in time $\text{poly}(m)$ with $\text{Adv}^{\mathcal{A}_{\text{per}}}(n) = 1$.*

Proof. By Claim 4.4 it suffices to prove the claim for $\Pi = \Pi_{\text{Alek}}$ and without loss of generality we may assume that the decryption function g is the identity function over \mathbb{F}_2 . By Claim 4.2 and the assumption that $\text{Adv}^{\text{Dec}}(n) = 1$ we have that $\langle \mu_{\text{sk}} \circ (-1), \mu_\xi \rangle \equiv \xi$ for every $\xi \in \{0, 1\}$. Let \hat{e} be an arbitrary element in $\text{supp}(\mu_{\text{sk}})$ and note that $\hat{e} \circ (-1) + \text{supp}(\mu_{\text{sk}} \circ (-1)) \subseteq (\text{supp}(\mu_\xi))^\perp$ for every $\xi \in \{0, 1\}$.

Our main observation is that if e' satisfies that $b = w' + e'$ for some $w' \in \text{Image}(G)$ and in addition $\hat{e} \circ (-1) + e' \circ (-1) \in (\text{supp}(\mu_\xi))^\perp$ for every $\xi \in \{0, 1\}$ then decrypting the ciphertext using e' as the private key achieves advantage 1. We find e' which satisfies the above by solving a system of linear equations.

Fix $\xi \in \{0, 1\}$. Let $d_\xi = \dim((\text{supp}(\mu_\xi))^\perp)$ and let $v_1^{(\xi)}, \dots, v_{m+1}^{(\xi)}$ be a basis for \mathbb{F}_2^{m+1} such that $v_1^{(\xi)}, \dots, v_{d_\xi}^{(\xi)}$ is a basis for $(\text{supp}(\mu_\xi))^\perp$. Let $L_\xi : \mathbb{F}_2^{m+1} \rightarrow \mathbb{F}_2^{m+1-d_\xi}$ be the linear transformation which

satisfies

$$L_\xi\left(\sum_{i=1}^{m+1} \alpha_i v_i^{(\xi)}\right) = (\alpha_{d_\xi+1}, \dots, \alpha_{m+1})$$

for every $\alpha_1, \dots, \alpha_{m+1} \in \mathbb{F}_2$. Note that $L_\xi(\mu_{\text{sk}} \circ (-1)) \equiv L_\xi(\hat{e} \circ (-1))$. Let $b' = b \circ (-1)$ and $G' = \begin{pmatrix} G \\ 0_n^T \end{pmatrix}$.

\mathcal{A}_{per}

- Let x' be an arbitrary joint solution of the systems of linear equations

$$L_0 b' = L_0 G' x' + L_0(\hat{e} \circ (-1)) \tag{13}$$
- and

$$L_1 b' = L_1 G' x' + L_1(\hat{e} \circ (-1)). \tag{14}$$
- Let $e' := P_m(b' - G' x')$.
- Output $\langle e' \circ (-1), c \rangle$.

We first show that the systems of linear equations (13) and (14) have a joint solution. Let $x \in \mathbb{F}_2^n$ be such that $w = Gx$. Then for every $\xi \in \{0, 1\}$ we have that

$$L_\xi G' x + L_\xi(\hat{e} \circ (-1)) = L_\xi(b' - e \circ (-1)) + L_\xi(\hat{e} \circ (-1)) = L_\xi b'.$$

It remains to show that if c is an encryption of a bit $\sigma \in \{0, 1\}$ then $\langle e' \circ (-1), c \rangle = \sigma$. To see this note first that $L_\sigma(b' - G' x') = L_\sigma(\hat{e} \circ (-1))$ and so $b' - G' x' = \hat{e} \circ (-1) + y$ for some $y \in \ker(L_\sigma) = (\text{supp}(\mu_\sigma))^\perp$. So we have

$$\begin{aligned} \langle e' \circ (-1), c \rangle &= \langle b' - G' x', \tilde{w} \rangle + \langle b' - G' x', \tilde{e} \rangle \\ &= \langle b', \tilde{w} \rangle - \langle x', (G')^T \tilde{w} \rangle + \langle y + \hat{e} \circ (-1), \tilde{e} \rangle \\ &= 0 - 0 + \langle y + \hat{e} \circ (-1), \tilde{e} \rangle \\ &= \langle \hat{e} \circ (-1), \tilde{e} \rangle \\ &= \sigma. \end{aligned}$$

□

8.2 Candidate over constant-size rings

In what follows we consider a generalization of the encryption scheme Π_{GPV} to arbitrary rings \mathbb{Z}_q . Over a general ring \mathbb{Z}_q the standard notions of matrix rank and Gaussian elimination do not apply. However, the formal definition of the encryption scheme Π_{GPV} still makes sense in this setting if we choose H to be a uniform random $m \times (m - n)$ matrix over \mathbb{Z}_q (so we do not require it to be full rank) and we generate \tilde{w} by choosing a uniform random vector $\tilde{x} \in \mathbb{Z}_q^{m-n}$ and letting $\tilde{w} = \tilde{H} \tilde{x}$.

Over sufficiently large constant-size rings, we propose to use the construction of *matching vector families* from [13], which builds on the constructions of [20, 15], as the noise distributions used in the key generation and encryption processes in order to obtain an instance of the abstract encryption scheme with

perfectly correct decryption that cannot be attacked in time $\text{poly}(m)$. We start with the definition of a matching vector family.

Definition 8.2 (Matching vector family). Let $S \subseteq \mathbb{Z}_q \setminus \{0\}$. A pair of subsets $\mathcal{U} = \{u_1, \dots, u_k\}, \mathcal{V} = \{v_1, \dots, v_k\} \subseteq \mathbb{Z}_q^m$ forms an S -matching vector family of size k if the following two conditions are satisfied:

- For all $i \in [k], \langle u_i, v_i \rangle = 0$.
- For all $i \neq j \in [k], \langle u_i, v_j \rangle \in S$.

We say that an S -matching vector family is b -bounded if $s < b$ for every $s \in S$ (where we identify \mathbb{Z}_q with the set of integers $\{0, 1, \dots, q-1\}$ and define the relation $<$ over \mathbb{Z}_q according to the corresponding ordering on the integers).

The following theorem is a consequence of Lemma 13 in [13] with the setting of parameters described in the proof of Lemma 16 in [13]. It shows that for every constant integer t there is a fixed integer q such that there is a bounded matching vector family with $k = 2^{\text{polylog}(m)}$ vectors over \mathbb{Z}_q^m , where the power in the polylog grows with t .

Theorem 8.3. For every integer $t \geq 2$ there exist integers $q = q(t)$ and $k_0 = k_0(t)$ such that for every $k \geq k_0$ there exists an $O(q/\log \log q)$ -bounded matching vector family $\mathcal{U}_k^{(t)}, \mathcal{V}_k^{(t)} \subseteq \mathbb{Z}_q^m$ of size k for $m = \exp((\log k)^{1/t}(\log \log k)^{1-1/t}t \ln t)$. Moreover, the uniform distribution over each of the subsets $\mathcal{U}_k^{(t)}$ and $\mathcal{V}_k^{(t)}$ is efficiently samplable.

Fix $t \geq 2$ and let $q = q(t), k_0 = k_0(t)$ be the integers guaranteed by Theorem 8.3. For every integer $k \geq k_0$ let $\Pi_k^{(t)}$ be the instance of Π_{GPV} over \mathbb{Z}_q in which $m = \exp((\log k)^{1/t}(\log \log k)^{1-1/t}t \ln t)$, $n = m/(1 + \epsilon)$, the distribution μ_{sk} is the uniform distribution over $\mathcal{U}_k^{(t)}$, the distribution μ_0 is the uniform distribution over all vectors of the form $v \circ 0$ and the distribution μ_1 is the uniform distribution over all vectors of the form $v \circ (-\lfloor \frac{q}{2} \rfloor)$ where $v \in \mathcal{V}_k^{(t)}$ and the decryption function g is given by

$$g(i) = \begin{cases} 1, & \lfloor \frac{q}{4} \rfloor \leq i \leq \lfloor \frac{3q}{4} \rfloor \\ 0, & \text{otherwise.} \end{cases}$$

Since the matching vector family $\mathcal{U}_k^{(t)}, \mathcal{V}_k^{(t)}$ is $O(q/\log \log q)$ -bounded, for sufficiently large q we have that $\text{Adv}^{\text{Dec}}(n) = 1$. Furthermore, we suggest that the 'non-linear' structure of matching vector families makes them resistance to polynomial time attacks.

Question 8.4. Does it hold for any integers $t \geq 2$ and $k \geq k_0(t)$ that any attack on $\Pi_k^{(t)}$ running in time $\text{poly}(m)$ will only be able to achieve a negligible advantage in n ?

Acknowledgements We thank Parikshit Gopalan, Elad Haramaty, Swastik Kopparty, Shachar Lovett, Oded Regev, Amir Shpilka, Shubhangi Saraf and Ben Lee Volk for useful conversations.

References

- [1] M. Ajtai. Generating hard instances of lattice problems (extended abstract). In *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing (STOC)*, pages 99–108. ACM Press, 1996.

- [2] M. Ajtai and C. Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing (STOC)*, pages 284–293. ACM Press, 1997.
- [3] M. Alekhnovich. More on average case vs approximation complexity. *Computational Complexity. Preliminary version in proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2003)*, 20(4):755–786, 2011.
- [4] B. Applebaum, D. Cash, C. Peikert, and A. Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In S. Halevi, editor, *CRYPTO*, volume 5677 of *Lecture Notes in Computer Science*, pages 595–618. Springer, 2009.
- [5] B. Applebaum, Y. Ishai, and E. Kushilevitz. Cryptography with constant input locality. *J. Cryptology*, 22(4):429–469, 2009.
- [6] E. Ben-Sasson, S. Lovett, and N. Ron-Zewi. An additive combinatorics approach relating rank to communication complexity. *Journal of the ACM*, to appear, 2013.
- [7] E. Ben-Sasson and N. Zewi. From affine to two-source extractors via approximate duality. In *proceedings of the 43rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 177–186. ACM Press, 2011.
- [8] A. Bhowmick, Z. Dvir, and S. Lovett. New bounds for matching vector families. In *Proceedings of the 47rd ACM Symposium on Theory of Computing (STOC)*, pages 823–832. ACM Press, 2013.
- [9] A. Blum, M. L. Furst, M. J. Kearns, and R. J. Lipton. Cryptographic primitives based on hard learning problems. In *CRYPTO*, pages 278–291, 1993.
- [10] A. Blum, A. Kalai, and H. Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. *J. ACM*, 50(4):506–519, 2003.
- [11] I. Damgård and S. Park. Is public-key encryption based on LPN practical? *IACR Cryptology ePrint Archive*, 2012:699, 2012.
- [12] W. Diffie and M. E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654, 1976.
- [13] Z. Dvir, P. Gopalan, and S. Yekhanin. Matching vector codes. *SIAM J. Comput. Preliminary version in proceedings of the IEEE 51st Annual Symposium on Foundations of Computer Science (FOCS 2010)*, 40(4):1154–1178, 2011.
- [14] C. Dwork, M. Naor, and O. Reingold. Immunizing encryption schemes from decryption errors. In *EUROCRYPT*, pages 342–360, 2004.
- [15] K. Efremenko. 3-query locally decodable codes of subexponential length. *SIAM J. Comput. Preliminary version in proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC 2009)*, 41(6):1694–1703, 2012.
- [16] T. ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Transactions on Information Theory*, 31(4):469–472, 1985.

- [17] C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*. ACM Press, 2008.
- [18] O. Goldreich, S. Goldwasser, and S. Halevi. Public-key cryptosystems from lattice reduction problems. In *CRYPTO*, pages 112–131, 1997.
- [19] B. Green. Finite field models in additive combinatorics. In *London Mathematical Society Lecture Note Series*, volume 324. Cambridge University Press, 2005.
- [20] V. Grolmusz. Superpolynomial size set-systems with restricted intersections mod 6 and explicit ramsey graphs. *Combinatorica*, 20:71–86, 2000.
- [21] J. Hoffstein, J. Pipher, and J. H. Silverman. Ntru: A ring-based public key cryptosystem. In *ANTS*, pages 267–288, 1998.
- [22] N. J. Hopper and M. Blum. Secure human identification protocols. In *ASIACRYPT*, pages 52–66, 2001.
- [23] A. T. Kalai, Y. Mansour, and E. Verbin. On agnostic boosting and parity learning. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*, pages 629–638. ACM Press, 2008.
- [24] S. Kopparty and S. Saraf. Local list-decoding and testing of random linear codes from high error. *SIAM J. Comput*, 42(3):1302–1326, 2013.
- [25] R. Lindner and C. Peikert. Better key sizes (and attacks) for LWE-based encryption. In A. Kiayias, editor, *CT-RSA*, volume 6558 of *Lecture Notes in Computer Science*, pages 319–339. Springer, 2011.
- [26] S. Lovett. Additive combinatorics and its applications in theoretical computer science, 2013.
- [27] S. Lovett. Communication is bounded by root of rank. In *Proceedings of the 46th ACM Symposium on Theory of Computing (STOC)*. ACM Press, 2014.
- [28] V. Lyubashevsky. The parity problem in the presence of noise, decoding random linear codes, and the subset sum problem. In *RANDOM: International Workshop on Randomization and Approximation Techniques in Computer Science*. LNCS, 2005.
- [29] V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. *J. ACM*, 60(6):43, 2013.
- [30] R. J. McEliece. A public-key cryptosystem based on algebraic coding theory. JPL DSN Progress Report, 1978.
- [31] D. Micciancio. Improving lattice based cryptosystems using the hermite normal form. In J. H. Silverman, editor, *CaLC*, volume 2146 of *Lecture Notes in Computer Science*, pages 126–145. Springer, 2001.
- [32] D. Micciancio. Duality in lattice-based cryptography. In *Public key cryptography (invited talk)*, 2010.

- [33] D. Micciancio and P. Mol. Pseudorandom knapsacks and the sample complexity of LWE search-to-decision reductions. In P. Rogaway, editor, *CRYPTO*, volume 6841 of *Lecture Notes in Computer Science*, pages 465–484. Springer, 2011.
- [34] D. Micciancio and O. Regev. Lattice-based cryptography. In *Post-quantum cryptography*, pages 147–192. Springer-Verlag, pub-SV:adr, 2009.
- [35] H. Niederreiter. Knapsack-type cryptosystems and algebraic coding theory. *Problems of Control and Information Theory (Problemy Upravlenija i Teorii Informacii)*, 15:159–166, 1986.
- [36] K. Pietrzak. Cryptography from learning parity with noise. In *SOFSEM*, pages 99–114, 2012.
- [37] M. Rabin. Digitalized signatures and public-key functions as intractable as factorization. MIT LCS TR-212, 1979.
- [38] O. Regev. New lattice-based cryptographic constructions. *J. ACM*, 51(6):899–942, 2004.
- [39] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM. Preliminary version in proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC 2005)*, 56(6):34:1–34:40, 2009.
- [40] O. Regev. The learning with errors problem (invited survey). In *IEEE Conference on Computational Complexity*, pages 191–204, 2010.
- [41] R. L. Rivest, A. Shamir, and L. M. Adleman. A method for obtaining digital signatures and public-key cryptosystems (reprint). *Commun. ACM*, 26(1):96–99, 1983.
- [42] A. Sahai and B. Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In *Proceedings of the 46th Annual ACM Symposium on the Theory of Computing (STOC)*, pages 475–484. ACM Press, 2014.
- [43] D. Stehlé, R. Steinfeld, K. Tanaka, and K. Xagawa. Efficient public key encryption based on ideal lattices. In M. Matsui, editor, *ASIACRYPT*, volume 5912 of *Lecture Notes in Computer Science*, pages 617–635. Springer, 2009.
- [44] T. Tao and V. Vu. *Additive Combinatorics*. Cambridge University Press, Cambridge, 2006.
- [45] S. P. Vadhan. Pseudorandomness. *Foundations and Trends in Theoretical Computer Science*, 7(1-3):1–336, 2012.

A Security of abstract encryption scheme – proof of Lemma 5.4

For the proof of Lemma 5.4 we shall use the following *decision to search reduction* which shows that in order to find the vector s it suffices to be able to distinguish the distribution $D_{s, \mu_{sk}}$ from the uniform distribution. Lemma 4.2 in [39] shows such a reduction for the case in which μ_{sk} is i.i.d in each coordinate and we observe that the same holds also when μ_{sk} is arbitrary.

Lemma A.1 (Decision to search). *Let $n \leq m$ be integers, q a prime and μ an arbitrary distribution over \mathbb{F}_q^m . Suppose that there exists a randomized procedure W' , running in time $t(n)$, which is given as input ℓ samples of the form $(G^{(i)}, b^{(i)})$ where $G^{(i)} \in \mathbb{F}_q^{m \times n}$ and $b^{(i)} \in \mathbb{F}_q^m$ for all $1 \leq i \leq \ell$ and which satisfies the following guarantees. For all $s \in \mathbb{F}_q^n$, the procedure W' accepts with probability at least $1 - \delta$ if all $(G^{(i)}, b^{(i)})$ are independent random samples from $D_{s, \mu}$ and the procedure W' rejects with probability at least $1 - \delta$ if all $G^{(i)}$ are independent uniform random $m \times n$ matrices over \mathbb{F}_q and all $b^{(i)}$ are independent uniform random vectors in \mathbb{F}_q^m . Then there exists a randomized algorithm W , running in time $t(n) \cdot \text{poly}(m, q)$, which for every $s \in \mathbb{F}_q^n$, given ℓ independent random samples from $D_{s, \mu}$, outputs s with probability at least $1 - qn\delta$.*

Proof. The algorithm W draws ℓ independent uniform random vectors $r^{(1)}, \dots, r^{(\ell)}$ in \mathbb{F}_q^m and for every $i \in [\ell]$ and $j \in [n]$ lets $R^{(i,j)}$ denote the $m \times n$ matrix which is all zero except for its j -th column which equals $r^{(i)}$. In order to find the j -th coordinate s_j of s , for any $\eta \in \mathbb{F}_q$ the algorithm W runs the procedure W' on the samples $(G^{(1)} + R^{(1,j)}, b^{(1)} + \eta \cdot r^{(1)}), \dots, (G^{(\ell)} + R^{(\ell,j)}, b^{(\ell)} + \eta \cdot r^{(\ell)})$. If W' accepts for some $\eta \in \mathbb{F}_q$ the algorithm W sets $s_j = \eta$. Otherwise, it sets s_j to an arbitrary value.

By union bound, it suffices to show that for every $j \in [n]$ the algorithm W succeeds in finding s_j with probability at least $1 - q\delta$. For this note that if $s_j = \eta$ then $(G^{(1)} + R^{(1,j)}, b^{(1)} + \eta \cdot r^{(1)}), \dots, (G^{(\ell)} + R^{(\ell,j)}, b^{(\ell)} + \eta \cdot r^{(\ell)})$ are ℓ independent random samples from $D_{s, \mu}$ and so W' accepts with probability at least $1 - \delta$. On the other hand, if $s_j \neq \eta$, we have that $(G^{(1)} + R^{(1,j)}, b^{(1)} + \eta \cdot r^{(1)}), \dots, (G^{(\ell)} + R^{(\ell,j)}, b^{(\ell)} + \eta \cdot r^{(\ell)})$ are ℓ independent uniform random samples and so W' rejects with probability at least $1 - \delta$. By union bound over all $\eta \in \mathbb{F}_q$ this implies in turn that W outputs s_j with probability at least $1 - q\delta$ which concludes the proof of the lemma. \square

For the proof of Lemma 5.4 we shall also need the following version of the leftover hash lemma which can be found for example in [45] (Theorem 6.18). We say that a family \mathcal{H} of hash functions from X to Y is *universal* if $\Pr_{h \in \mathcal{H}}[h(x_1) = h(x_2)] \leq 1/|Y|$ for all $x_1 \neq x_2 \in X$.

Lemma A.2 (Leftover hash lemma). *Let $\mathcal{H} \subseteq \{h : X \rightarrow Y\}$ be a universal family of hash functions from X to Y and let μ be a distribution over X with $H_\infty(\mu) \geq k$. Then for all but a $\sqrt[4]{|Y|2^{-k}}$ -fraction of $h \in \mathcal{H}$ it holds that the distribution $h(\mu)$ is $\sqrt[4]{|Y|2^{-k}}$ -close to the uniform distribution over Y .*

We proceed to the proof of Lemma 5.4. f

Proof of Lemma 5.4. By Lemma A.1 it suffices to show a randomized procedure W' , running in time $(t(n)\ell/\epsilon^2) \cdot \text{poly}(m, q)$, which given as input ℓ samples of the form $(G^{(1)}, b^{(1)}), \dots, (G^{(\ell)}, b^{(\ell)})$, for all $s \in \mathbb{F}_q^n$, accepts with probability at least $1 - \delta^\ell - \ell \exp(-n)$ if all $(G^{(i)}, b^{(i)})$ are independent random samples from $D_{s, \mu_{\text{sk}}}$ and rejects with probability at least $1 - \delta^\ell - \ell \exp(-n)$ if all $(G^{(i)}, b^{(i)})$ are independent uniform random samples.

The idea of the proof is as follows. Using the samples $(G^{(1)}, b^{(1)}), \dots, (G^{(\ell)}, b^{(\ell)})$ the procedure W' will generate ℓ public keys $\tilde{G}^{(1)}, \dots, \tilde{G}^{(\ell)}$ and for each of these keys it will generate $\ell' \approx n/\epsilon^2$ independent random encryptions of 0 and 1 using the key. The procedure W' will then run the attack \mathcal{A} on each of the public keys and corresponding ciphertexts and estimate the advantage of the attack \mathcal{A} in guessing the message bit. The procedure W' will accept if and only if this advantage is sufficiently high. The main observation is that if the samples $(G^{(i)}, b^{(i)})$ are drawn from the distribution $D_{s, \mu_{\text{sk}}}$ then the attack \mathcal{A} will be able to distinguish between encryptions of 0 and 1, while if the samples $(G^{(i)}, b^{(i)})$ are uniform random the public keys and corresponding encryptions would be almost uniform random (thanks to the leftover hash lemma) and therefore \mathcal{A} will not be able to distinguish between encryptions of 0 and 1. Details follow.

$W'((G^{(1)}, b^{(1)}), \dots, (G^{(\ell)}, b^{(\ell)}))$

- For $i = 1, \dots, \ell$:

- Draw a uniform random vector $y \in \mathbb{F}_q^n$ and let

$$\tilde{G}^{(i)} = \begin{pmatrix} G^{(i)} & b^{(i)} + G^{(i)}y \\ 0_n^T & -1 \end{pmatrix}.$$

- For $\ell' = n/(\epsilon/8)^2$:

- Draw ℓ' independent random samples $e_1^{(0)}, \dots, e_{\ell'}^{(0)} \sim \mu_0$.

- Draw ℓ' independent random samples $e_1^{(1)}, \dots, e_{\ell'}^{(1)} \sim \mu_1$.

- For every $\xi \in \{0, 1\}$ and $j \in [\ell']$ run the attack \mathcal{A} with the public key $\tilde{G}^{(i)}$ and the ciphertext $(\tilde{G}^{(i)})^T \cdot e_j^{(\xi)}$.

- If

$$\Pr_{j \in [\ell']} \left[\mathcal{A} \left(\tilde{G}^{(i)}, (\tilde{G}^{(i)})^T \cdot e_j^{(1)} \right) = 1 \right] - \Pr_{j \in [\ell']} \left[\mathcal{A} \left(\tilde{G}^{(i)}, (\tilde{G}^{(i)})^T \cdot e_j^{(0)} \right) = 1 \right] \geq \frac{3\epsilon}{4}, \quad (15)$$

output accept.

- Else if (15) does not hold for every $i = 1, \dots, \ell$, output reject.

Inspection reveals that the running time of W' is $(t(n)\ell/\epsilon^2) \cdot \text{poly}(m, q)$. It remains to analyze the acceptance probability of W' on samples from $D_{s, \mu_{\text{sk}}}$ and the rejection probability on uniform random samples.

We start with analyzing the acceptance probability on samples from $D_{s, \mu_{\text{sk}}}$. Observe that in this case $\tilde{G}^{(1)}, \dots, \tilde{G}^{(\ell)}$ are independent random public keys in Π_{Reg}^{ℓ} . Since $\text{Adv}_{\delta}^{\mathcal{A}}(n) \geq \epsilon$ this implies in turn that with probability at least $1 - \delta^{\ell}$ there exists $i \in [\ell]$ such that

$$\Pr[\mathcal{A}(\tilde{G}^{(i)}, \text{Enc}_{\tilde{G}^{(i)}}(1)) = 1] - \Pr[\mathcal{A}(\tilde{G}^{(i)}, \text{Enc}_{\tilde{G}^{(i)}}(0)) = 1] \geq \epsilon.$$

Conditioned on this, the Hoeffding bound for sampling (Lemma 3.1) implies that (15) holds with probability at least $1 - 4 \cdot 2^{-2n}$. So in this case the procedure W' accepts with probability at least $1 - \delta^{\ell} - 4 \cdot 2^{-2n} = 1 - \delta^{\ell} - \exp(-n)$.

Next we analyze the rejection probability on independent uniform random samples. Fix $i \in [\ell]$. Applying the leftover hash lemma (Lemma A.2) with $X = \mathbb{F}_q^{m+1}$, $Y = \mathbb{F}_q^{n+1}$,

$$\mathcal{H} = \left\{ \tilde{G}^T \in \mathbb{F}_q^{(n+1) \times (m+1)} \mid \tilde{G} = \begin{pmatrix} G & b \\ 0_n^T & -1 \end{pmatrix} \text{ for } G \in \mathbb{F}_q^{m \times n} \text{ and } b \in \mathbb{F}_q^m \right\}$$

and $\mu = \mu_0$ we obtain that with probability at least $1 - 2^{-\alpha(n+1) \log q/4}$, the distribution $(\tilde{G}^{(i)})^T \cdot \mu_0$ is $2^{-\alpha(n+1) \log q/4}$ -close to the uniform distribution over \mathbb{F}_q^{n+1} . Similarly, with probability at least $1 - 2^{-\alpha(n+1) \log q/4}$, the distribution $(\tilde{G}^{(i)})^T \cdot \mu_1$ is $2^{-\alpha(n+1) \log q/4}$ -close to the uniform distribution over \mathbb{F}_q^{n+1} .

So with probability at least $1 - 2 \cdot 2^{-\alpha(n+1) \log q/4}$ the distributions $(\tilde{G}^{(i)})^T \cdot \mu_0$ and $(\tilde{G}^{(i)})^T \cdot \mu_1$ are $2 \cdot 2^{-\alpha(n+1) \log q/4}$ -close. Conditioned on this, the Hoeffding bound for sampling implies that (15) holds with probability at most $4 \cdot 2^{-2n}$ due to our assumption that $2 \cdot 2^{-\alpha(n+1) \log q/4} \leq \epsilon/2$. So by union bound over all $i \in [\ell]$, in this case the procedure W' rejects with probability at least $1 - \ell(2 \cdot 2^{-\alpha(n+1) \log q/4} + 4 \cdot 2^{-2n}) = 1 - \ell \exp(-n)$. \square