

Chosen Ciphertext Secure Pseudorandom Codes in the Standard Model

Nico Döttling¹, Antoine Joux¹, Venkata Koppula²,
Mahesh Sreekumar Rajasree¹, and Hendrik Waldner¹

¹ CISA Helmholtz Center for Information Security, Saarbrücken, Germany
{doettling, joux, hendrik.waldner}@cispa.de, srmahesh1994@gmail.com

² Indian Institute of Technology Delhi, New Delhi, India
kvenkata@iitd.ac.in

Abstract. Pseudorandom codes (PRCs), recently proposed by Christ and Gunn (CRYPTO’24), are encryption schemes that have pseudorandom ciphertexts and a decryption algorithm which is resilient against a bounded number of Hamming errors. This notion provides a significant strengthening over standard PKE and has exciting applications in, e.g., watermarking LLMs. The recent work of Alrabiah et al. (STOC’25) initiated the study of CCA-secure public-key PRCs, where the adversary is additionally given access to a decoding oracle. In terms of realizations, they provide one that can be proven secure in the random oracle model. Constructing CCA-secure public-key PRCs in the standard model remained an open problem.

In this work, we resolve this problem and provide the first construction of CCA-secure public-key PRCs in the standard model. Our construction achieves a constant rate and can decode from a constant fraction of adversarial errors. In fact, our construction is a general blueprint that can be instantiated from a broad range of standard cryptographic assumptions.

As an additional contribution, we construct a strong adaptively robust public-key pseudorandom code with conjectured sub-exponential security based on a new family of assumptions we call Noisy McEliece. In a nutshell, these assumptions mask a scrambled generator matrix from an efficiently decodable inner-code family with sparse Bernoulli noise; this additional error is meant to obscure the algebraic structure targeted by known attacks and thereby permits candidate instantiations from broader classes of codes.

1 Introduction

1.1 Pseudorandom Codes

Error-correcting codes and pseudorandomness are two of the most fundamental notions in theoretical computer science and cryptography. Error-correcting codes enable reliable communication over noisy channels, while pseudorandom generators (PRGs) are indispensable for derandomization and cryptographic security. While error-correction requires redundancy and hence structure, pseudorandomness makes any structure undetectable.

The notion of pseudorandom codes (PRCs), recently proposed by Christ and Gunn [CG24], connects these seemingly contradictory notions. In a nutshell, a PRC is a family of keyed error-correcting codes whose codewords are computationally indistinguishable from uniformly random ones to an outside observer. But, a legitimate receiver who is in possession of the secret key can efficiently decode the codeword even in the presence of a large fraction of Hamming errors. While pseudorandom codes were originally motivated by the problem of watermarking large language models (LLMs), they have emerged as an exciting cryptographic notion in their own right.

However, constructing pseudorandom codes has turned out to be a highly non-trivial task. Current constructions [CG24, GG24, GM24, AAC⁺25, CGG⁺25] are based on hardness assumptions relating to coding theory, and there is evidence that even constructing just the weakest private-key notions cannot be achieved from *generic* hardness assumptions [DMR25, GGW25]. In particular, [GGW25] argues that even the most powerful tools available in cryptography such as (indistinguishability) obfuscation do not help when attempting to construct pseudorandom codes with generic recipes.

1.2 Adaptive Robustness and CCA Security

Adaptive Robustness: Early PRC constructions addressed errors introduced obliviously by a memoryless channel. Subsequent works [CHS25, AAC⁺25] argued that, in realistic settings, robustness must be defined against *adaptive* adversaries that have access to an encoding oracle (or to the public key in the public-key setting). Such an adversary must output a message m , randomness r , and a c such that encoding m with randomness r yields a codeword close to c , yet decoding c produces a different message. In this work, we rely on this stronger notion of robustness.

CCA security: The standard security notion for pseudorandom codes [CG24] mirrors the basic *chosen plaintext* security notion (CPA) for encryption schemes with pseudorandom ciphertexts: the adversary is given access to an oracle $\mathcal{O}_{Enc}$ that produces strings c_i (as well as potentially the PRC public key) and must decide whether the strings produced by the oracle are well-formed codes of the PRC or just uniformly random strings. Several real-world attacks [JSV24, PHZS24] exploit access to both the encoding and decoding oracles in order to remove/counterfeit watermarks. Alrabiah et al. [AAC⁺25] introduced stronger security notions in the private key and public key setting, where the adversary also has access to the decoding oracle. In the private key setting, they introduced the notion of *ideal pseudorandom codes* for pseudorandom codes. In this notion, the adversary is additionally given access to a *decoding* oracle $\mathcal{O}_{Dec}$ which provides the adversary access to the decoding algorithm. Implementing this oracle for the case that $\mathcal{O}_{Enc}$ actually implements the encoding algorithm is easy. However, if $\mathcal{O}_{Enc}$ produces uniformly random strings, a more refined argument is required. Specifically, in this case $\mathcal{O}_{Enc}$ and $\mathcal{O}_{Dec}$ must share a joint state, and $\mathcal{O}_{Dec}$ needs to handle strings produced by $\mathcal{O}_{Enc}$ differently from codewords produced by the adversary.

In the public key setting, which is the focus of this work, [AAC⁺25] proposed a simple *chosen ciphertext* security notion (CCA) and provided a construction in the random oracle model [BR93]. Specifically, in this security notion the adversary is given the public key, as well as access to a decoding oracle $\mathcal{O}_{Dec}$ which implements the decoding algorithm. At some point, the adversary asks for a challenge and is either given a well-formed codeword, or a uniformly random string. As in the standard CCA experiment, access to the decoding oracle is now restricted. While, at this point in the CCA experiment for public key encryption, the decryption oracle accepts every query *except* the challenge ciphertext, this formulation is insufficient for pseudorandom codes. Specifically, the adversary can always introduce a small amount of errors to the challenge and obtain a different string. Hence, in the PRC-setting the post-challenge decoding oracle must be restricted in a more stringent way. Roughly speaking, the post-challenge decoding oracle now summarily rejects any input within (say) the decoding-radius of the challenge.

1.3 Constructing CCA-Secure Pseudorandom Codes

However, in light of mounting impossibility results for PRCs [DMR25, GGW25], and given the growing body of random oracle uninstantiability results [Nie02, GK03, BBP04, LN09, Zha22, BDD22], the very existence of the notion of CCA-secure PRC rests on uncertain grounds. This naturally raises the following question:

Can CCA-secure pseudorandom codes be realized in the standard model under standard assumptions?

1.4 Our Results

In this work, we present the first construction of public-key pseudorandom codes that achieves indistinguishability against chosen-ciphertext attacks in the standard model. We also present a new candidate construction of an adaptively robust PRC.

New Candidates for Robust PRCs. Existing constructions of robust public-key PRCs need to assume sub-exponential hardness of LPN [AAC⁺25], while only providing quasi-polynomial security. Furthermore, constructions attaining sub-exponential security are known only in the secret-key setting [CGG⁺25]. As an additional contribution, we propose a new family of assumptions which we call *Noisy McEliece* assumptions.

Recall that the decisional McEliece assumption [McE78] conjectures that the generator matrices of *scrambled Goppa codes* are pseudorandom. Specifically, if $\mathbf{G} \in \mathbb{F}_2^{m \times n}$ is the generator matrix of a binary Goppa code, $\mathbf{\Pi} \in S_m$ is a random permutation matrix, and $\mathbf{T} \in \text{GL}_n(\mathbb{F}_2)$ is a random invertible matrix, then we set $\mathbf{A} = \mathbf{\Pi} \cdot \mathbf{G} \cdot \mathbf{T}$. Under the decisional McEliece assumption this matrix $\mathbf{A}$ is pseudorandom. Notoriously, virtually all attempts to replace binary Goppa codes in this construction with other efficiently decodable linear codes have led to insecure constructions [SS92, vT94, CC98, GOT12, CGGU⁺14, COTGU15, Mce26]. These attacks show that the choice of the code family is central. In our Noisy McEliece assumptions, one fixes a specified code family and masks the scrambled generator matrix by Bernoulli noise, that is, we consider matrices of the form $\mathbf{A} = \mathbf{\Pi} \cdot \mathbf{G} \cdot \mathbf{T} + \mathbf{E}$, where $\mathbf{E} \leftarrow \text{Ber}_\rho^{m \times n}$ is chosen from a $\rho = O(1/\sqrt{n})$ -noise Bernoulli distribution. The hope is that the noise disrupts the exact algebraic invariants exploited by existing attacks, but this depends on the chosen code family. In Section 4 we discuss candidate code families, insecure regimes, and the need for further cryptanalysis, and provide a construction of an adaptively robust pseudorandom code with conjectured sub-exponential security from this assumption.

Public-Key Pseudorandom Codes with CCA Security. In Section 5, we describe a generic transformation for constructing an IND-CCA-secure PRC scheme using a primitive called *hinting pseudorandom generator* (HPRG) [KW19].

Theorem 1 (Informal). *Assuming the existence of an IND-CPA-secure PRC scheme with adaptive robustness, and a secure hinting PRG, there exists an IND-CCA-secure PRC scheme.*

Instantiations. Since every PRC is a PKE scheme, the required building blocks can be instantiated using several well-studied hardness assumptions:

- **Adaptively Robust PRC:** Under sub-exponential Learning Parity with noise (exLPN) [AAC⁺25], or our noisy McEliece assumption.
- **Hinting PRGs:** Constructions are known from Learning with Errors (LWE) [KW19], Computational Diffie–Hellman (CDH) [KW19], Decisional Diffie–Hellman (DDH) [AP22].

These instantiations yield the following corollary.

Corollary 1. *Assuming the hardness of $\mathcal{X}$ and $\mathcal{Y}$, where $\mathcal{X} \in \{\text{LWE}, \text{DDH}, \text{CDH}\}$ and $\mathcal{Y} \in \{\text{noisy McEliece}, \text{exLPN}\}$, there exists an IND-CCA-secure PRC scheme.*

Remark 1 (Instantiating our results). The above construction relies on a specialized form of hinting PRG, which we refer to as an *entropic hinting PRG*. This is essentially a hinting PRG that remains secure even when its seed is not sampled uniformly at random, but instead drawn from a distribution with sufficiently high entropy. Notably, most of the known constructions of hinting PRGs employ the *leftover hash lemma* in their security analysis, thereby ensuring the entropic property. For further details, see Section 3.5.

One exception to the above is the Linear Hidden Shift (LHS)-based construction of [AP22]. In the proof of this construction, the pseudorandomness of the output is directly argued by relying on the LHS assumption. While this is out of scope of our current work, we believe that the same proof strategy as in [AP22] gives rise to entropic hinting PRGs, if one is willing to rely on a high entropy version of the LHS assumption. If such an assumption is plausible, we would obtain an additional instantiation of the result above.

1.5 Related Works

Christ and Gunn [CG24] introduced the notion of PRCs and gave constructions in both, the secret-key and public-key settings. The subsequent work [AAC⁺25] strengthened the security model by allowing the adversary to have access to the decoding algorithm, introducing the corresponding CCA notion. More recently, [CGG⁺25] constructed secret-key PRCs that, for the first time, simultaneously achieve plausible subexponential pseudorandomness security, robustness to worst-case edit errors over a binary alphabet, and resilience even against computationally unbounded adversaries equipped with the decoding key. Achieving this combination of

properties requires overcoming several barriers, as prior constructions neither attained subexponential security nor provided strong resilience guarantees. Zero-bit PRCs³ were constructed in [GM24] under the assumption of local weak pseudorandom functions, which in turn follow from the hardness of learning $\log(n)$ -juntas over the uniform distribution on $\{0, 1\}^n$. Additional constructions of pseudorandom codes were explored in [GG24]. In particular, they construct public-key PRCs based on the planted hyperloop assumption [BKR23] and pseudorandom generators, and also obtain statistically secure PRCs against space-bounded adversaries. On the contrary, two independent works [GGW25, DMR25] establish strong impossibility results for black-box constructions of pseudorandom codes from a broad class of strong cryptographic primitives including virtual black-box obfuscation.

Next, we give a brief overview of existing IND-CPA to IND-CCA transformations. None of these transformations consider efficient error correction, which makes them unsuitable for amplifying the security of pseudorandom codes directly. We categorize the existing transformations in the following way:

1. **Proof-Carrying Ciphertexts:** This is the classic paradigm for constructing CCA-secure PKE. In this paradigm, multiple ciphertexts of a CPA secure encryption scheme, generated under different public keys, are *glued together* using a NIZK proof of well-formedness. Such proofs can be publicly verifiable (e.g. [NY90, DDN91, Sah99]) or designated verifier (e.g. [CS98, CS02, QRW19, LQR⁺19]). Evidently, using publicly verifiable NIZK proofs is incompatible with the goal of constructing CCA encryption with pseudorandom ciphertexts. By contrast, schemes based on designated verifier proofs can indeed yield CCA-secure encryption schemes with pseudorandom ciphertexts, as exemplified by the well-known Cramer-Shoup encryption scheme [CS98, BBDP01].
2. **Tag/Label-based Encryption and the CHK-Transform:** An alternative standard-model approach to constructing CCA-secure PKE follows the CHK-paradigm [CHK03, BCHK07, BK05] to transform an identity-based encryption (IBE) scheme or, more generally, a tag-based encryption (TBE) scheme [Kil06] by embedding a one-time signature verification key into the identity/tag, and by including a signature of the IBE/TBE ciphertext into the CCA ciphertext. This blueprint can be instantiated from a broad range of assumptions, including pairings [CHK03], lossy trapdoor functions [PW08], lattice assumptions [MP12], coding assumptions [DMN12, KMP14], hinting PRGs [KW19] and injective trapdoor functions [HKW20]. This paradigm does not typically yield CCA-encryption with pseudorandom ciphertexts, since the ciphertext includes both, a signature and a corresponding verification key, which obey a publicly testable relation. One exception here is the transformation of Boneh and Katz [BK05], which has been shown in [Xag22]. In this work, the author shows that the Boneh-Katz transformation achieves pseudorandom ciphertexts if all of the underlying primitives are pseudorandom as well.
3. **CCA-security via KEM:** In [KMT19], besides multiple other results, the authors show how to realize an IND-CCA secure key encapsulation mechanism (KEM) from an IND-CPA secure KEM as well as a one-time key-dependent-message (KDM) secure SKE scheme. Relying on the IND-CCA secure KEM, an IND-CCA secure PKE scheme can then be obtained using a hybrid encryption approach. To obtain the IND-CPA secure KEM, the authors follow the framework of Koppula and Waters [KW19], besides also mentioning other possible instantiations [BH08, ACPS09, BG10, BLSV18, DGHM18]. Additionally, the authors also show how their IND-CCA transformation can be adapted to obtain IND-CCA secure KEMs⁴. To achieve this result, the presented compiler is instantiated using an IND-CPA secure KEM, instead of an IND-CPA secure one, as well as a one-time randomness-recoverable KDM secure SKE with pseudorandom ciphertexts. Achieving IND-CPA secure KEM, as well as amplifying an existing one-time KDM secure SKE by adding randomness recoverability and pseudorandom ciphertexts, can be easily achieved. On a more technical level, in this transformation, the authors first realize an adaptively secure TDF which directly implies the desired IND-CCA secure KEM. Given that the authors of [KMT19] generalized their transformation for adaptively secure TDFs, it is also possible to use the work of [KM25], in which the authors construct an adaptively secure TDF with pseudorandom strings based on a standard TDF and a PRG, to obtain IND-CCA secure KEMs.

³A zero-bit PRC can encode only a single message $\text{msg} = 0$.

⁴An IND-CCA secure KEM can be turned into an IND-CCA secure PKE using the same hybrid encryption approach as for the IND-CCA case.

4. **Constructions in the Random Oracle Model:** The random oracle model (ROM) [BR93] greatly simplifies the construction of CCA-secure encryption schemes as it facilitates the simulation of the decryption oracle. Specifically, the random oracle enables extraction of the ciphertext randomness of adversarially generated ciphertexts, hence decryption queries can be answered without access to the secret key. Constructing CCA-secure encryption schemes in the ROM is typically practically motivated, as the resulting CCA-secure encryption schemes usually achieve essentially the same efficiency as the underlying CPA-secure encryption schemes. In fact, the random oracle model was formalized precisely for this purpose [BR93]. The Fujisaki-Okamoto transform [FO99] allows to transform any CPA-secure PKE into a CCA-secure encryption scheme without changing the ciphertext structure. Hence, in particular, it enables the construction of CCA-secure encryption schemes with pseudorandom ciphertexts from any CPA-secure encryption scheme with this property. On a conceptual level, while constructions in the ROM yield practically efficient schemes, they do not provide deep foundational insights.

As already mentioned in Section 1.4, in this work, we rely on the hinting PRG framework of Koppula-Waters [KW19]. We leave it as an interesting open problem if any of the other mentioned CCA transformations can also be applied to handle efficient error correction and obtain IND-CCA secure PRCs.

2 Technical Overview

In this section, we give a high-level overview of our construction. We begin by recalling the Koppula-Waters (KW) [KW19] construction of an IND-CCA-2 scheme.

2.1 The KW construction.

The KW construction gave a template for upgrading any IND-CPA PKE scheme to a CCA secure one, using a special kind of pseudorandom generators called *hinting PRGs*. A hinting PRG is a deterministic function $\text{HPRG} : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{\ell + \lambda^2}$ (where ℓ is polynomial in λ) with the following security property: for any PPT adversary, the following distribution looks uniform:

- sample seed $s \leftarrow \{0, 1\}^\lambda$ and compute $\text{HPRG}(s) = (r_0, r_1, \dots, r_\lambda)$. Here $r_0 \in \{0, 1\}^\ell$ and each $r_i \in \{0, 1\}^\lambda$.
- For each i , set $y_{i, s_i} = r_i$, and sample $y_{i, \bar{s}_i} \leftarrow \{0, 1\}^\lambda$.
- Output $(r_0, (y_{i, b})_{i, b})$.

Additionally, the KW construction uses *perfectly binding tag-based equivocal bit commitments*. In a tag-based commitment scheme, we have a setup algorithm that outputs the public parameters. The commitment algorithm Com takes as input the public parameters, the message bit and the tag, and produces a commitment string. For perfect binding, we require that for any public parameters pp , any tag , there does not exist randomness v_0, v_1 such that $\text{Com}(\text{pp}, 0, \text{tag}; v_0) = \text{Com}(\text{pp}, 1, \text{tag}; v_1)$. The equivocal property states that it is possible to generate ‘fake’ public parameters pp^* together with randomness v_0^*, v_1^* for a challenge tag^* such that pp^* is computationally indistinguishable from the regular public parameters, but $\text{Com}(\text{pp}^*, 0, \text{tag}^*; v_0^*) = \text{Com}(\text{pp}^*, 1, \text{tag}^*; v_1^*)$. Almost perfectly binding tag-based equivocal bit commitment schemes can be constructed using any (regular) PRG [Nao90, KW19].

The KW transformation also requires the underlying PKE scheme to satisfy the following two properties:

- *almost perfect correctness* - for almost all public keys pk , for all messages m and randomness r , $\text{Dec}(\text{sk}, \text{Enc}(\text{pk}, m; r)) = m$; and
- *decryption using randomness* - ciphertexts can be decrypted given the encryption randomness.

Both these properties can be achieved for free. Dwork, Naor and Reingold [DNR04a] showed that any IND-CPA PKE scheme can be transformed into another PKE scheme that satisfies almost perfect correctness. Likewise, any PKE scheme can be modified to also support decryption using randomness [KW19]. Now, we describe the setup, encryption and decryption algorithms of the KW scheme:

Key generation. The setup algorithm generates key pairs $(\mathbf{pk}_{i,b}, \mathbf{sk}_{i,b})$ for all $i \in [\lambda]$, $b \in \{0, 1\}$, where λ is the seed length of HPRG. The public key is the collection $(\mathbf{pk}_{i,b})_{i,b}$, while the secret key consists of $(\mathbf{sk}_{i,0})_i$.⁵

Encryption. To encrypt a message m , the sender samples a random seed $s \in \{0, 1\}^\lambda$ and computes $(r_0, r_1, \dots, r_\lambda) = \text{HPRG}(s)$. Next, it samples a one-time signing/verification key $(\mathbf{ds.sk}, \mathbf{ds.vk})$ and randomness $(v_i)_i$ for the commitment scheme. The ciphertext consists of two parts:

- The payload $c = m \oplus r_0$.
- For each $i \in [\lambda]$, it computes a commitment to the bit s_i with tag $\mathbf{ds.vk}$ and randomness v_i (that is, $\text{com}_i \leftarrow \text{Com}(s_i, \mathbf{ds.vk}; v_i)$). Then, it sets

$$c_{i,s_i} = \text{Enc}(\mathbf{pk}_{i,s_i}, 1\|v_i; \tilde{r}_i), \quad c_{i,\bar{s}_i} \leftarrow \text{Enc}(\mathbf{pk}_{i,\bar{s}_i}, 0^\lambda)$$

Finally, the encryption algorithm signs $\mathbf{ds.x} = (c, (c_{i,b}, \text{com}_i)_{i,b})$ using $\mathbf{ds.sk}$, resulting in the signature $\mathbf{ds.\sigma}$. The full ciphertext is $(\mathbf{ds.vk}, \mathbf{ds.x}, \mathbf{ds.\sigma})$.

Decryption. First, the decryption algorithm verifies the signature using the verification key $\mathbf{ds.vk}$. If the signature verifies, the decryption algorithm recovers a candidate seed s , performs consistency checks using s , and finally recovers the message from the payload. More formally,

1. *Recovering the seed.* For each i , it uses $\mathbf{sk}_{i,0}$ to decrypt $c_{i,0}$. If the result is of the form $1\|v_i$ and v_i is a valid opening of com_i to 0 (with respect to tag $\mathbf{ds.vk}$), it sets $s_i = 0$. Otherwise, it sets $s_i = 1$. This yields a candidate seed s .
2. *Verification and message recovery.* Using s , it recomputes $(r_0, (r_i)_i) = \text{HPRG}(s)$. For each $i \in [\lambda]$,
 - it checks that decrypting c_{i,s_i} with randomness r_i ⁶ yields a message of the form $1\|w_i$, and that w_i is a valid opening of com_i for bit s_i (with respect to $\mathbf{ds.vk}$)
 - it checks that re-encrypting $1\|w_i$ using randomness r_i yields c_{i,s_i} .
If all checks pass, it outputs $m = c \oplus r_0$; otherwise, it outputs $\perp$.

The KW scheme, like many other CCA secure encryption schemes, uses the *secret hiding principle* [FS90], where we have multiple secret keys for the same public key. In the security proof, we switch from using one secret key to another. This is a key technical step in the security proof of KW. Note that in the above described construction, for any $\mathbf{p} \in \{0, 1\}^\lambda$, we could have used the secret keys $(\mathbf{sk}_{i,\mathbf{p}_i})_i$ for decryption, and correctness would still hold. KW showed that even for maliciously crafted ciphertexts, decryption of ct using $(\mathbf{sk}_{i,0})_i$ is equal to the decryption of ct using $(\mathbf{sk}_{i,\mathbf{p}_i})_i$. Below, we briefly discuss why this property holds (assuming almost-perfect correctness of PKE and almost-perfect binding of the commitment scheme).

Let $\text{Dec}^\dagger$ denote the decryption using $(\mathbf{sk}_{i,\mathbf{p}_i})_{i \in [\lambda]}$. Fix any ciphertext ct , and let s (resp. $s^\dagger$) denote the candidate seed found in the first step of Dec (resp. $\text{Dec}^\dagger$). Clearly, $s \neq s^\dagger$, otherwise both decryption routines produce the same output. Suppose $s_i = 0$ and $s_i^\dagger = 1$, and suppose Dec does not output $\perp$. First, note that since $s_i \neq s_i^\dagger$, $\mathbf{p}_i$ must be 1. Next, since Dec does not output $\perp$, it follows that $\text{ct}_{i,0}$ is an encryption of $1\|w_i$, and w_i is a valid opening for certifying that com_i is a commitment to bit 0. Finally, let us see why $\text{Dec}^\dagger$ set $s_i^\dagger = 1$: it decrypted $\text{ct}_{i,1}$ and obtained $1\|w_i^\dagger$ such that $w_i^\dagger$ is a valid decommitment of com_i for bit 1. But this is a contradiction since we assumed perfect binding for the commitment scheme. The other case (where $s_i = 1$ and $s_i^\dagger = 0$) relies on the perfect correctness of PKE and is handled similarly.

2.2 IND-CCA PRC with low error rate.

In this section, we present our first CCA-PRC construction. This construction is adaptively α -robust, where $\alpha = \Theta(1/\lambda)$. In the next section, we discuss how to achieve constant rate robustness.

⁵Note that only the $\mathbf{sk}_{i,0}$ secret keys are included, the remaining are discarded.

⁶Here, we use the decryption-using-randomness property of the base PKE scheme.

Making KW ciphertexts pseudorandom: Note that the ciphertexts in the KW encryption scheme are neither pseudorandom, nor error tolerant. We first discuss how to make the ciphertexts pseudorandom. There are three obstacles:

1. *PKE ciphertexts.* The ciphertext components $c_{i,b}$ may not be pseudorandom. This can be fixed by requiring the underlying PKE scheme to have pseudorandom ciphertexts.
2. *Commitments.* The commitments used in the construction may also fail to be pseudorandom. This can be addressed by employing a commitment scheme with pseudorandom commitments. For instance, Naor's commitment scheme (based solely on the existence of a PRG) suffices.
3. *Verification key.* The most crucial issue is that the verification key vk , together with a valid signature, is never pseudorandom, since the verification algorithm itself acts as a distinguisher. However, Lamport's signature scheme provides signatures that are indistinguishable from random strings in the absence of vk . To hide vk , we can embed it inside one of the PKE encryptions, e.g., in both $c_{1,0}$ and $c_{1,1}$. Including it in both components is necessary for the security proof.

We note that CCA-secure schemes with pseudorandom ciphertexts have already been constructed in prior works.

Adding robustness: A first attempt is to replace the CPA-PKE ciphertexts in the previous construction with CPA-PRC ciphertexts. Since each of the ciphertexts are individually pseudorandom, and the payload is also pseudorandom, one can hope that this directly results in an IND-CCA-secure PRC scheme. Unfortunately, this natural attempt is not compatible with the KW proof strategy. In particular, the decryptor must be able to detect if the PRC ciphertexts contain too much noise (otherwise the adversary can create ciphertexts which can detect the set of secret keys being used for decryption), and this is not possible if the decryptor recovers only half the randomness used for the PRC encoding.

We address this issue by using a two-layered approach — first encrypting the commitment openings using a PKE scheme, and then encrypting the resulting ciphertexts and commitments using a CPA-PRC scheme. The key generation algorithm samples $2n$ public/secret key pairs for the PKE scheme $(\text{pke.pk}_{i,b}, \text{pke.sk}_{i,b})_{i,b}$. It also samples $2n$ PRC encoding/decoding keys $(\text{prc.pk}_{i,b}, \text{prc.sk}_{i,b})$, and the public parameters com.pp for the commitment scheme,⁷ and the public parameters hprg.pp for the hinting PRG scheme. We also sample a number ζ uniformly at random (from a sufficiently large field $\mathbb{F}$). The public and secret key for our CCA-PRC scheme are as follows:

$$\text{PK} = \left(\zeta, \text{com.pp}, \text{hprg.pp}, (\text{pke.pk}_{i,b}, \text{prc.pk}_{i,b})_{i,b} \right), \quad \text{SK} = (\text{pke.sk}_{i,b}, \text{prc.sk}_{i,b})_{i,b}.$$

Encryption. The encryption algorithm samples a HPRG seed s and signing/verification keys $(\text{ds.sk}, \text{ds.vk})$. It computes $(r_0, r_1, \dots, r_\lambda) = \text{HPRG}(s)$ and parses $r_0 = (r_{00}, r_{01})$, where both r_{00}, r_{01} are λ bits long. Next, it samples commitment openings $(v_i)_i$, computes $\text{com}_i = \text{Com}(\text{com.pp}, s_i, \text{ds.vk}; v_i)$, and follows the KW template for computing the PKE ciphertexts. For each i , it sets

$$\begin{aligned} \text{pke.ct}_{i,s_i} &= \text{PKE.Enc}(\text{pke.pk}_{i,s_i}, 1 \| v_i \| \text{ds.vk}; r_i) \\ \text{pke.ct}_{i,\overline{s_i}} &\leftarrow \text{PKE.Enc}(\text{pke.pk}_{i,\overline{s_i}}, 0^{2\lambda+1}) \end{aligned}$$

Next, it samples randomness $(z_{i,b})_{i,b}$ and computes PRC encodings of each ciphertext and commitment. That is, for each i, b , it sets

$$c_{i,b} = \text{PRC.Encode}(\text{prc.pk}_{i,b}, (\text{pke.ct}_{i,b}, \text{com}_i); z_{i,b}).$$

It then computes the payload c , which contains an encoding of the message m , randomness $(z_{i,b})_{i,b}$ used for PRC encoding, the commitments and the PKE ciphertexts. The payload is

$$c = \text{PRG}(r_{00}) + \zeta \cdot \text{ECC.Encode}\left(m, (z_{i,b}, \text{pke.ct}_{i,b}, \text{com}_i)_{i,b}\right).$$

⁷In our construction, we also require an additional property from the commitment scheme - the commitment to bit 0 should not depend upon the tag.

Here, PRG is a regular pseudorandom generator (with sufficient stretch) and the output of the PRG is a number in a large field $\mathbb{F}$, and ECC is a regular error correcting code (over the binary alphabet). The output of the ECC will be interpreted as a number in $\mathbb{F}$. Looking ahead, we will argue that with overwhelming probability over the choice of ζ , once the payload is fixed, the underlying components $(m, (z_{i,b}, \text{pke.ct}_{i,b}, \text{com}_i)_{i,b})$ are fixed.

Finally, the encryption algorithm computes a signature on the payload and PRC codewords. It computes $\text{ds.}\sigma = \text{Sign}(\text{ds.sk}, (c, (c_{i,b})_{i,b}))$, and a masked version $\sigma = \text{PRG}(r_{01}) \oplus \text{ECC.Encode}(\text{ds.}\sigma)$. The final ciphertext is $(c, (c_{i,b})_{i,b}, \sigma)$.

Information rate of our PRC scheme: Our PRC scheme preserves the information rate of the underlying error correcting code. Suppose ECC maps t bit messages to $k_{\text{ecc}} \cdot t$ bits for some $k_{\text{ecc}} > 1$, and let ℓ_{msg} denote the length of the message being encoded by our PRC scheme. The payload size is $k_{\text{ecc}} \cdot (\ell_{\text{msg}} + \text{poly}(\lambda))$, since each CPA-PKE ciphertext encrypts $O(\lambda)$ bit strings, and the commitments are also $\text{poly}(\lambda)$ bits long. The CPA-PRC codewords encode $\text{poly}(\lambda)$ bit strings, and therefore have size $\text{poly}(\lambda)$. Finally, the signature is only $O(\lambda)$ bits long. Putting everything together, the ciphertext size is dominated by $k_{\text{ecc}} \cdot \ell_{\text{msg}}$ (for long messages).

Decryption. The decryption algorithm uses the secret keys corresponding to the 0^n path (i.e., $(\text{prc.sk}_{i,0}, \text{pke.sk}_{i,0})_i$), and proceeds as follows:

- *Recovering a candidate seed.*
 1. First, it decrypts each $c_{i,0}$ using $\text{prc.sk}_{i,0}$, recovering $(\text{pke.ct}_{i,0}, \text{com}_i)$. If any of the decryptions fail, it is set to a default value (but do not abort the decryption).
 2. Next, for each i , it decrypts $\text{pke.ct}_{i,0}$. If the decrypted message is of the form $1\|\text{op}_i\|\text{ds.vk}_i$ and $\text{com}_i = \text{Com}(0, \text{ds.vk}_i; \text{op}_i)$, then it guesses that $t_i = 0$, else it guesses $t_i = 1$. Again, decryption does not abort if the decryption of $\text{pke.ct}_{i,0}$ fails, it simply guesses $t_i = 1$.
- *Checking noise in the ciphertext components*
 1. Having obtained a candidate seed t , the decryption algorithm computes $(r_{00}, r_{01}, r_1, \dots, r_\lambda) = \text{HPRG}(t)$.
 2. It then recovers the (error-free) message, the randomness used for the PRC encodings, the commitments, and PKE ciphertexts from the payload. It computes $(\text{msg}, (z_{i,b}, \text{pke.ct}_{i,b}, \text{com}_i)_i) = \text{ECC.Decode}(\zeta^{-1}(\text{PRG}(r_{00}) \oplus c))$.
 3. Similarly, it recovers the error-free digital signature from σ — $\text{ds.}\sigma = \text{ECC.Decode}(\text{PRG}(r_{01}) \oplus \sigma)$.
 4. It re-encrypts each PRC ciphertext using the error-free PKE ciphertexts and commitments, and the randomness obtained from the payload. For each i, b , it computes $c'_{i,b} = \text{PRC.Encode}(\text{prc.pk}_{i,b}, (\text{pke.ct}_{i,b}, \text{com}_i); z_{i,b})$ and checks that $\Delta(c_{i,b}, c'_{i,b})$ is small (which ensures that the amount of noise in each PRC codeword is bounded).
 5. Similarly, it checks that the noise in c and σ is bounded. It then computes $c' = \text{PRG}(r_{00}) \oplus (\text{msg}, (z_{i,b}, \text{pke.ct}_{i,b}, \text{com}_i)_{i,b})$, $\sigma' = \text{PRG}(r_{01}) \oplus \text{ds.}\sigma$, and checks that $\Delta(c, c')$ and $\Delta(\sigma, \sigma')$ are small.
- *Comparing the commitments and PKE ciphertexts recovered via decryption, with the commitments and PKE ciphertexts recovered from the payload.*
 1. For each i such that t_i is set to 0, we recovered a PKE ciphertext and commitment from $c_{i,0}$. The decryption algorithm checks that these are consistent with the PKE ciphertexts and commitments in the payload. If $t_i = 0$, we check that $\text{pke.ct}_{i,0} = \text{pke.ct}_{i,0}$ and $\text{com}_i = \text{com}_i$.
- *Performing the KW checks*
 1. For each i , it decrypts $\text{pke.ct}_{i,t_i}$ using the randomness r_i , recovers $1\|v_i\|\text{ds.vk}_i$, checks that v_i is a valid opening of com_i for the bit s_i (w.r.t. ds.vk_i), and that re-encryption of $1\|v_i\|\text{ds.vk}_i$ with randomness r_i gives $\text{pke.ct}_{i,s_i}$. If any of these checks fail, it outputs $\perp$.
- *Checking the digital signature:* Finally, the decryption algorithm checks that all the verification keys ds.vk_i are equal, and that the signature $\text{ds.}\sigma$ (obtained from σ) on $(c', (c'_{i,b})_{i,b})$ is valid (with respect to the common verification key).

Security of our construction. At a high level, our proof follows the KW proof structure. Let t^* denote the seed, and ds.vk^* the verification key used in the challenge ciphertext. First, using the security of the signature scheme, we can argue that for all decryption queries not close to the challenge ciphertext, if the verification key extracted from the ciphertext is equal to ds.vk^* , then such decryption queries can be rejected. Next, we switch the commitment scheme's public parameters to make them equivocal at ds.vk^* , this follows from the security of the commitment scheme. After this, we decrypt using the secret keys $(\text{prc.sk}_{i,t_i^*}, \text{pke.sk}_{i,t_i^*})_i$. Using the perfect correctness of the PKE scheme, the perfect binding of the commitment scheme and adaptive robustness of the base PRC scheme, we argue that this switch in the decryption oracle goes unnoticed by the adversary. This indistinguishability proof is one of the key technical challenges in our proof, and we discuss this separately in the next paragraph. Assuming the decryption switch works, we then switch the PKE ciphertexts in the challenge ciphertext that are not on the t^* path. Instead of encrypting $\mathbf{0}$, we encrypt a commitment opening for bit $\overline{t_i^*}$ wrt ds.vk^* (recall, since our commitments are equivocal for tag ds.vk^* , we can produce openings for both 0 and 1). This switch follows from the CPA security of the PKE scheme since the corresponding secret keys are not used in the decryption oracle. We now switch back the decryption oracle to use $(\text{prc.sk}_{i,0}, \text{pke.sk}_{i,0})_i$. At this point, note that the HPRG seed is only used to compute the randomness $(r_{00}, r_{01}, r_1, \dots, r_\lambda)$. We can now use the HPRG security to switch all these to uniformly random, and then use the regular PRG security to make the payload c and signature component σ uniformly random. Next, we use the PRC security to replace the PRC ciphertexts $c_{i,1}$ with uniformly random strings (note that the corresponding PRC secret keys are not needed for decryption, and we already switched the payload c to a uniformly random string). At this point, the only components that are not uniformly random are the PRC ciphertexts $c_{i,0}$. We now switch the decryption oracle again, using the secret keys $(\text{prc.sk}_{i,1}, \text{pke.sk}_{i,1})_i$. Now, we can finally replace the codewords $(c_{i,0})_i$ with uniformly random strings. Therefore, all that remains is to show that if an adversary can distinguish between various modes of decryption, then this adversary can break adaptive robustness of the base PRC.

Switching the decryption oracle. Let Dec denote the regular decryption, and $\text{Dec}^\dagger$ denote the decryption along path $\mathbf{p}$. We want to argue that no p.p.t. adversary can produce a ciphertext $\text{ct} = (c, (c_{i,b})_{i,b})$ such that the outputs of Dec and $\text{Dec}^\dagger$ on ct are different (and this argument does not rely on the hiding of $\mathbf{p}$). We will prove the following statements:

1. Suppose Dec finds a candidate seed t , and finally produces a non- $\perp$ message, then $\text{Dec}^\dagger$ must also produce the same candidate t .⁸
2. If Dec and $\text{Dec}^\dagger$ find the same candidate seed t , then they must produce the same output.

Let us consider the first statement. Suppose, on the contrary, Dec uses a candidate seed t and produces a non- $\perp$ output, and $\text{Dec}^\dagger$ produces a different candidate seed $t^\dagger$ such that $t_i \neq t_i^\dagger$. As in the KW decryption switching proof, we will consider two sub-cases: $t_i = 0$ and $t_i^\dagger = 1$ and $t_i = 1$ and $t_i^\dagger = 0$.

- $t_i = 0$ and $t_i^\dagger = 1$: For the first sub-case, we will arrive at a contradiction by relying on the perfect binding of the commitments and the adaptive robustness. First, note that if $t_i \neq t_i^\dagger$, then $\mathbf{p}_i$ must be equal to 1 (otherwise the strings t and $t^\dagger$ cannot differ at the i^{th} index. Next, from the structure of our payload, it follows that with all but negligible probability over the choice of ζ , there exists at most one r_{00} , $\text{ds.x} = (\text{msg}, (z_{i,b}, \text{pke.ct}_{i,b}, \text{com}_i)_i)$ such that $c = \text{PRG}(r_{00}) + \zeta \cdot \text{ECC}(\text{ds.x})$.

Let $(\text{pke.ct}_{i,0}, \text{com}_i)$ be the ciphertext and commitment recovered from $c_{i,0}$. Since $t_i = 0$ and Dec does not output $\perp$, we get that

- $\text{com}_i = \text{com}_i$, $\text{pke.ct}_{i,0} = \text{pke.ct}_{i,0}$
- decryption of $\text{pke.ct}_{i,0}$ produces $1\|v_i\|\text{ds.vk}_i$ such that v_i is an opening of com_i for bit 0
- $\Delta(c_{i,0}, c'_{i,0})$ and $\Delta(c_{i,1}, c'_{i,1})$ are bounded

⁸We observe that the following is also true, i.e., suppose $\text{Dec}^\dagger$ finds a candidate seed t , and finally produces a non- $\perp$ message, then Dec must also produce the same candidate t .

Next, we use the fact that $t_i^\dagger$ is set to 1. This means that $c_{i,1}$ decrypts to some $(\text{pke.ct}_{i,1}^\dagger, \text{com}_{i,1}^\dagger)$ and $\text{com}_i^\dagger$ is a commitment to 1 (with respect to some $\text{ds.vk}^\dagger$). Now, if $\text{com}_i^\dagger = \text{com}_i$, then we break the perfect binding property of the commitment scheme. This is because $\text{com}_i = \text{com}_i$ is a commitment to the bit 0 (and commitments to the bit 0 do not depend on the tag), and also a commitment to the bit 1 wrt $\text{ds.vk}^\dagger$ if $\text{com}_i^\dagger = \text{com}_i$. On the other hand, if $\text{com}_i^\dagger \neq \text{com}_i$, then we break the adaptive robustness of the base PRC scheme. This is because $c_{i,1}$ decrypts to $(\dots, \text{com}_{i,1}^\dagger)$, and $c'_{i,1}$ (which is a PRC encryption of $(\text{pke.ct}_{i,1}, \text{com}_i)$ with randomness $z_{i,1}$) is close to $c_{i,1}$.

- $t_i = 1, t_i^\dagger = 0$: For the second sub-case, we will show a contradiction by relying on the perfect correctness of the PKE and the adaptive robustness of the PRC.

Next, let us consider the second case. If both Dec and $\text{Dec}^\dagger$ recover the same candidate seed t and Dec outputs $\text{non-}\perp$, then $\text{Dec}^\dagger$ must also output the same $\text{non-}\perp$ message. This again follows from the fact that c is binding to $(\text{msg}, (z_{i,b}, \text{pke.ct}_{i,b}, \text{com}_i)_i)$, and since both recover the same guess t , both modes of decryption would output msg .

2.3 Improving the error rate of our CCA-PRC construction

One limitation of the previous construction is that it cannot handle errors at a constant fraction of the positions, as this would mean that some of the PRC ciphertexts can be completely lost. To achieve better error tolerance, the decryption algorithm must be able to recover the hinting PRG seed s even if some PRC components are corrupted. A natural fix is to use an error-correcting code: rather than feeding s directly into the hinting PRG, we encode it as $t = \text{ECC.Encode}(s)$ and evaluate the HPRG on t . This way, even if several PRC ciphertexts are corrupted, the decryption algorithm can still reconstruct s and proceed with decryption. This gives us a PRC with constant rate adaptive robustness, assuming the base PRC and ECC have constant rate robustness.

Security proof. One nuance that arises is with respect to the security of hinting PRGs. Existing hinting PRG constructions prove security when the seed is sampled uniformly at random. However, it is easy to check that the proofs of security for standard hinting PRG constructions rely only on min-entropy of the input. Since ECC is a deterministic encoding, the entropy of t is identical to that of s . Hence, this modification is sound.

The security proof for the constant rate construction follows the same proof steps as the low-rate construction. The only step that requires some attention is the oracle switching step. To address this, we give a *lifting lemma* (in Lemma 12). This lemma essentially shows that if an adversary produces a ciphertext ct such that it can distinguish between Dec and $\text{Dec}^\dagger$, then there exists an index i such that $\text{p}_i = 1$, $\tilde{t}_i$ (the noisy candidate seed recovered by Dec) is not equal to the corresponding bit $t_i^\dagger$ recovered by $\text{Dec}^\dagger$, and both $c_{i,0}$ and $c_{i,1}$ have bounded noise. Note that this is all that we required for the oracle switching proof for the low-rate scheme.

3 Preliminaries

Notations: Let PPT denote probabilistic polynomial-time. We denote the set of all positive integers up to n as $[n] := \{1, \dots, n\}$. For any set S , we denote $\mathcal{P}(S)$ as the power set of S . We denote vectors/binary strings using bold font. For any two binary strings $\mathbf{x}, \mathbf{y}$, we use the notation $\mathbf{x} \parallel \mathbf{y}$ to denote $\mathbf{x}$ concatenated with $\mathbf{y}$. Here, x_i denotes the i^{th} bit of $\mathbf{x}$. For any $\mathbf{x} \in \mathbb{F}^n$, we denote $\text{Ham}(\mathbf{x})$ as the Hamming weight of $\mathbf{x}$, i.e., the number of non-zero entries in $\mathbf{x}$. And, $\text{Ham}(\mathbf{x}, \mathbf{y})$ denotes the Hamming distance between the binary strings $\mathbf{x}$ and $\mathbf{y}$, i.e., the number of indices where $\mathbf{x}$ and $\mathbf{y}$ differ. We also write $\Delta(\mathbf{x}, \mathbf{y})$ for Hamming distance. We denote $\mathcal{B}_\delta(\mathbf{x}, \mathbb{F}^n)$ as the set $\{\mathbf{x}' \in \mathbb{F}^n \mid \text{Ham}(\mathbf{x}, \mathbf{x}') \leq \delta \cdot n\}$, i.e., the set of all strings with Hamming distance of r from $\mathbf{x}$. Furthermore, we define $\mathcal{B}_\delta(\mathbb{F}^n) := \mathcal{B}_\delta(\mathbf{0}, \mathbb{F}^n)$. We denote by $\text{Ber}_\rho^{m \times n}$ the distribution over $m \times n$ matrices defined as follows: each entry is sampled independently, taking a value uniformly at random from $\mathbb{F} \setminus \{0\}$ with probability ρ , and 0 otherwise. For any finite set S , $x \leftarrow S$ denotes a uniformly random element

x from the set S . Similarly, for any distribution $\mathcal{D}$, $x \leftarrow \mathcal{D}$ denotes an element x drawn from the distribution $\mathcal{D}$.

Lemma 1 (Leftover Hash Lemma). *Let $\mathcal{D}$ be a distribution over $\mathbb{Z}_q^n$ with min-entropy k . For any $\epsilon > 0$ and $\ell \leq (k - 2\log(1/\epsilon) - O(1))/\log(q)$, the joint distribution of (C, Cs) where $C \leftarrow \mathbb{Z}_q^{\ell \times n}$ is uniformly random and $s \in \mathbb{Z}_q^n$ is drawn from the distribution $\mathcal{D}$ is ϵ close to the uniform distribution over $\mathbb{Z}_q^{\ell \times n} \times \mathbb{Z}_q^\ell$.*

3.1 Error-Correcting Codes

Definition 1 (Error-Correcting Code). *A binary error-correcting code is defined by three efficient algorithms (Setup, Encode, Decode) with the following description.*

Setup(1^γ) : *The setup algorithm takes as input the error-rate 1^γ and outputs a public parameter pp .*
Encode(pp, msg) : *The encoding algorithm takes as input a public parameter pp and a message $\text{msg} \in \{0, 1\}^{\ell_{\text{msg}}}$ and outputs a codeword $\text{cw} \in \{0, 1\}^{\ell_{\text{cw}}}$.*
Decode(pp, cw) : *The decoding algorithm takes as input a public parameter pp and a codeword $\text{cw} \in \{0, 1\}^{\ell_{\text{cw}}}$ and outputs either a message $\text{msg} \in \{0, 1\}^{\ell_{\text{msg}}}$ or $\perp$.*

An error-correcting code must satisfy the following properties.

- **Correctness:** *A code C is correct if for all $\text{pp} \leftarrow \text{Setup}(1^\gamma)$, $\text{msg} \in \{0, 1\}^{\ell_{\text{msg}}}$, $\text{Decode}(\text{pp}, \text{Encode}(\text{pp}, \text{msg})) = \text{msg}$.*
- **Robustness:** *A code C is robust for up to γ fraction of errors if for all $\text{pp} \leftarrow \text{Setup}(1^\gamma)$, $\text{msg} \in \{0, 1\}^{\ell_{\text{msg}}}$ and $\mathbf{e} \in \{0, 1\}^{\ell_{\text{cw}}}$ such that $\text{Ham}(\mathbf{e})$ is less than $\gamma \times \ell_{\text{cw}}$, it holds that $\text{Decode}(\text{pp}, \text{Encode}(\text{pp}, \text{msg}) + \mathbf{e}) = \text{msg}$.*

The information rate of a code is the ratio between ℓ_{cw} and ℓ_{msg} .

Theorem 2 ([JST21]). *Fix $\gamma \in (0, 1/4)$. There exist binary error-correcting codes that are robust to error-rate γ with constant information rate.*

3.2 Pseudorandom Functions

Definition 2 (Pseudorandom Functions). *A family of pseudorandom functions is a pair of algorithms $\text{PRF} = (\text{Setup}, \text{Eval})$ with the following description.*

- **Setup(1^λ)** : *The key generation algorithm is a randomized algorithm that takes as input the security parameter 1^λ and outputs a key k .*
- **Eval(k, x)** : *The evaluation algorithm is a deterministic algorithm that takes as input a key k and x and outputs y .*

Definition 3 (Security). *The PRF scheme PRF is secure if for all PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$,*

$$|\Pr[\mathcal{A}^{\text{Eval}(k, \cdot)}(1^\lambda) = 1 : k \leftarrow \text{Setup}(1^\lambda)] - \Pr[\mathcal{A}^{R(\cdot)}(1^\lambda) = 1 : R \leftarrow \mathcal{U}_\lambda]| \leq \text{negl}(\lambda)$$

where $\mathcal{U}_\lambda$ is the set of all functions from the domain space of Eval to the codomain space of Eval .

3.3 Equivocal Statistically Binding Tag-based Set Commitments

We will follow the abstraction of set commitments as defined in [HKW20]. Here, the commitment algorithm takes as input a string $S \in \{0, 1\}^N$ which represents the set $\{(i, S_i) : i \in [N]\}$.

Syntax. A tag-based set-commitment scheme for tag space $\mathcal{T}$, with commitment space $\mathcal{C}$ and opening space $\mathcal{O}$ is described by algorithms $\text{Com} = (\text{Setup}, \text{Commit}, \text{Open}, \text{Equiv})$ with the following description.

Setup($1^\lambda, 1^N$): The setup algorithm is a randomized algorithm that takes as input a unary representation of the security parameter λ and the set size N , and outputs public parameters pp .

Commit(pp, S, tag): The commit algorithm is a randomized algorithm that takes as input a public parameter pp , a set-representation $S \in \{0, 1\}^N$ and $\text{tag} \in \mathcal{T}$, and outputs a commitment vector $(\text{com}_1, \dots, \text{com}_N) \in \mathcal{C}^N$ and an opening vector $(\text{op}_1, \dots, \text{op}_N) \in \mathcal{O}^N$.

Open($\text{pp}, b, \text{com}_i, \text{tag}, \text{op}_i$): The opening algorithm is a deterministic algorithm that takes as input public parameters pp , a bit b , a commitment $\text{com}_i \in \mathcal{C}$, $\text{tag} \in \mathcal{T}$ and an opening $\text{op}_i \in \mathcal{O}$ and outputs $\top/\perp$.

Equiv($1^\lambda, 1^N, \text{tag}^*$): The alternate setup algorithm takes as input the security parameter λ , the set size N and tag^* , and outputs public parameters pp together with commitment vector $(\text{com}_i)_{i \in [N]}$ and openings $(\text{op}_{i,0}, \text{op}_{i,1})_{i \in [N]}$.

An equivocal tag-based commitment scheme Com is *correct* if the following conditions hold:

- for all $\lambda, N \in \mathbb{N}$, $\text{pp} \leftarrow \text{Setup}(1^\lambda, 1^N)$, for all $\text{tag} \in \mathcal{T}$ and $S \in \{0, 1\}^N$, for all $((\text{com}_i, \text{op}_i)_{i \in [N]}) \leftarrow \text{Commit}(\text{pp}, S, \text{tag})$, for all $i \in [N]$, $\text{Open}(\text{pp}, S_i, \text{com}_i, \text{tag}, \text{op}_i) = \top$.
- for all $\lambda, N \in \mathbb{N}$, $\text{tag}^* \in \mathcal{T}$, for all $((\text{pp}, (\text{com}_i, \text{op}_{i,0}, \text{op}_{i,1})_i) \leftarrow \text{Equiv}(1^\lambda, 1^N, \text{tag}^*)$, for all $i \in [N]$ and $b \in \{0, 1\}$, $\text{Open}(\text{pp}, b, \text{com}_i, \text{tag}^*, \text{op}_{i,b}) = \top$.

Definition 4 (Statistical Binding). A tag-based commitment scheme Com is *statistical binding* if for any adversary $\mathcal{A}$ and polynomial p , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$ and $N = p(\lambda)$, the advantage of $\mathcal{A}$ in the following security game is bounded by $\text{negl}(\lambda)$:

1. $\mathcal{A}$ chooses tag^* and sends it to the challenger.
2. The challenger computes $((\text{pp}, (\text{com}_i, \text{op}_{i,0}^*, \text{op}_{i,1}^*)_i) \leftarrow \text{Equiv}(1^\lambda, 1^N, \text{tag}^*)$ and sends pp to $\mathcal{A}$.
3. $\mathcal{A}$ outputs $\text{tag} \neq \text{tag}^*, \text{com}, \text{op}_0, \text{op}_1$ and wins if $\text{Open}(\text{pp}, b, \text{com}, \text{tag}, \text{op}_b) = \top$ for both $b \in \{0, 1\}$.

Definition 5 (Equivocality: Indistinguishability of Setup and Equiv). A tag-based commitment scheme Com is said to be *equivocal* if for any PPT adversary $\mathcal{A}$, any $\text{tag}^* \in \mathcal{T}$ and any $S^* \in \{0, 1\}^N$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$ and n that is polynomial in λ , the advantage of $\mathcal{A}$ is bounded by $\text{negl}(\lambda)$.

1. The challenger samples $b \leftarrow \{0, 1\}$.
If $b = 0$, it generates $\text{pp} \leftarrow \text{Setup}(1^\lambda)$ and $((\text{com}_i, \text{op}_i)_{i \in [N]}) \leftarrow \text{Commit}(\text{pp}, S^*, \text{tag}^*)$. It sends $(\text{pp}, (\text{com}_i, \text{op}_i)_{i \in [N]})$ to $\mathcal{A}$.
If $b = 1$ it generates $((\text{pp}, (\text{com}_i, \text{op}_{i,0}, \text{op}_{i,1})_i) \leftarrow \text{Equiv}(1^\lambda, 1^N, \text{tag}^*)$. It sends $(\text{pp}, (\text{com}_i, \text{op}_{i,S_i})_{i \in [N]})$ to $\mathcal{A}$.
2. $\mathcal{A}$ outputs $b' \in \{0, 1\}$ and wins if $b = b'$, and the advantage of $\mathcal{A}$ is $\Pr[\mathcal{A} \text{ wins}] - 1/2$.

Definition 6 (Pseudorandomness). A tag-based commitment scheme Com is said to have *pseudorandom commitments* if, for any PPT adversary $\mathcal{A}$ and polynomial p , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $n = p(\lambda)$, $\mathcal{A}$ wins the following game with at most $\text{negl}(\lambda)$ advantage.

1. The challenger samples $b \leftarrow \{0, 1\}$. It samples $\text{pp} \leftarrow \text{Setup}(1^\lambda, 1^N)$ and sends pp to the adversary.
2. The adversary makes polynomially many queries. In each query, it sends a set $S^{(i)}$ and $\text{tag}^{(i)}$.
– If $b = 0$, challenger samples $((\text{com}_j^{(i)}, \text{op}_j^{(i)})_{j \in [N]}) \leftarrow \text{Commit}(\text{pp}, S^{(i)}, \text{tag}^{(i)})$ and sends $(\text{com}_j^{(i)})_{j \in [N]}$ to $\mathcal{A}$.
– If $b = 1$, challenger samples $u_j^{(i)} \leftarrow \mathcal{C}$ uniformly at random and sends $(u_j^{(i)})_{j \in [N]}$ to $\mathcal{A}$.
3. $\mathcal{A}$ finally sends its guess b' and wins if $b = b'$. The advantage of $\mathcal{A}$ is $\Pr[b = b'] - 1/2$.

Although tag-based commitment schemes were formally introduced in [HKW20], the concept was employed by Koppula and Waters [KW19] in their construction of chosen-ciphertext secure encryption schemes. In our work, we rely on the construction from [KW19] rather than the later construction in [HKW20]. The reason is that in [HKW20] the size of the commitment grows with the size of the set S , whereas in [KW19] the commitment size depends only on the fixed parameter N . This fixed-size property is crucial for our applications.

Theorem 3 ([Nao90], [KW19]). *Assuming the existence of secure PRGs, there exists equivocal, statistically binding, tag-based commitment schemes with pseudorandom commitments for tag space $\{0, 1\}^\lambda$. The commitment space is $\{0, 1\}^{4\ell}$ and the opening space is $\{0, 1\}^\lambda$. Furthermore, for any pp sampled by Setup or Equiv, for any $S \in \{0, 1\}^N$, for any tag, tag', for any i such that $S_i = 0$, the following distributions are identical:*

$$\left\{ \text{com}_i : (\text{com}_j, \text{op}_j)_{j \in [N]} \leftarrow \text{Commit}(\text{pp}, S, \text{tag}) \right\} \equiv \left\{ \text{com}_i : (\text{com}_j, \text{op}_j)_{j \in [N]} \leftarrow \text{Commit}(\text{pp}, S, \text{tag}') \right\}$$

3.4 Public-key Pseudorandom Codes

A public-key PRC with κ -correctness α -robustness is described by three algorithms (Setup, Encode, Decode), satisfying the following criteria:

Syntax: There exist functions $s, p, \ell_{\text{cw}}, \ell_{\text{msg}}, r, \ell_{\text{Enc}}, r'' : \mathbb{N} \rightarrow \mathbb{N}$ such that for all $\lambda \in \mathbb{N}$, **Setup** : $1^\lambda \times \{0, 1\}^{r''(\lambda)} \rightarrow \{0, 1\}^{p(\lambda)} \times \{0, 1\}^{s(\lambda)}$, **Encode** : $\{0, 1\}^{p(\lambda)} \times \{0, 1\}^{\ell_{\text{msg}}(\lambda)} \times \{0, 1\}^{\ell_{\text{Enc}}(\lambda)} \rightarrow \{0, 1\}^{\ell_{\text{cw}}(\lambda)}$, and **Decode** : $\{0, 1\}^{s(\lambda)} \times \{0, 1\}^{\ell_{\text{cw}}(\lambda)} \times \{0, 1\}^{r(\lambda)} \rightarrow \{0, 1\}^{\ell_{\text{msg}}(\lambda)} \cup \perp$.

Error correction, or robustness: For any $\lambda \in \mathbb{N}$, message $\text{msg} \in \{0, 1\}^{\ell_{\text{msg}}(\lambda)}$,

$$\Pr \left[\text{Decode}(\text{sk}, \text{c} + \text{e}) \neq \text{msg} \mid \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{Setup}(1^\lambda) \\ \text{c} \leftarrow \text{Encode}(\text{pk}, \text{msg}) \\ \text{e} \leftarrow \mathcal{B}_\alpha^{\ell_{\text{cw}}} \end{array} \right] \leq \kappa(\lambda).$$

where the probability is over the randomness of Setup, Encode and error e .

Pseudorandomness: For any polynomial-time adversary $\mathcal{A}$,

$$\Pr_{(\text{pk}, \text{sk}) \leftarrow \text{Setup}(1^\lambda)} [\mathcal{A}^{\text{Opk}(\cdot)}(1^\lambda, \text{pk}) = 1] - \Pr_{(\text{pk}, \text{sk}) \leftarrow \text{Setup}(1^\lambda)} [\mathcal{A}^{U(1^\lambda, \cdot)}(1^\lambda, \text{pk}) = 1] \leq \text{negl}(\lambda),$$

where Opk on any (even previously queried) input msg samples a random $r \in \{0, 1\}^{\ell_{\text{Enc}}(\lambda)}$ and outputs $\text{Encode}(\text{pk}, \text{msg}; r)$ and U denotes an oracle that on any (even previously queried) input returns a fresh uniform string in $\{0, 1\}^{\ell_{\text{cw}}}$.

CCA-Security: For any stateful polynomial time adversary $\mathcal{A}$, we require the following to hold:

$$\Pr \left[b' = b \mid \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{Setup}(1^\lambda) \\ m \leftarrow \mathcal{A}^{\text{Decode}(\text{sk}, \cdot)}(\text{pk}) \\ b \leftarrow \{0, 1\} \\ \text{if } b = 0 : \text{ct} \leftarrow \text{Encode}(\text{pk}, m) \\ \text{if } b = 1 : \text{ct} \leftarrow \{0, 1\}^{\ell_{\text{cw}}(\lambda)} \\ b' \leftarrow \mathcal{A}^{\mathcal{O}(\cdot)}(\text{ct}) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda).$$

where

$$\mathcal{O}(\text{ct}') = \begin{cases} \text{Decode}(\text{sk}, \text{ct}') & \text{if } \forall \text{e} : \text{ct}' \oplus \text{e} \neq \text{ct}, \text{ where } \text{Ham}(\text{e}) \leq \alpha \ell_{\text{cw}}(\lambda) \\ \perp & \text{otherwise} . \end{cases}$$

where $\ell_{\text{cw}}(\lambda)$ is the length of the ciphertext and α is the robustness of the PRC.

The information rate of a PRC is the ratio between ℓ_{cw} and ℓ_{msg} .

Definition 7 (Soundness). A public-key PRC with codeword space $\mathcal{Y}_\lambda$ is sound if $\Pr [\text{Decode}(\text{sk}, \mathbf{u}) \neq \perp \mid (\text{pk}, \text{sk}) \leftarrow \text{Setup}(1^\lambda)] \leq \text{negl}(\lambda)$. For the binary syntax above, $\mathcal{Y}_\lambda = \{0, 1\}^{\ell_{\text{cw}}(\lambda)}$.

Definition 8 (Adaptive Robustness). A stronger notion of robustness called adaptive robustness is defined via the following game parameterized by a constant $\alpha \in (0, 1/2)$.

1. The challenger samples $(\text{pk}, \text{sk}) \leftarrow \text{Setup}(1^\lambda)$ and sends pk to the adversary.
2. The adversary sends (m, r, x) to the challenger.
3. The challenger computes $c = \text{Encode}(\text{pk}, m; r)$ and $m' \leftarrow \text{Decode}(\text{sk}, x)$. If $\text{Ham}(c, x) \leq \alpha n$ and $m \neq m'$, then output the adversary wins. Otherwise, it loses.

We say that a public-key pseudorandom code PRC is adaptively α -robust if, for any PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $\Pr[\mathcal{A} \text{ wins in the above game}] \leq \text{negl}(\lambda)$.

Theorem 4 ([AAC⁺25]). Assuming $2^{O(\sqrt{\lambda})}$ hardness of LPN, there exists an adaptively robust public key PRC with error rate $\delta/8$ for any $\delta \in (0, 1/4)$.

If, in addition, the adversary is given the secret key sk in the first step of the above game, then the notion is referred to as *strong adaptive robustness* [CGG⁺25].

3.5 Hinting Pseudorandom Generator

A hinting PRG (HPRG) is a pair of algorithms $(\text{Setup}, \text{Eval})$ such that:

$\text{Setup}(1^\lambda, 1^\ell)$: The generation algorithm is a randomized algorithm that takes as input a unary representation of the security parameter λ , a length parameter ℓ and outputs a public parameter pp and input length $n = n(\lambda, \ell)$.

$\text{Eval}(\text{pp}, s, i)$: The evaluation algorithm is a deterministic algorithm that takes as input a public parameter pp , a seed $s \in \{0, 1\}^n$ and an index $i \in \{0\} \cup [n]$ and outputs a bit string $y \in \{0, 1\}^\ell$.

Security: A hinting PRG scheme $(\text{Setup}, \text{Eval})$ is said to be secure if for any PPT adversary $\mathcal{A}$, polynomial $\ell(\cdot)$ there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, the following holds:

$$\left| \Pr \left[\beta \leftarrow \mathcal{A} \left(\text{pp}, y_0^\beta, \{y_{i,b}^\beta\} \right) \mid \begin{array}{l} (\text{pp}, n) \leftarrow \text{Setup}(1^\lambda, 1^\ell), s \leftarrow \{0, 1\}^n, \\ \beta \leftarrow \{0, 1\}, y_0^0 \leftarrow \{0, 1\}^\ell, y_0^1 = \text{Eval}(\text{pp}, s, 0), \\ y_{i,b}^0 \leftarrow \{0, 1\}^\ell \forall i \in [n], b \in \{0, 1\}, \\ y_{i,s_i}^1 = \text{Eval}(\text{pp}, s, i), y_{i,\overline{s_i}}^1 \leftarrow \{0, 1\}^\ell \forall i \in [n] \end{array} \right] - \frac{1}{2} \right| \leq \text{negl}(\lambda)$$

Theorem 5 ([KW19, AP22]). Assuming hardness of $\mathcal{X}$ where $\mathcal{X} \in \{\text{LWE}, \text{CDH}, \text{DDH}, \text{LHS}\}$, there exist secure hinting PRGs.

We additionally require a security notion that we refer to as entropic security.

Entropic Security: A hinting PRG scheme $(\text{Setup}, \text{Eval})$ is said to be k -entropic secure if for any PPT adversary $\mathcal{A}$, polynomial $\ell(\cdot)$ there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, for all distributions $\mathcal{D}$ over $\{0, 1\}^n$ such that $\mathcal{H}_\infty(\mathcal{D}) \geq k$ the following holds:

$$\left| \Pr \left[\beta \leftarrow \mathcal{A} \left(\text{pp}, y_0^\beta, \{y_{i,b}^\beta\} \right) \mid \begin{array}{l} (\text{pp}, n) \leftarrow \text{Setup}(1^\lambda, 1^\ell), s \leftarrow \mathcal{D}, \\ \beta \leftarrow \{0, 1\}, y_0^0 \leftarrow \{0, 1\}^\ell, y_0^1 = \text{Eval}(\text{pp}, s, 0), \\ y_{i,b}^0 \leftarrow \{0, 1\}^\ell \forall i \in [n], b \in \{0, 1\}, \\ y_{i,s_i}^1 = \text{Eval}(\text{pp}, s, i), y_{i,\overline{s_i}}^1 \leftarrow \{0, 1\}^\ell \forall i \in [n] \end{array} \right] - \frac{1}{2} \right| \leq \text{negl}(\lambda)$$

Remark 2 (On the instantiation of Entropic HPRGs). Both the CDH and LWE based constructions of hinting PRGs given by Koppula and Waters [KW19] remain secure even when the seed is not uniform but only has sufficiently high min-entropy. The only point in their analysis where the uniformity of the seed s is invoked is in a hybrid that applies a randomness-extraction style argument. In the CDH instantiation, this hybrid replaces $h = \prod_{i=1}^n g_{i,s_i}$ with a uniformly random group element h , and statistical closeness is justified using the Leftover Hash Lemma (see Lemma 1). Likewise, in the LWE construction, the analogous hybrid replaces $a = \sum_{i=1}^n a_{i,s_i}$ with a uniform vector a .

This argument remains valid verbatim when the seed s is sampled from any distribution s of sufficiently large min-entropy. Specifically, if $H_\infty(s) \geq m + 2 \log(1/\varepsilon)$, where m is the number of bits that must be statistically close to uniform in the hybrid, then the Leftover Hash Lemma guarantees that the corresponding hybrid remains indistinguishable up to error ε . Consequently, by ensuring that the seed distribution provides this entropy margin, every hybrid in the Koppula-Waters proof continues to hold unchanged. Since both the CDH and LWE based instantiations in [KW19] implement the same extractor-like structure required for this step, their constructions immediately yield an *entropic* hinting PRG.

In the DDH-based construction of the hinting PRG by Alapati and Patranabis [AP22], the only place where the uniformity of the seed s is required is in establishing the pseudorandomness property of the generator. Somewhat surprisingly, their proof of the *hinting* security itself does not rely on any stronger guarantee about the distribution of s , such as full uniformity. In particular, they invoke an application of the Leftover Hash Lemma to argue that $([M], [Ms]) \approx_c ([M], [u])$, where u is uniformly distributed. Consequently, if the uniform seed is replaced by any distribution D with sufficiently high min-entropy so that the above extraction step continues to hold, the pseudorandomness property remains valid. With the parameters adjusted to satisfy this entropy requirement, their construction therefore yields an *entropic hinting PRG*.

Looking ahead, in our construction we will use an error-correction code applied on a randomly sampled k -bit string as our distribution $\mathcal{D}$. Given the correctness of the decoder of the error-correcting code, we will obtain as an output a string of length n which has at least min-entropy k . This fulfills the entropy requirement in the above definition.

3.6 Public Key Encryption

Definition 9 (Public-Key Encryption). A public-key encryption scheme (PKE) is a triple of PPT algorithms $\text{PKE} = (\text{Setup}, \text{Enc}, \text{Dec})$:

$\text{Setup}(1^\lambda, 1^n)$: The setup algorithm takes as input a unary representation of the security parameter λ and the message length n and outputs a public key pk and secret key sk .

$\text{Enc}(\text{pk}, m)$: The encryption algorithm takes as input a public key pk and a message $m \in \{0, 1\}^n$ and outputs a ciphertext ct .

$\text{Dec}(\text{sk}, \text{ct})$: The decryption algorithm takes as input a secret key sk and a ciphertext ct and outputs either a message $m \in \{0, 1\}^n$ or $\perp$.

A scheme PKE is correct if there exists a function $\epsilon(\lambda) = \text{negl}(\lambda)$ such that for all $\lambda \in \mathbb{N}, n \in \mathbb{N}, m \in \{0, 1\}^n$ and pk, sk in the support of Setup ,

$$\Pr[\text{Dec}(\text{sk}, \text{Enc}(\text{pk}, m)) = m] = 1 - \epsilon(\lambda),$$

where the probability is taken over the randomness of Enc . We say that PKE is perfectly correct if $\epsilon(\lambda) = 0$. We say that a scheme is almost perfectly correct if there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda, n \in \mathbb{N}$, for a $1 - \text{negl}(\lambda)$ fraction of keys pk, sk in the support of Setup , for all $m \in \{0, 1\}^n$, it holds that

$$\Pr[\text{Dec}(\text{sk}, \text{Enc}(\text{pk}, m)) = m] = 1,$$

where the probability is taken over the randomness of Enc .

Theorem 6 ([DNR04b]). Assuming the existence of PRGs, correct PKE implies almost perfectly correct PKE.

Definition 10 (Decryption using Randomness). A PKE scheme $\text{PKE} = (\text{Setup}, \text{Enc}, \text{Dec})$ is decryptable using encryption randomness if there exists an algorithm Recover that takes as input the public key pk , a ciphertext ct and a randomness r and outputs a message m . This algorithm has the following property:

$$\text{Recover}(\text{pk}, \text{Enc}(\text{pk}, m; r), r) = m$$

for all $\lambda \in \mathbb{N}, n \in \mathbb{N}, m \in \{0, 1\}^n$ and pk, sk in the support of Setup .

Next, we define the IND-CPA security notion for public-key encryption.

Definition 11 ([GM84]). A PKE scheme is said to be IND-CPA secure if for all stateful PPT adversaries $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}, n \in \mathbb{N}$,

$$\Pr \left[b' = b \mid \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{Setup}(1^\lambda, 1^n) \\ (m_0, m_1) \leftarrow \mathcal{A}(\text{pk}) \\ b \leftarrow \{0, 1\} \\ \text{ct} \leftarrow \text{Enc}(\text{pk}, m_b) \\ b' \leftarrow \mathcal{A}(\text{ct}) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda).$$

3.7 Digital Signatures

A one-time signature scheme consists of the following algorithms:

Setup(1^λ): The setup algorithm takes as input a unary representation of the security parameter λ and outputs a verification key vk and a signing key sk .

Sign(sk, m): The signing algorithm takes as input sk and a message $m \in \{0, 1\}^*$ and outputs a signature $\sigma \in \{0, 1\}^n$, where $n = n(\lambda)$ is the signature length.

Verify(vk, m, σ): The verification algorithm takes as input vk , m , σ and outputs 1 (accept) or 0 (reject).

Correctness. The scheme is correct if for all $\lambda, \ell \in \mathbb{N}$, all $(\text{vk}, \text{sk}) \leftarrow \text{Setup}(1^\lambda)$, all $m \in \{0, 1\}^\ell$, and $\sigma \leftarrow \text{Sign}(\text{sk}, m)$, we have

$$\text{Verify}(\text{vk}, m, \sigma) = 1,$$

except with negligible probability in λ over the coins of Setup and Sign . If the probability is 1, then the scheme is perfect correct.

Definition 12 (One-Time Strong Unforgeability). Consider the following game between a challenger and an adversary $\mathcal{A}$:

1. The challenger samples $(\text{vk}, \text{sk}) \leftarrow \text{Setup}(1^\lambda)$ and gives vk to $\mathcal{A}$.
2. $\mathcal{A}$ makes one signing query $m \in \{0, 1\}^\ell$ and receives $\sigma \leftarrow \text{Sign}(\text{sk}, m)$.
3. $\mathcal{A}$ outputs (m^*, σ^*) . It wins if $\text{Verify}(\text{vk}, m^*, \sigma^*) = 1$ and $\sigma^* \neq \sigma$.

The scheme is one-time strongly unforgeable if for all PPT $\mathcal{A}$, $\Pr[\mathcal{A} \text{ wins the above game}] \leq \text{negl}(\lambda)$.

Theorem 7 ([Lam79]). Assuming the existence of one-way functions and collision-resistant hash functions, there exists a one-time strong unforgeable digital signature scheme with $|\sigma| = O(\lambda)$.

4 Candidate Robust PRC from Noisy-McEliece Type Assumptions

In this section, we propose a new family of hardness assumptions which give rise to adaptively robust pseudorandom codes. Recall the decisional McEliece assumption: let $\mathbf{G} \in \mathbb{F}_2^{m \times n}$ be a generator matrix of a binary Goppa code, let $\Pi \leftarrow S_m$ ⁹ be a uniformly random permutation on m elements, and let $\mathbf{T} \leftarrow \text{GL}_n(\mathbb{F}_2)$ be a uniformly random invertible matrix in $\mathbb{F}_2^{n \times n}$. The decisional McEliece assumption [McE78, DDMN12] then posits that the matrix

$$\mathbf{A} = \Pi \cdot \mathbf{G} \cdot \mathbf{T}$$

is pseudorandom. Binary Goppa codes can only decode from a sublinear number of Hamming errors, specifically $O(n/\log(n))$ Hamming errors. Hence a tremendous research effort has been devoted to finding alternatives for binary Goppa codes in this construction. Notoriously, though, virtually all attempts to replace the binary Goppa code with other families of efficiently decodable linear codes have turned out to be insecure [Nie86, SS92, vT94, CC98, LS02, BLP08, FOPT10, GOT12, CGGU⁺14, COTGU15, COT16, FOP⁺16, CMCP17, KM23, CMT23, BC24, Mce26]. Furthermore, even the classical McEliece scheme was recently targeted by an algebraic distinguishing attack [Ran25]. A fundamental reason for the presence of such algebraic attacks is that the generator matrix $\mathbf{A}$ shares many algebraic invariants with the generator $\mathbf{G}$. We observe that all currently known algebraic attacks against McEliece-type assumptions fail if we only provide the matrix $\mathbf{A}$ masked by a suitable amount of Bernoulli noise. Indeed, the problem of solving algebraic systems of equations is mostly considered in the presence of exact equations. Hence, we propose the following general assumptions.

Assumption 1 (McEliece with Noise). *Let $\mathbb{F}$ be a finite field, let $m = m(\lambda)$, $k = k(\lambda)$, $n = n(\lambda)$ with $m \geq k \geq 2n > 0$ and $\ell = \ell(\lambda)$ be integers that are polynomial in λ , and let $\rho = \rho(\lambda) = O(1/\sqrt{n})$ and $\epsilon \in [0, 1]$ be a constant. Further let $(\mathbf{C}_\kappa)_{\kappa \in K}$, where $\mathbf{C}_\kappa \subseteq \mathbb{F}^m$, be a family of linear codes of length m and dimension k specified by a set $K = K(\lambda)$. For each $\kappa \in K$, let $\mathbf{G}_\kappa \in \mathbb{F}^{m \times k}$ be a generator matrix for $\mathbf{C}_\kappa$. Let $\Pi \leftarrow S_m$, $\kappa \leftarrow K$, let $\mathbf{T} \leftarrow (\mathbb{F}^{k \times 2n})_{\text{full}}$ be a uniformly random full-column-rank matrix, and let $\mathbf{E} \leftarrow \text{Ber}_\rho^{m \times 2n}$. Let*

$$\mathbf{A} = \Pi \cdot \mathbf{G}_\kappa \cdot \mathbf{T} + \mathbf{E}.$$

Now parse $\mathbf{A} = (\mathbf{A}_1 \| \mathbf{A}_2)$ where $\mathbf{A}_1, \mathbf{A}_2 \in \mathbb{F}^{m \times n}$. Further let $\mathbf{s}_1, \dots, \mathbf{s}_\ell \leftarrow \mathcal{B}_\rho(\mathbb{F}^n)$ and $\mathbf{e}_1, \dots, \mathbf{e}_\ell \leftarrow \mathcal{B}_\epsilon(\mathbb{F}^m)$. Then we conjecture that

$$(\mathbf{A}, \mathbf{A}_1 \mathbf{s}_1 + \mathbf{e}_1, \dots, \mathbf{A}_1 \mathbf{s}_\ell + \mathbf{e}_\ell) \approx_c (\mathbf{A}, \mathbf{u}_1, \dots, \mathbf{u}_\ell),$$

where the $\mathbf{u}_1, \dots, \mathbf{u}_\ell \leftarrow \mathbb{F}^m$ are chosen uniformly at random. That is, we conjecture that the $\mathbf{A}_1 \mathbf{s}_1 + \mathbf{e}_1, \dots, \mathbf{A}_1 \mathbf{s}_\ell + \mathbf{e}_\ell$ are pseudorandom given $\mathbf{A}$.

We remark that if a matrix $\mathbf{T}$ has t columns and $k \geq t + \lambda$, then a uniformly random matrix $\mathbf{T} \leftarrow \mathbb{F}^{k \times t}$ has full column rank except with negligible probability. In this case, we can drop the requirement that $\mathbf{T}$ be sampled conditioned on full column rank.

The choice of the code family $(\mathbf{C}_\kappa)_\kappa$ is part of the security assumption: the assumptions below should be read for explicitly specified, efficiently decodable families, not as consequences of generic coding-theoretic parameters. In particular, high dual distance is only a necessary sanity check against simple distinguishers, not a hardness guarantee. With this caveat, and for families that do not have an unusually small dual distance, we can, from a cryptanalytic perspective, make the following somewhat stronger assumption, which, similar to the decisional McEliece assumption, requires that the matrix $\mathbf{A} = \Pi \cdot \mathbf{G}_\kappa \cdot \mathbf{T} + \mathbf{E}$ itself is pseudorandom.

Assumption 2 (Strong McEliece with Noise). *Let $\mathbb{F}$ be a finite field, let $m = m(\lambda)$, $k = k(\lambda)$, $n = n(\lambda)$ with $m \geq k \geq n > 0$ be integers that are polynomial in λ , and let $\rho = \rho(\lambda) = O(1/\sqrt{n})$. Further let $(\mathbf{C}_\kappa)_{\kappa \in K}$, where $\mathbf{C}_\kappa \subseteq \mathbb{F}^m$, be a family of linear codes of length m and dimension k specified by a set $K = K(\lambda)$. For each $\kappa \in K$, let $\mathbf{G}_\kappa \in \mathbb{F}^{m \times k}$ be a generator matrix for $\mathbf{C}_\kappa$. Let $\Pi \leftarrow S_m$, $\kappa \leftarrow K$, let $\mathbf{T} \leftarrow (\mathbb{F}^{k \times n})_{\text{full}}$ be a uniformly random full-column-rank matrix, and let $\mathbf{E} \leftarrow \text{Ber}_\rho^{m \times n}$. Then we conjecture that the matrix*

$$\mathbf{A} = \Pi \cdot \mathbf{G}_\kappa \cdot \mathbf{T} + \mathbf{E}$$

⁹We denote the symmetric group on $\{1, \dots, m\}$ by S_m .

is pseudorandom in $\mathbb{F}^{m \times n}$. That is, for any PPT distinguisher $\mathcal{D}$ it holds that $|\Pr[\mathcal{D}(\Pi \cdot \mathbf{G}_\kappa \cdot \mathbf{T} + \mathbf{E}) = 1] - \Pr[\mathcal{D}(\mathbf{U}) = 1]| < \text{negl}(\lambda)$, where $\mathbf{U} \leftarrow \mathbb{F}^{m \times n}$ is uniformly random.

We do not present the strong McEliece with Noise assumption as a new route to ordinary public-key encryption. In our use of this assumption, it is combined with the standard $1/\sqrt{n}$ -noise LPN assumption to obtain the McEliece with Noise assumption above, and LPN is already known to imply public-key encryption [Ale03].

Candidate code families and limitations. We view the noisy McEliece assumptions as candidate assumption families rather than as fully cryptanalyzed concrete proposals. A plausible instantiation should use a code family with efficient decoding, good relative distance, sufficiently large dual distance, and no known algebraic structure that remains exploitable after coordinate permutation and the addition of noise. Our main candidate is obtained by taking suitable algebraic-geometric codes over $\mathbb{F}_{16}$, for example from the Garcia–Stichtenoth tower, and expanding them to binary codes. Over $\mathbb{F}_{16}$, this tower has genus-to-length ratio tending to $1/3$; by the standard AG-code bounds, choosing a divisor of relative degree $3/4$ yields asymptotically good families with rate $R \geq 3/4 - 1/3 = 5/12$, relative distance $\delta \geq 1/4$, dual-distance bound $\delta^\perp \geq 3/4 - 2/3 = 1/12$, and efficient unique decoding from half the designed distance, i.e., a $1/8$ fraction of errors [GS95, Sti09]. After binary expansion these constants degrade by at most a constant factor.

Concretely, for an $\mathbb{F}_{16}$ -linear code $\mathbf{C}$ one can write each symbol as a vector in $\mathbb{F}_2^4$, optionally composing the embedding in each coordinate with a public random invertible $\mathbb{F}_2$ -linear map. The McEliece permutation is then applied bit-wise rather than symbol-wise. This is intended to obscure the extension-field structure used by Schur-product and square-code attacks against algebraic-code variants of McEliece [CGGU⁺14, CMCP17]. A second natural cryptanalytic target is given by the binary Raw Reed–Solomon codes of [SKS19]. These codes have very strong distance properties and polynomially large dual distance, and they share the same high-level rationale of viewing an algebraic code over an extension field as a binary code. However, because they remain Reed–Solomon-derived, they are close to regimes where square-code and Schur-product distinguishers are especially relevant. We therefore regard Raw Reed–Solomon codes as an important target for further cryptanalysis rather than as our primary proposed instantiation. More generally, further work is needed to analyze whether noisy permuted binary expansions resist adaptations of known algebraic attacks, including the attacks and parameter regimes discussed in [CGG⁺25].

One basic necessary condition is that the dual code of $\mathbf{C}_\kappa$ does not contain non-zero vectors of very low Hamming weight. In particular, if there are vectors of constant Hamming weight in the dual, then a simple polynomial-time distinguisher exists. To avoid this basic vulnerability, we require the codes used in any proposed instantiation to have high dual distance, following [CGG⁺25], who impose a similar condition to guarantee the absence of low-weight parity constraints and to rule out low-degree statistical or algebraic distinguishers in their assumptions. For completeness, we briefly describe the attack that arises when this condition is not enforced.

Assume that the dual distance $\bar{d}$ of the code generated by $\mathbf{G}_\kappa$ is very small (e.g. constant). Let $\mathbf{H}$ be a parity check matrix for this code, i.e., it holds that $\mathbf{H} \cdot \mathbf{G}_\kappa = 0$. Then there is a vector $\mathbf{t}$ in the row-span of $\mathbf{H}$ with Hamming weight $\bar{d}$. The attacker does not know $\mathbf{H}$ or $\mathbf{t}$, but can enumerate all public-coordinate vectors $\mathbf{h} \in \mathbb{F}^m$ of Hamming weight $c = \bar{d}$. Since Π is a permutation, this enumeration includes $\mathbf{h} = \mathbf{t} \cdot \Pi^{-1}$ and takes time m^c for constant c . For this vector, the attacker calculates:

$$\begin{aligned} \mathbf{h} \cdot \mathbf{A} &= \mathbf{h} \cdot (\Pi \cdot \mathbf{G}_\kappa \cdot \mathbf{T} + \mathbf{E}) \\ &= \mathbf{h} \cdot \Pi \cdot (\mathbf{G}_\kappa \cdot \mathbf{T} + \Pi^{-1} \cdot \mathbf{E}) \\ &= \mathbf{t} \cdot (\mathbf{G}_\kappa \cdot \mathbf{T} + \Pi^{-1} \cdot \mathbf{E}) \\ &= \underbrace{\mathbf{t} \cdot \mathbf{G}_\kappa \cdot \mathbf{T}}_{=0} + \underbrace{\mathbf{t} \cdot \Pi^{-1} \cdot \mathbf{E}}_{=\mathbf{h}'} \end{aligned}$$

Indeed, if $\text{Ham}(\mathbf{t}) = c$, then each coordinate of $\mathbf{t} \cdot \Pi^{-1} \mathbf{E}$ is nonzero with probability at most $c\rho$, and hence the resulting vector has sublinear weight with overwhelming probability when c is constant and $\rho = O(1/\sqrt{n})$. By contrast, for a random matrix and any non-zero choice of $\mathbf{h}$, the output of this procedure has high weight with overwhelming probability. Therefore, the attacker easily distinguishes between the two cases.

We can also show that the strong McEliece with Noise assumption implies the McEliece with Noise assumption under the additional standard $1/\sqrt{n}$ -noise LPN assumption.

Lemma 2. *Assume that the strong McEliece with Noise assumption and the $1/\sqrt{n}$ -noise LPN assumption hold. Then the McEliece with Noise assumption holds.*

Proof (Sketch). The proof is a simple hybrid argument. First, we use the strong McEliece with Noise assumption to replace the matrix $\mathbf{A}$ with a uniformly random matrix. Then, in the following hybrids, we successively replace the values $\mathbf{A}_1 \mathbf{s}_i + \mathbf{e}_i$ by uniformly random $\mathbf{u}_i$, with $\mathbf{A} := (\mathbf{A}_1 \| \mathbf{A}_2)$.

For the latter step we use a standard consequence of low-noise LPN. Starting from LPN with a uniform secret and Bernoulli noise of rate $\rho = O(1/\sqrt{n})$, one first puts the sample matrix in systematic form and applies the corresponding parity-check transformation. This eliminates the original secret and yields an LPN instance in which both the new secret and the new error are distributed according to Ber_ρ . Since our assumption uses a constant error rate ϵ and $\rho = o(1)$, we can add an independent smoothing error to the right-hand side, chosen so that the sum of the low-noise error and the smoothing error is distributed as Ber_ϵ ; this operation is oblivious and preserves indistinguishability. Finally, the leakage lemma of [Döt15, Lemma 6] allows us to reveal the Hamming weights of the new secret and of the smoothed error, losing only a polynomial factor. Conditioning on these weights gives the fixed-weight, and hence Hamming-ball, distributions used here. The cited lemma is stated over $\mathbb{F}_2$, but the argument is linear-algebraic and extends directly to any fixed finite field. $\square$

Construction We will now show that the noisy McEliece assumption for a suitable family of codes $(\mathcal{C}_\kappa)_\kappa$, together with the standard $1/\sqrt{n}$ -noise LPN assumption, gives rise to adaptively robust pseudorandom codes. Importantly, this PRC construction enjoys essentially the same robustness properties as $\mathcal{C}_\kappa$.

Now fix a family of codes $(\mathcal{C}_\kappa)_\kappa$ which can efficiently decode from an α fraction of errors. Fix constants $\epsilon, \delta > 0$. Let $m = O(n)$ and ρ be such that $4\rho^2 \cdot n + \epsilon + \delta \leq \alpha$ (e.g. $\rho = \sqrt{(\alpha - \epsilon - \delta)/(4n)} = O(1/\sqrt{n})$). For soundness, we additionally require $2n + H_q(\alpha)m \leq m - \omega(\log_q \lambda)$, where $q = |\mathbb{F}|$ and H_q denotes the q -ary entropy function. The pseudorandom code $\text{PRC} = (\text{Setup}, \text{Encode}, \text{Decode})$ is given as follows.

- **Setup**(1^λ): Choose $\mathbf{\Pi} \leftarrow S_m$, $\kappa \leftarrow K$ and $\mathbf{T} \leftarrow (\mathbb{F}^{k \times 2n})_{\text{full}}$ uniformly at random, and sample $\mathbf{E} \leftarrow \text{Ber}_\rho^{m \times 2n}$. Set $\mathbf{A} = \mathbf{\Pi} \mathbf{G}_\kappa \mathbf{T} + \mathbf{E}$. Output $\text{pk} = \mathbf{A}$ and $\text{sk} = (\mathbf{\Pi}, \kappa, \mathbf{T})$.
- **Encode**($\text{pk} = \mathbf{A}$, $\text{msg} \in \mathcal{B}_\rho(\mathbb{F}^n)$): Sample $\mathbf{s} \leftarrow \mathcal{B}_\rho(\mathbb{F}^n)$ uniformly at random, set $\mathbf{x} = (\mathbf{s}, \text{msg}) \in \mathcal{B}_\rho(\mathbb{F}^{2n})$ and sample $\mathbf{e} \leftarrow \mathcal{B}_\epsilon(\mathbb{F}^m)$. Compute and output $\mathbf{c} = \mathbf{A} \cdot \mathbf{x} + \mathbf{e}$.
- **Decode**($\text{sk} = (\mathbf{\Pi}, \kappa, \mathbf{T})$, $\tilde{\mathbf{c}}$): Compute $\mathbf{c}' = \mathbf{\Pi}^{-1} \tilde{\mathbf{c}}$ and then $\mathbf{y} = \mathcal{C}_\kappa.\text{Decode}(\mathbf{c}')$. If this outputs $\perp$, output $\perp$. If $\text{Ham}(\mathbf{c}' - \mathbf{G}_\kappa \mathbf{y}) > \alpha m$, output $\perp$. Otherwise, compute the unique $\mathbf{x}$ such that $\mathbf{T} \cdot \mathbf{x} = \mathbf{y}$, parse $\mathbf{x} = (\mathbf{s}, \text{msg})$ and check if $\mathbf{s} \in \mathcal{B}_\rho(\mathbb{F}^n)$ and $\text{msg} \in \mathcal{B}_\rho(\mathbb{F}^n)$. If any of these steps fail, output $\perp$; otherwise, output msg .

We note that in the above construction, we require the message msg to be sparse. To encode a dense message $\text{msg}' \in \mathbb{F}^\ell$, given ρ , we choose n such that $|\mathbb{F}^\ell| \leq |\mathcal{B}_\rho(\mathbb{F}^n)|$.

Pseudorandomness The pseudorandomness property of PRC follows immediately from (in fact, is equivalent to) the noisy McEliece assumption.

Lemma 3. *The code PRC has pseudorandom codewords, given that the noisy McEliece assumption holds.*

Soundness We first record that the construction satisfies the usual soundness property: a uniformly random word is rejected except with negligible probability.

Lemma 4. *Assume that $2n + H_q(\alpha)m \leq m - \omega(\log_q \lambda)$, where $q = |\mathbb{F}|$ and H_q is the q -ary entropy function. Then PRC is sound.*

Proof. Fix any secret key $\text{sk} = (\mathbf{\Pi}, \kappa, \mathbf{T})$. Since $\mathbf{G}_\kappa$ has rank k and $\mathbf{T}$ has full column rank, the effective code

$$\mathcal{C}_{\kappa, \mathbf{T}} = \{\mathbf{\Pi G}_\kappa \mathbf{T} \mathbf{x} : \mathbf{x} \in \mathbb{F}^{2n}\}$$

has size q^{2n} . By the explicit distance check in `Decode`, if a word $\mathbf{u} \in \mathbb{F}^m$ is accepted, then $\mathbf{u}$ lies in an αm -Hamming ball around some codeword in $\mathcal{C}_{\kappa, \mathbf{T}}$. Hence

$$\Pr_{\mathbf{u} \leftarrow \mathbb{F}^m} [\text{Decode}(\text{sk}, \mathbf{u}) \neq \perp] \leq \frac{|\mathcal{C}_{\kappa, \mathbf{T}}| \cdot |\mathcal{B}_\alpha(\mathbb{F}^m)|}{q^m} \leq q^{2n + H_q(\alpha)m - m} \leq q^{-\omega(\log_q \lambda)} = \text{negl}(\lambda).$$

The bound is uniform over the choice of sk , and therefore also holds over the randomness of `Setup`. $\square$

Adaptive Robustness We will now argue adaptive robustness of PRC. We will use the following simple facts about Bernoulli distributed matrices in the proof. For a matrix $\mathbf{M} = (\mathbf{m}_1, \dots, \mathbf{m}_N) \in \mathbb{F}^{m \times N}$, we define the induced matrix weight $\text{Ham}(\mathbf{M})$ via $\text{Ham}(\mathbf{M}) = \max_{i \in [N]} \text{Ham}(\mathbf{m}_i)$. This matrix weight is sub-multiplicative, that is, it holds for every $\mathbf{x} \in \mathbb{F}^N$ that $\text{Ham}(\mathbf{M} \cdot \mathbf{x}) \leq \text{Ham}(\mathbf{M}) \cdot \text{Ham}(\mathbf{x})$. If $\mathbf{E} \leftarrow \text{Ber}_\rho^{m \times N}$, then a simple Chernoff bound and a union bound imply that $\text{Ham}(\mathbf{E}) \leq 2\rho m$ except with probability $N \cdot e^{-\frac{1}{3}\rho m}$, which is negligible for our parameter choices.

Lemma 5. *Assume that $\mathcal{C}_\kappa$ can efficiently uniquely decode from up to an $\alpha < 1/4$ fraction of Hamming errors and that $4\rho^2 n + \epsilon + \delta \leq \alpha$. Then PRC is a strong adaptively robust PRC against a δ -fraction of adversarial errors, except with negligible probability over the choice of the public key $\mathbf{A}$.*

Proof. As discussed above, it follows by a Chernoff bound and a union bound that $\text{Ham}(\mathbf{E}) \leq 2\rho m$, except with negligible probability $2n \cdot e^{-\frac{1}{3}\rho m}$. Hence, fix a public key $\text{pk} = \mathbf{A}$ where $\mathbf{A} = \mathbf{\Pi} \cdot \mathbf{G}_\kappa \cdot \mathbf{T} + \mathbf{E}$, such that $\text{Ham}(\mathbf{E}) \leq 2\rho m$. Consider a codeword

$$\begin{aligned} \mathbf{c} &= \mathbf{A} \mathbf{x} + \mathbf{e} \\ &= \mathbf{\Pi G}_\kappa \mathbf{T} \mathbf{x} + \mathbf{E} \mathbf{x} + \mathbf{e}. \end{aligned}$$

It holds that

$$\text{Ham}(\mathbf{E} \mathbf{x} + \mathbf{e}) \leq \text{Ham}(\mathbf{E}) \cdot \text{Ham}(\mathbf{x}) + \text{Ham}(\mathbf{e}) \leq 2\rho \cdot m \cdot 2\rho \cdot n + \epsilon \cdot m = (4\rho^2 n + \epsilon)m.$$

As we require $4\rho^2 n + \epsilon \leq \alpha - \delta$, it follows that PRC is adaptively robust against errors of weight at most $\delta \cdot m$: fix any adversarial error $\mathbf{e}^*$ with Hamming weight at most $\delta \cdot m$. Then

$$\hat{\mathbf{c}} = \mathbf{c} + \mathbf{e}^* = \mathbf{\Pi G}_\kappa \mathbf{T} \mathbf{x} + \mathbf{E} \mathbf{x} + \mathbf{e} + \mathbf{e}^* = \mathbf{\Pi G}_\kappa \mathbf{T} \mathbf{x} + \mathbf{d},$$

where we set $\mathbf{d} = \mathbf{E} \mathbf{x} + \mathbf{e} + \mathbf{e}^*$. As $\text{Ham}(\mathbf{E} \mathbf{x} + \mathbf{e}) \leq (\alpha - \delta)m$ and $\text{Ham}(\mathbf{e}^*) \leq \delta m$, it follows by the triangle inequality that $\text{Ham}(\mathbf{d}) \leq \alpha m$. Hence

$$\begin{aligned} \mathbf{c}' &= \mathbf{\Pi}^{-1}(\hat{\mathbf{c}}) \\ &= \mathbf{G}_\kappa \mathbf{T} \mathbf{x} + \mathbf{\Pi}^{-1} \mathbf{d}. \end{aligned}$$

Thus we get that $\mathbf{c}'$ has distance at most αm from $\mathcal{C}_\kappa$ as $\text{Ham}(\mathbf{\Pi}^{-1} \mathbf{d}) = \text{Ham}(\mathbf{d}) \leq \alpha m$. Hence $\mathcal{C}_\kappa.\text{Decode}$ will uniquely recover $\mathbf{T} \mathbf{x}$. Since $\mathbf{T}$ has full column rank, the final linear solve recovers $\mathbf{x}$ and therefore `msg`. $\square$

5 CCA secure Pseudorandom Codes

In this section, we present our CCA PRC transformation. In general, constructing PRCs becomes easier when the underlying field is large [DMR25]. As in prior work, we focus exclusively on the binary case for clarity of presentation. Nevertheless, we expect that our transformation extends naturally to arbitrary finite fields, provided that the underlying primitives are replaced by their corresponding counterparts over such fields, for example by relying on the noisy McEliece assumption (Section 4) defined over general finite fields.

Our construction requires the following primitives:

- $\text{DS} = (\text{DS.Setup}, \text{DS.Sign}, \text{DS.Verify})$: a one-time strong unforgeable signature scheme with message space $\{0, 1\}^*$, verification key space $\{0, 1\}^{\ell_{\text{vk}}}$ and signature space $\{0, 1\}^\lambda$.
- $\text{Com} = (\text{Com.Setup}, \text{Com.Commit}, \text{Com.Open}, \text{Com.Equiv})$: a statistically binding equivocal tag-based set-commitment scheme with pseudorandom commitments, with tag space $\{0, 1\}^{\ell_{\text{vk}}}$, commitment space $\{0, 1\}^{\ell_{\text{com}}}$ and opening space $\{0, 1\}^{\ell_{\text{op}}}$.
- $\text{PKE} = (\text{PKE.Setup}, \text{PKE.Enc}, \text{PKE.Dec})$: an almost perfect-correctness IND-CPA secure PKE with message space $\{0, 1\}^{\ell_{\text{op}} + \ell_{\text{vk}} + 1}$ and randomness $\{0, 1\}^\lambda$. The ciphertext size is $\ell_{\text{pke.ct}}$.
- $\text{ECC} = (\text{ECC.Setup}, \text{ECC.Encode}, \text{ECC.Decode})$: a γ -robust error-correcting code over alphabet $\{0, 1\}$, with message-rate c_{ecc} (that is, for binary strings of length k , it produces codewords of length $c_{\text{ecc}} \cdot k$).
- $\text{HPRG} = (\text{HPRG.Setup}, \text{HPRG.Eval})$: a hinting pseudorandom generator that takes $n = c_{\text{ecc}} \cdot \lambda$ bits as input, and outputs $n + 2$ blocks where each block of the output has λ bits.
- $\text{PRF} = (\text{PRF.Setup}, \text{PRF.Eval})$: a pseudorandom function with key space, input and output space $\{0, 1\}^\lambda$.
- PRG : a (regular) pseudorandom generator that maps λ bits to ℓ_{PRG} bits, where $\ell_{\text{PRG}} = 2c_{\text{ecc}} \cdot (\ell_{\text{msg}} + \lambda + 2n \cdot \ell_{\text{pke.ct}} + n \cdot (\ell_{\text{com}})) + 3\lambda$. We will interpret the output as a number in $\mathbb{F}_{2^{\ell_{\text{PRG}}}}$.
- $\text{PRC} = (\text{PRC.Setup}, \text{PRC.Encode}, \text{PRC.Decode})$: an α adaptively-robust CPA secure public-key pseudorandom code with message space $\{0, 1\}^{\ell_{\text{pke.ct}} + \ell_{\text{com}}}$, randomness $\{0, 1\}^\lambda$. The codeword length is $\ell_{\text{cw}} = \ell_{\text{PRG}} / (2 \cdot n)$.¹⁰

Setup(1^λ): The setup algorithm performs the following steps:

1. Sample $\text{ecc.pp} \leftarrow \text{ECC.Setup}(1^\gamma)$.
2. Sample $\text{hprg.pp} \leftarrow \text{HPRG.Setup}(1^\lambda)$.
3. Sample $2n$ different PKE keys, i.e., $(\text{pke.pk}_{i,b}, \text{pke.sk}_{i,b}) \leftarrow \text{PKE.Setup}(1^\lambda)$ for all $b \in \{0, 1\}, i \in [n]$.
4. Sample $2n$ different PRC keys, i.e., $(\text{prc.pk}_{i,b}, \text{prc.sk}_{i,b}) \leftarrow \text{PRC.Setup}(1^\lambda)$ for all $b \in \{0, 1\}, i \in [n]$.
5. Sample $\text{com.pp} \leftarrow \text{Com.Setup}(1^\lambda, 1^n)$.
6. Sample $B_{\text{pl}} \leftarrow \{0, 1\}^{\ell_{\text{PRG}}}$.
7. Output the public key $\text{pk} = (\text{ecc.pp}, \text{hprg.pp}, B_{\text{pl}}, \text{com.pp}, (\text{pke.pk}_{i,b}, \text{prc.pk}_{i,b})_{i \in [n], b \in \{0, 1\}})$ and the secret key $\text{sk} = (\text{pke.sk}_{i,0}, \text{prc.sk}_{i,0})_{i \in [n]}$.

Enc(pk, m): Parse the public key $\text{pk} = (\text{ecc.pp}, \text{hprg.pp}, B_{\text{pl}}, \text{com.pp}, (\text{pke.pk}_{i,b}, \text{prc.pk}_{i,b})_{i \in [n], b \in \{0, 1\}})$.

1. Sampling the various random strings needed for encryption:
 - (a) Sample $s \leftarrow \{0, 1\}^\lambda$, set $t = \text{ECC.Encode}(\text{ecc.pp}, s)$.
 - (b) Sample $(\text{ds.vk}, \text{ds.sk}) \leftarrow \text{DS.Setup}(1^\lambda)$.
 - (c) Sample $K \leftarrow \text{PRF.Setup}(1^\lambda)$.
 - (d) For each $i \in [n]$, choose $\tilde{r}_i \leftarrow \{0, 1\}^\lambda$.
 - (e) Let $(r_{00}, r_{01}, r_1, \dots, r_n) = \text{HPRG.Eval}(\text{hprg.pp}, t)$.
2. Compute the PKE ciphertexts and the commitments.
First compute $(\text{com}_i, \text{op}_i)_i \leftarrow \text{Commit}(\text{com.pp}, t, \text{ds.vk})$ (here, ds.vk is the tag used for the commitment).
For each $i \in [n]$,

$$\begin{aligned} \text{pke.ct}_{i,t_i} &:= \text{PKE.Enc}(\text{pke.pk}_{i,t_i}, 1 \parallel \text{op}_i \parallel \text{ds.vk}; r_i) \\ \text{pke.ct}_{i,1-t_i} &:= \text{PKE.Enc}(\text{pke.pk}_{i,1-t_i}, 0^{\ell_{\text{op}} + \ell_{\text{vk}} + 1}; \tilde{r}_i) \end{aligned}$$

3. Compute the PRC ciphertexts and the header.
 - (a) For each $i \in [n]$ and $b \in \{0, 1\}$, $c_{i,b} := \text{PRC.Enc}(\text{prc.pk}_{i,b}, (\text{pke.ct}_{i,b}, \text{com}_i); \text{PRF}(K, (i, b)))$.
 - (b) Set $c := \text{PRG}(r_{00}) + \text{ECC.Encode}(m, K, (\text{pke.ct}_{i,b}, \text{com}_i)_{i,b}) \cdot B_{\text{pl}}$.
4. Computing the digital signature:
 - (a) Let $\text{ds.x} = (c, (c_{i,0}, c_{i,1})_{i \in [n]})$. Compute $\text{ds.}\sigma \leftarrow \text{DS.Sign}(\text{ds.sk}, \text{ds.x})$. Pad the signature so that the length is ℓ_{PRG} .
 - (b) Compute $\sigma = \text{PRG}(r_{01}) \oplus \text{ECC.Encode}(\text{ecc.pp}, \text{ds.}\sigma)$.
5. Output $\text{ct} = (\text{ds.x}, \sigma)$.

¹⁰If needed, the message is padded so that the PRC codeword length is $\ell_{\text{PRG}}/2n$.

Note that the final ciphertext consists of three parts, each of length ℓ_{PRG} : c , $(c_{i,b})_{i,b}$ and σ .

Dec(sk, ct): Parse the secret key $\text{sk} := (\text{pke.sk}_{i,0}, \text{prc.sk}_{i,0})_{i \in [n]}$ and the ciphertext $\text{ct} = (\text{ds.x}, \sigma)$ where $\text{ds.x} = (c, \{c_{i,0}, c_{i,1}\}_{i \in [n]})$:

1. Compute $t = \text{PRC.Find}(\text{sk}, (c_{i,0})_i)$ (Figure 1).
2. Compute $(\text{msg}, \text{K}, c', (c'_{i,b}, \text{pke.ct}_{i,b}, \text{com}_i)_{i,b}, \text{ds.}\sigma) \leftarrow \text{PRC.CheckNoise}(c, (c_{i,b})_{i,b}, \sigma, t, (\text{prc.sk}_{i,0})_i)$ (Figure 2).
3. Let $\text{ds.x}' = (c', (c'_{i,b})_{i,b})$. Compute $\text{ds.vk} \leftarrow \text{PRC.CheckPKE}(t, (\text{pke.ct}_{i,t_i}, \text{com}_i)_i)$ (Figure 3).
4. If $\text{DS.Verify}(\text{ds.vk}, \text{ds.x}', \text{ds.}\sigma) = 1$, output msg , else output $\perp$.

PRC.Find	
Inputs:	Secret Key $\text{sk} = (\text{pke.sk}_{i,0}, \text{prc.sk}_{i,0})_{i \in [n]}$ PRC ciphertexts $(c_{i,0})_{i \in [n]}$
Output: t	1. For each $i \in [n]$, do the following: (a) Let $(\text{pke.ct}_{i,0}, \text{com}_i) = \text{PRC.Decode}(\text{prc.sk}_{i,0}, c_{i,0})$ (b) Let $\text{m}_i = \text{PKE.Dec}(\text{pke.sk}_{i,0}, \text{pke.ct}_{i,0})$. (c) If $\text{m}_i = 1 \parallel \text{op}_i \parallel \text{ds.vk}_i$ and $\text{Open}(\text{com.pp}, 0, \text{com}_i, \text{op}_i, \text{ds.vk}_i) = \top$, set $\tilde{t}_i = 0$. Else set $\tilde{t}_i = 1$. 2. Compute $s = \text{ECC.Decode}(\text{ecc.pp}, \tilde{t})$ and $t = \text{ECC.Encode}(\text{ecc.pp}, s)$. 3. Output t.

Fig. 1. Routine PRC.Find.

Information rate: Suppose that ECC encodes t -bit messages into codewords of length $k_{\text{ecc}} \cdot t$ for some constant $k_{\text{ecc}} > 1$, and let ℓ_{msg} denote the message length in our PRC scheme. Since both $|\text{op}_i|$ and $|\text{ds.vk}|$ are polynomial in the security parameter, it follows that $|\text{pke.ct}_{i,t_i}| = \text{poly}(\lambda)$. Likewise, because $|\text{com}_i| = \text{poly}(\lambda)$, we have $|c_{i,b}| = \text{poly}(\lambda)$. As $n = \text{poly}(\lambda)$, this further implies that $|\{c_{i,b}\}| = \text{poly}(\lambda)$. Since, $|\text{K}| = \text{poly}(\lambda)$, total payload size is therefore $|c| = k_{\text{ecc}} \cdot (\ell_{\text{msg}} + \text{poly}(\lambda))$, while the signature contributes $O(\lambda)$ additional bits. Consequently, the information rate of the scheme is

$$\frac{\ell_{\text{msg}}}{k_{\text{ecc}} \cdot (\ell_{\text{msg}} + \text{poly}(\lambda)) + \text{poly}(\lambda)} = \left(k_{\text{ecc}} + \frac{k_{\text{ecc}} \cdot \text{poly}(\lambda)}{\ell_{\text{msg}}} + \frac{\text{poly}(\lambda)}{\ell_{\text{msg}}} \right)^{-1}.$$

In particular, as ℓ_{msg} grows large, the information rate approaches k_{ecc}^{-1} .

Theorem 8 (Adaptive Robustness). *Assuming PRC is α -robust, ECC is γ -robust, then the above scheme is $\frac{\alpha \cdot \gamma}{6}$ - adaptively robust.*

Proof. Let $\ell_c := |c|$, $\ell_{\text{cw}} = |c_{i,b}|$, $\ell_\sigma = |\sigma|$, we have $|\ell_c| = |\ell_\sigma| = 2n \cdot \ell_{\text{cw}}$, and $\ell_{\text{ct}} := \ell_c + 2n \cdot \ell_{\text{cw}} + \ell_\sigma$. Let $\delta = \alpha\gamma/6$. We then need to argue, to prove adaptive corruptness, that an adversary cannot output a tuple (m, r, x) where $\text{Ham}(\text{Enc}(\text{pk}, m; r), x) \leq \delta \ell_{\text{ct}}$ while $m \neq \text{Dec}(\text{sk}, x)$ with non-negligible probability. To prove this, we argue something stronger, we consider an adversary that submits a message-randomness pair (m, r)




Fig. 2. Routine PRC.CheckNoise.




Fig. 3. Routine PRC.CheckPKE.

as well as three error vectors $\mathbf{e}_1 \in \{0, 1\}^{\ell_c}$, $\mathbf{e}_2 \in \{0, 1\}^{2n\ell_{\text{cw}}}$, $\mathbf{e}_3 \in \{0, 1\}^{\ell_\sigma}$ with $\text{Ham}(\mathbf{e}_i) \leq \delta \ell_{\text{ct}}$ for all $i \in [3]$.

Now, we say that such an adversary wins the game if it holds that

$$\text{Dec}\left(\text{sk}, \left(c, (c_{i,b})_{i \in [n], b \in \{0,1\}}, \sigma\right)\right) \neq \text{Dec}\left(\text{sk}, \left(c \oplus \mathbf{e}_1, (c_{i,b})_{i \in [n], b \in \{0,1\}} \oplus \mathbf{e}_2, \sigma \oplus \mathbf{e}_3\right)\right)$$

with $\text{ct} = \left(c, (c_{i,b})_{i,b}, \sigma\right) = \text{Enc}(\text{pk}, m; r)$. We can directly see that such a game achieves stronger security than the adaptive robustness game since the adaptive robustness adversary is only allowed to submit a single vector $\mathbf{e} \in \{0,1\}^{\ell_c + 2n\ell_{\text{cw}} + \ell_\sigma}$ with $\text{Ham}(\mathbf{e}) \leq \delta\ell_{\text{ct}}$. Let $c' := c \oplus \mathbf{e}_1$, $c'_{i,b} := c_{i,b} \oplus \mathbf{e}_2$, $\sigma' := \sigma \oplus \mathbf{e}_3$, and let $\text{ct}' = \left(c', (c'_{i,b})_{i,b}, \sigma'\right)$.

To prove the above equality, it suffices to show that the procedures **Find** and **CheckNoise** output the same for the ciphertexts ct and ct' . Now, we are walking through the different steps of the two procedures to bound these probabilities.

Claim. $\text{Find}(\text{sk}, \text{ct}) = \text{Find}(\text{sk}, \text{ct}')$.

Proof. First, note that if t is the HPRG seed used to encrypt m , then Step 1 of **Find**(sk, ct) computes $\tilde{t} = t$, and as a result, **Find**(sk, ct) outputs t . This follows from the perfect correctness of the PKE and commitment schemes. We will now argue that if $\text{Ham}(\mathbf{e}_2) \leq \delta \cdot \ell_{\text{ct}}$, then Step 1 of **Find**(sk, ct') produces a $\tilde{t}'$ such that $\Delta(\tilde{t}', t) \leq \gamma \cdot n$, assuming α -adaptive robustness of the base PRC scheme.

Parse $\mathbf{e}_2 = (\mathbf{e}_{2,i,b})_{i \in [n], b \in \{0,1\}}$, and let $\mathcal{I} = \{i : \text{Ham}(\mathbf{e}_{2,i,b}) \leq \alpha \cdot \ell_{\text{cw}} \text{ for both } b \in \{0,1\}\}$. Since $\delta = \alpha \cdot \gamma/6$, it follows that $|\mathcal{I}| \geq (1 - \gamma) \cdot n$.¹¹ Using the adaptive robustness of the base PRC scheme (and the perfect correctness of the PKE scheme and the commitment scheme), we get that $\tilde{t}'_i = t_i$ for all $i \in \mathcal{I}$. In particular, if $c_{i,b}$ is the PRC encoding of $(\text{pke.ct}_{i,b}, \text{com}_i)$ with randomness $z_{i,b}$ and $\text{Dec}(\text{prc.sk}_{i,b}, (c_{i,b} \oplus \mathbf{e}_{2,i,b})) \neq (\text{pke.ct}_{i,b}, \text{com}_i)$, then the reduction algorithm can break α -adaptive robustness of the base PRC scheme by sending $(\text{pke.ct}_{i,b}, \text{com}_i)$ as the message, $z_{i,b}$ as the randomness and $(c_{i,b} \oplus \mathbf{e}_{2,i,b})$ as the codeword. $\square$

It remains to analyze the behavior of **CheckNoise** on the inputs ct and ct' where both procedures take as an input the same t (which follows from the above analysis of **Find**). Given that both procedures take as an input the same t , it follows that step 1 (a) will also output the same values. Next, we rely on the robustness of ECC. In more detail, since $\delta\ell_{\text{ct}} \leq \gamma\ell_c = \gamma\ell_\sigma$, for the steps 1. (b) and (c), it must hold that

$$\begin{aligned} \text{ECC.Decode}\left(\text{ecc.pp}, \left(c - \text{PRG}(r_{00}) \cdot B_{\text{pl}}^{-1}\right)\right) &= \text{ECC.Decode}\left(\text{ecc.pp}, \left((c \oplus \mathbf{e}_1) - \text{PRG}(r_{00}) \cdot B_{\text{pl}}^{-1}\right)\right) \\ \text{ECC.Decode}(\text{ecc.pp}, (\text{PRG}(r_{01}) \oplus \sigma)) &= \text{ECC.Decode}(\text{ecc.pp}, (\text{PRG}(r_{01}) \oplus \sigma \oplus \mathbf{e}_3)), \end{aligned}$$

If the above holds, then the steps 1 (d)-(f) directly have the same output as well. The bounds in step 2. are all satisfied due to our restriction $\text{Ham}(\mathbf{e}_i) \leq \delta\ell_{\text{ct}}$. To argue that step 3 yields the same output, we can, again, rely on the adaptive robustness as it has been done in the beginning of the proof. In more detail, the decoding output computed in step 3 (a) is the same as the decoding that is computed as the 1. step of **Find**. Only violating adaptive robustness would lead to a failing check in the step 3 (b), given the above analysis involving the error-correcting codes. $\square$

Theorem 9. *Assuming that PKE is an IND-CPA secure PKE scheme with almost perfect-correctness, PRC is an α adaptively robust IND-CPA secure PRC scheme, ECC is a γ -robust error-correcting code, Com is a commitment scheme satisfies statistical binding, equivocality, pseudorandomness of commitments and independence of tag for 0-commitments, HPRG is a secure hinting PRG, PRG is a secure PRG, PRF is a secure PRF and DS is a secure signature scheme. Then the above construction is a IND-CCA secure PRC scheme.*

Proof. Our proof proceeds via a sequence of hybrids. First, we will describe all hybrids, and then show that the hybrids are computationally indistinguishable.

Hybrid $\mathcal{H}_0$: The corresponds to the original security game with a few modifications as follows.

¹¹If $|\mathcal{I}| < (1 - \gamma) \cdot n$, then there exist at least $\gamma \cdot n$ indices i where either $\text{Ham}(\mathbf{e}_{2,i,0}) > \alpha \cdot \ell_{\text{cw}}$ or $\text{Ham}(\mathbf{e}_{2,i,1}) > \alpha \cdot \ell_{\text{cw}}$. Therefore, $\text{Ham}(\mathbf{e}) > \gamma \cdot n \cdot (\alpha \cdot \ell_{\text{cw}})$. But this is a contradiction, since we assumed $\text{Ham}(\mathbf{e}) \leq \delta \cdot \ell_{\text{ct}}$, $\delta = \alpha \cdot \gamma/6$ and $\ell_{\text{ct}} = 6n \cdot \ell_{\text{cw}}$

Description of Hybrid $\mathcal{H}_0$

– **Initialization Phase:** The challenger performs the following steps:

1. Sample $\text{ecc.pp} \leftarrow \text{ECC.Setup}(1^\lambda)$.
2. Sample $\text{hprg.pp} \leftarrow \text{HPRG.Setup}(1^\lambda)$.
3. Sample $2n$ different PKE keys, i.e., $(\text{pke.pk}_{i,b}, \text{pke.sk}_{i,b}) \leftarrow \text{PKE.Setup}(1^\lambda)$ for all $b \in \{0, 1\}, i \in [n]$.
4. Sample $2n$ different PRC keys, i.e., $(\text{prc.pk}_{i,b}, \text{prc.sk}_{i,b}) \leftarrow \text{PRC.Setup}(1^\lambda)$ for all $b \in \{0, 1\}, i \in [n]$.
5. Sample $\text{com.pp} \leftarrow \text{Com.Setup}(1^\lambda, 1^n)$.
6. Sample $B_{\text{pl}} \leftarrow \{0, 1\}^{\ell_{\text{PRG}}}$.
7. It sets $\text{pk} = (\text{ecc.pp}, \text{hprg.pp}, B_{\text{pl}}, \text{com.pp}, (\text{pke.pk}_{i,b}, \text{prc.pk}_{i,b})_{i \in [n], b \in \{0, 1\}})$ and $\text{sk} = (\text{pke.sk}_{i,0}, \text{prc.sk}_{i,0})_{i \in [n]}$.
8. It sends pk to the adversary.
9. Sample $s^* \leftarrow \{0, 1\}^\lambda$.
10. Set $t^* = \text{ECC.Encode}(\text{ecc.pp}, s^*)$.
11. Sample $K^* \leftarrow \text{PRF.Setup}(1^\lambda)$.
12. Sample $(\text{ds.vk}^*, \text{ds.sk}^*) \leftarrow \text{DS.Setup}(1^\lambda)$.
13. Compute $(\text{com}_i^*, \text{op}_i^*)_i \leftarrow \text{Commit}(\text{com.pp}, t^*, \text{ds.vk}^*)$.
14. Let $(r_{00}^*, r_{01}^*, r_1^*, \dots, r_n^*) = \text{HPRG.Eval}(\text{hprg.pp}, t^*)$. For each $i \in [n]$, choose $\tilde{r}_i^* \leftarrow \{0, 1\}^\ell$.
15. For each $i \in [n]$,
 - Compute $\text{pke.ct}_{i,t_i^*}^* := \text{PKE.Enc}(\text{pke.pk}_{i,t_i^*}, 1 \parallel \text{op}_i^* \parallel \text{ds.vk}^*; r_i^*)$
 - Compute $\text{pke.ct}_{i,1-t_i^*}^* := \text{PKE.Enc}(\text{pke.pk}_{i,1-t_i^*}, 0^{\ell_{\text{op}} + \ell_{\text{vk}} + 1}, \tilde{r}_i^*)$
 - Compute $c_{i,b}^* := \text{PRC.Enc}(\text{prc.pk}_{i,b}, (\text{pke.ct}_{i,b}^*, \text{com}_i^*); \text{PRF}(K^*, (i, b)))$.

– **Pre-Challenge Phase:** The adversary can adaptively query the decryption oracle polynomially many times. For each query $\text{ct} = (\text{ds.x}, \sigma)$ where $\text{ds.x} = (c, (c_{i,0}, c_{i,1})_i)$, the challenger computes

- $t = \text{PRC.Find}(\text{sk}, (c_{i,0})_i)$,
- $(\text{msg}, K, (\text{pke.ct}_{i,b}, \text{com}_i)_i, \text{ds.}\sigma) \leftarrow \text{CheckNoise}(c, (c_{i,b})_{i,b}, \sigma, t)$,
- $\text{ds.vk} \leftarrow \text{CheckPKE}(t, (\text{pke.ct}_{i,t_i}, \text{com}_i)_i)$,
- outputs msg if $\text{DS.Verify}(\text{ds.vk}, \text{ds.x}, \text{ds.}\sigma) = 1$.

– **Challenge phase:** The adversary sends m^* to the challenger who performs the following steps (recall, some of the components used in this phase were already computed in the Setup Phase).

1. It sets $c^* := \text{PRG}(r_{00}^*) + \text{ECC.Encode}(\text{ecc.pp}, (m^*, K^*, (\text{pke.ct}_{i,b}^*, \text{com}_i^*)_{i \in [n], b \in \{0, 1\}})) \cdot B_{\text{pl}}$.
2. It sets $\text{ds.x}^* = (c^*, (c_{i,0}^*, c_{i,1}^*)_{i \in [n]})$.
3. It computes $\text{ds.}\sigma^* \leftarrow \text{DS.Sign}(\text{ds.sk}^*, \text{ds.x}^*)$.
4. It compute $\sigma^* = \text{PRG}(r_{01}^*) \oplus \text{ECC.Encode}(\text{ecc.pp}, \text{ds.}\sigma^*)$.
5. It returns $\text{ct}^* = (\text{ds.x}^*, \sigma^*)$.

– **Post-Challenge Phase:** The adversary can adaptively query the decryption oracle polynomially many times. These queries are handled similar to the pre-challenge decryption queries.

– **Response Phase:** The adversary outputs a bit β .

Hybrid $\mathcal{H}_1$: This hybrid proceeds as the previous one, with the only difference that the decryption oracle rejects queries where the underlying verification key ds.vk is equal to the verification key ds.vk^* used in the challenge encryption query. The indistinguishability of this and the previous hybrid can be argued using the unforgeability of the signature scheme DS by showing that the probability of an adversary submitting such a query is negligible. The details on this are given in Lemma 6.

Hybrid $\mathcal{H}_2$: In this hybrid, during the initialization phase, the commitment scheme is in equivocal mode, i.e., $(\text{com.pp}, (\text{com}_i^*, \text{op}_{i,0}^*, \text{op}_{i,1}^*)_i) \leftarrow \text{Equiv}(1^\lambda, 1^n, \text{ds.vk}^*)$ and $\text{op}_i^* := \text{op}_{i,t_i^*}^*$. The indistinguishability of this and the previous hybrid can be argued using the equivocality of Com. The details on this are given in Lemma 7.

Hybrid $\mathcal{H}_3$: In this hybrid, the decryption queries are handled using the procedures $\text{PRC.Find}^\dagger$ (see Fig. 4) and $\text{PRC.CheckNoise}^\dagger$ (see Fig. 5) on path t^* , instead of PRC.Find and PRC.CheckNoise . To argue the indistinguishability between this and the previous hybrid, we rely on the perfect correctness of PKE and ECC, and the adaptive robustness of PRC. This transition is quite delicate, and we prove it in detail in Lemma 8.




Fig. 4. Routine $\text{PRC.Find}^\dagger$.

Hybrid $\mathcal{H}_4$: In this hybrid, the ciphertexts $\text{pke.ct}_{i,\overline{t_i^*}}^*$ are generated as $\text{PKE.Enc}(\text{pke}_{i,\overline{t_i^*}}, 1 \parallel \text{op}_{i,\overline{t_i^*}} \parallel \text{ds.vk}^*; r_i^*)$ for all $i \in [n]$. The indistinguishability of this and the previous hybrid can be argued using the IND-CPA security of PKE. The details on this are given in Lemma 13.

Hybrid $\mathcal{H}_5$: In this hybrid, the procedure PRC.Find instead of $\text{PRC.Find}^\dagger$ is executed in the pre- and post-challenge phases. The transition from the previous hybrid to this hybrid proceeds analogously to the transition from $\mathcal{H}_2$ to $\mathcal{H}_3$, therefore, we refer to Lemma 8 for more details.

Hybrid $\mathcal{H}_6$: In this hybrid, the values $\tilde{r}_i^*, c^*$ and σ^* are chosen uniformly at random for all $i \in [n]$. The indistinguishability of this and the previous hybrid can be argued using the security of the hinting PRG HPRG and the standard PRG PRG. The details on this are given in Lemma 14.

Hybrid $\mathcal{H}_7$: In this hybrid, the ciphertexts $c_{i,b}^*$ for all $i \in [n], b \in \{0, 1\}$ are generated as $c_{i,b}^* := \text{PRC.Enc}(\text{prc.pk}_{i,b}, (\text{pke.ct}_{i,b}^*, \text{com}_i^*); u_{i,b})$ where $u_{i,b}$ are sampled uniformly at random. The indistinguishability of this and the previous hybrid can be argued using the pseudorandomness of PRF. The details on this are given in Lemma 15.

Hybrid $\mathcal{H}_8$: In this hybrid, the challenger chooses $c_{i,1}^*$ uniformly at random for all $i \in [n]$. The indistinguishability of this and the previous hybrid can be argued using the pseudorandomness of PRC. The details on this are given in Lemma 16.

Hybrid $\mathcal{H}_9$: In this hybrid, the challenger uses $\text{PRC.Find}^\dagger$ on path $\bar{1}$ instead of PRC.Find in the pre-challenge and post-challenge phases. The transition from the previous hybrid to this hybrid proceeds analogously to the transition from $\mathcal{H}_2$ to $\mathcal{H}_3$, therefore, we refer to Lemma 8 for more details.

Hybrid $\mathcal{H}_{10}$: In this hybrid, the challenger chooses $c_{i,0}^*$ uniformly at random for all $i \in [n]$. The transition from the previous hybrid to this hybrid proceeds analogously to the transition from $\mathcal{H}_7$ to $\mathcal{H}_8$, relying on the pseudorandomness of PRC. We refer to Lemma 16 for more details.




Fig. 5. Routine PRC.CheckNoise[†].

Hybrid $\mathcal{H}_{11}$: In this hybrid, the challenger uses PRC.Find instead of PRC.Find[†] in the pre-challenge and post-challenge phases. The transition from the previous hybrid to this hybrid proceeds analogously to the transition from $\mathcal{H}_2$ to $\mathcal{H}_3$, therefore, we refer to Lemma 8 for more details.

Hybrid $\mathcal{H}_{12}$: In this hybrid, during the initialization phase, the setup of the commitment scheme is changed back to the standard setup, i.e., $\text{com.pp} \leftarrow \text{Com.Setup}(1^\lambda, 1^n)$. The transition from the previous hybrid to this hybrid proceeds analogously to the transition from $\mathcal{H}_1$ to $\mathcal{H}_2$, relying on the equivocality of Com. We refer to Lemma 7 for more details.

Hybrid $\mathcal{H}_{13}$: In this hybrid, during the decryption queries, the challenger does not reject queries when $\text{ds.vk} = \text{ds.vk}^*$. The transition from the previous hybrid to this hybrid proceeds analogously to the transition from $\mathcal{H}_0$ to $\mathcal{H}_1$, relying on the unforgeability of DS. We refer to Lemma 6 for more details. Observe that hybrid $\mathcal{H}_{13}$ corresponds to the original security game, in which the challenger outputs a truly random string during the challenge phase.

Combining the below lemmas, using the triangle inequality, we conclude that for every PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $|\Pr[\mathcal{H}_0(\lambda, \mathcal{A}) = 1] - \Pr[\mathcal{H}_{12}(\lambda, \mathcal{A}) = 1]| \leq \text{negl}(\lambda)$, where $\mathcal{H}_i(\lambda, \mathcal{A})$ denotes the output of $\mathcal{A}$ in hybrid $\mathcal{H}_i$. $\square$

Lemma 6. Assuming DS is a strongly unforgeable one-time signature scheme, then for all PPT adversaries $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all λ , $|\Pr[\mathcal{H}_0(\lambda, \mathcal{A}) = 1] - \Pr[\mathcal{H}_1(\lambda, \mathcal{A}) = 1]| \leq \text{negl}(\lambda)$.

Proof. Suppose there exists an adversary $\mathcal{A}$ that is able to distinguish between the hybrids $\mathcal{H}_0$ and $\mathcal{H}_1$. Recall, in the CCA game, the adversary is not allowed decryption queries that are close to the challenge

ciphertext. Therefore, the only way $\mathcal{A}$ can distinguish between the two hybrids is if it makes a decryption query $\text{ct} = \left(c, (c_{i,b})_{i,b}, \sigma\right)$, that is not close to the challenge ciphertext $\text{ct}^* = \left(c^*, (c_{i,b}^*)_{i,b}, \sigma^*\right)$, where the underlying verification key is equal to ds.vk^* , and the decryption results in a non- $\perp$ output (in particular the signature verification in the last step of decryption passes). We turn $\mathcal{A}$ into an adversary $\mathcal{B}$ that violates the unforgeability of DS. In the first step, the reduction $\mathcal{B}$ obtains a verification key vk^* from the DS challenger. Afterwards, $\mathcal{B}$ proceeds as described in the description of $\mathcal{H}_0$ with the difference that it queries the challenger to generate the signature $\text{ds}.\sigma^*$ for the challenge ciphertext. Next, we specify how the adversary $\mathcal{B}$ will obtain its forgery. Let $\text{ct}' = \left(c', (c'_{i,b})_{i,b}, \sigma'\right)$ be the noise-free values recovered at the end of Step 1 in **CheckNoise**. Let $\text{ds}.x' = \left(c', (c'_{i,b})_{i,b}\right)$ (and recall $\text{ds}.x^* = \left(c^*, (c_{i,b}^*)_{i,b}\right)$ is the string that is signed in the challenge phase). Since $\Delta(\text{ct}, \text{ct}^*) > \delta\ell_{\text{ct}}$ and $\Delta(\text{ct}, \text{ct}') \leq \delta\ell_{\text{ct}}$,¹² we get that $\text{ct}' \neq \text{ct}^*$. We have two cases:

- $\text{ds}.x' \neq \text{ds}.x^*$. In this case, the reduction outputs $\text{ds}.x'$ and $\text{ds}.\sigma$ (output by **CheckNoise**). This signature verifies (since $\text{Dec}(\text{sk}, \text{ct}) \neq \perp$), and this is a valid forgery since $\text{ds}.x' \neq \text{ds}.x^*$.
- $\text{ds}.x' = \text{ds}.x^*$ but $\sigma' \neq \sigma^*$. We get a strong forgery if we can show that $\text{ds}.\sigma \neq \text{ds}.\sigma^*$. We will show in the claim below that if $\text{ds}.x' = \text{ds}.x^*$ and decryption of ct does not output $\perp$, then the candidate seed t extracted by **Find** is equal to t^* . From here, we can conclude that $\text{ds}.\sigma \neq \text{ds}.\sigma^*$.¹³

Claim. Let $\text{ct}^* = (\text{ds}.x^*, \sigma^*)$ be the challenge ciphertext (with HPRG seed t^*), and $\text{ct} = (\text{ds}.x, \sigma)$ be any ciphertext such that $\text{Dec}(\text{sk}, \text{ct}) \neq \perp$. Let $t = \text{Find}(\text{sk}, \text{ct})$, and $\text{ds}.x', \sigma'$ are the error-free terms computed in Step 1 of **CheckNoise**. If $\text{ds}.x^* = \text{ds}.x'$, then $t = t^*$ (assuming α -adaptive robustness for the base PRC scheme and perfect correctness for the KE scheme and the commitment scheme).

Proof. $\text{Dec}(\text{sk}, \text{ct})$ does not output $\perp$, therefore $\Delta\left((c_{i,b})_{i,b}, (c'_{i,b})_{i,b}\right) \leq \delta\ell_{\text{ct}}$. Consider the set

$$\mathcal{I} = \{i : \Delta(c_{i,b}, c'_{i,b}) \leq \alpha \cdot \ell_{\text{cw}} \text{ for both } b \in \{0, 1\}\}.$$

Since $\delta = \alpha \cdot \gamma / 6$, $|\mathcal{I}| \geq (1 - \gamma) \cdot n$.¹⁴ For each of these indices, $\Delta(c_{i,b}, c'_{i,b}) = \Delta(c_{i,b}, c_{i,b}^*) \leq \alpha \cdot \ell_{\text{cw}}$. Therefore, by adaptive robustness of the PRC,¹⁵ $\tilde{t}_i = t_i^*$ for all $i \in \mathcal{I}$. Since $\Delta(\tilde{t}, t^*) \leq \gamma \cdot n$, $t = t^*$. $\square$

Lemma 7. *Assuming Com satisfies equivocality, then for all PPT adversaries $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all λ , $|\Pr[\mathcal{H}_1(\lambda, \mathcal{A}) = 1] - \Pr[\mathcal{H}_2(\lambda, \mathcal{A}) = 1]| \leq \text{negl}(\lambda)$.*

Proof. The proof follows directly from the equivocality of the commitment scheme. The only distinction between the two hybrids lies in the way com.pp , commitments $(\text{com}_i^*, \text{op}_i^*)_i$ are generated. In $\mathcal{H}_1$, com.pp is generated using **Setup** and the commitments and openings are generated using the commitment algorithm, while in $\mathcal{H}_2$, the challenger samples $(\text{com.pp}, (\text{com}_i^*, \text{op}_{i,0}^*, \text{op}_{i,1}^*)_i) \leftarrow \text{Equiv}(1^\lambda, 1^n, \text{ds.vk}^*)$ and sets $\text{op}_i^* = \text{op}_{i,t_i^*}^*$. Therefore a PPT adversary that can distinguish between $\mathcal{H}_1$ and $\mathcal{H}_2$ can be used to break the equivocality of Com. $\square$

¹²This is because $\text{Dec}(\text{sk}, \text{ct})$ does not output $\perp$ in Step 2 of **CheckNoise**.

¹³If $t = t^*$, then $\sigma^* = \text{PRG}(r_{01}) \oplus \text{ECC.Encode}(\text{ecc.pp}, \text{ds}.\sigma^*)$ and $\sigma' = \text{PRG}(r_{01}) \oplus \text{ECC.Encode}(\text{ecc.pp}, \text{ds}.\sigma)$, where $\text{HPRG}(t) = (r_{00}, r_{01}, r_1, \dots, r_n)$.

¹⁴Suppose, on the contrary, there are more than $\gamma \cdot n$ indices where at least one out of the two codewords has greater than $\alpha \cdot \ell_{\text{cw}}$ noise. Then the total noise is more than $\gamma \cdot n \cdot (\alpha \cdot \ell_{\text{cw}})$. But this is a contradiction, as we assumed that the total noise is at most $\delta \cdot \ell_{\text{ct}} = (\alpha \cdot \gamma / 6)(3 \cdot (2 \cdot n \cdot \ell_{\text{cw}})) = \alpha \cdot \gamma \cdot \ell_{\text{cw}}$. Here, we are using the fact that $|c| = 2 \cdot n \cdot |c_{i,b}| = |\sigma|$.

¹⁵

Lemma 8. Assuming perfect correctness of PKE and ECC, and α -adaptive robustness of PRC, then for all PPT adversaries $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all λ , $|\Pr[\mathcal{H}_2(\lambda, \mathcal{A}) = 1] - \Pr[\mathcal{H}_3(\lambda, \mathcal{A}) = 1]| \leq \text{negl}(\lambda)$.

Proof. Before proving the transition from $\mathcal{H}_2$ to $\mathcal{H}_3$, we prove the following helping lemma.

Lemma 9. Assuming the commitment scheme satisfies statistical binding, and the PKE scheme satisfies almost perfect-correctness, if $\text{Dec}(\text{sk}, \text{ct}) \neq \perp$, then for all $i \in [n]$ with $\Delta(c_{i,b}, c'_{i,b}) \leq \alpha \ell_{\text{cw}}$ for both $b \in \{0, 1\}$, it holds that $\tilde{t}_i = t_i$.¹⁶ Therefore, if $\text{Dec}(\text{sk}, \text{ct}) \neq \perp$, then $\Delta(t, \tilde{t}) \leq \gamma \cdot n$.

High level proof sketch. Suppose $\text{Dec}(\text{sk}, \text{ct}) \neq \perp$, $\Delta(c_{i,b}, c'_{i,b}) \leq \alpha \ell_{\text{cw}}$ for both $b \in \{0, 1\}$, but $\tilde{t}_i = 0$ and $t_i = 1$. Since $\tilde{t}_i$ is set to 0, this means that com_i extracted from $c_{i,0}$ is a commitment of 0, and the corresponding opening is present in $\text{pke.ct}_{i,0}$ (which is also extracted from $c_{i,0}$). However, since $t_i = 1$ and $\text{Dec}(\text{sk}, \text{ct}) \neq \perp$, this means that com_i (which is extracted from the header c) is a commitment of 1 (this is checked in **CheckPKE**). Finally, note that since $\Delta(c_{i,b}, c'_{i,b}) \leq \alpha \ell_{\text{cw}}$, we also check that the com_i extracted from $c_{i,b}$ is equal to the com_i extracted from the header c . Therefore, if $t_i \neq \tilde{t}_i$, then one must break the statistical binding property of the commitment scheme. The other case ($t_i = 0$, $\tilde{t}_i = 1$) is handled similarly.

Proof. We prove this lemma by contradiction. Suppose there exists an $i \in [n]$ with $\Delta(c_{i,b}, c'_{i,b}) \leq \alpha \ell_{\text{cw}}$ for both $b \in \{0, 1\}$ such that $\tilde{t}_i \neq t_i$. We distinguish between two cases: first, $\tilde{t}_i = 0$ and $t_i = 1$ and, second, $\tilde{t}_i = 1$ and $t_i = 0$.

In the first case, since $\tilde{t}_i = 0$, **PRC.Find** has checked that $\text{Open}(\text{com.pp}, 0, \text{com}_i, \text{op}_i, \text{ds.vk}_i) = \top$ where $(\text{pke.ct}_{i,0}, \text{com}_i) = \text{PRC.Decode}(\text{prc.sk}_{i,0}, c_{i,0})$, $\underline{m}_i = \text{PKE.Dec}(\text{pke.sk}_{i,0}, \text{pke.ct}_{i,0})$, and $\underline{m}_i = 1 \parallel \text{op}_i \parallel \text{ds.vk}_i$. Since we assume that $\Delta(c_{i,b}, c'_{i,b}) \leq \alpha \ell_{\text{cw}}$ for both $b \in \{0, 1\}$, during the **PRC.CheckNoise** procedure it is checked that $\text{com}_i = \underline{\text{com}}_i$ where $\text{pke.ct}_{i,b}, \text{com}_i$ are obtained from c . Next, we proceed with **PRC.CheckPKE**, which, since $t_i = 1$, checks that $\text{Open}(\text{com.pp}, 1, \text{com}_i, \text{op}_i, \text{ds.vk}_i) = \top$ where com_i is obtained as part of the output of **PRC.CheckNoise**. Now, from these checks, we get a contradiction (due to the binding property of the commitment scheme). This concludes the proof of the first case.

Next, we analyze the second case. Since $\tilde{t}_i = 1$, the check in **PRC.Find** resulted in $\underline{m}_i \neq 1 \parallel \text{op}_i \parallel \text{ds.vk}_i$ or $\text{Open}(\text{com.pp}, 0, \text{com}_i, \text{op}_i, \text{ds.vk}_i) = \perp$, where $(\text{pke.ct}_{i,0}, \text{com}_i)$ is obtained from $\text{PRC.Decode}(\text{prc.sk}_{i,0}, c_{i,0})$ and $\underline{m}_i = \text{PKE.Dec}(\text{pke.sk}_{i,0}, \text{pke.ct}_{i,0})$. Since we assume that $\Delta(c_{i,b}, c'_{i,b}) \leq \alpha \ell_{\text{cw}}$ for both $b \in \{0, 1\}$, during the **PRC.CheckNoise** procedure it is checked that $\text{pke.ct}_{i,0} = \underline{\text{pke.ct}}_{i,0}$ and $\text{com}_i = \underline{\text{com}}_i$ where $\text{pke.ct}_{i,b}, \text{com}_i$ is obtained from c . Next, we proceed with **PRC.CheckPKE**, which, since $t_i = 0$, checks that $m_i = 1 \parallel \text{op}_i \parallel \text{ds.vk}_i$, $\text{pke.ct}_{i,0} = \text{PKE.Enc}(\text{pke.pk}_{i,0}, m_i; r_i)$ and $\text{Open}(\text{com.pp}, 0, \text{com}_i, \text{op}_i, \text{ds.vk}_i) = \top$. Assuming the perfect correctness of PKE, we reach a contradiction since the check in **Find** resulted in $\underline{m}_i \neq 1 \parallel \text{op}_i \parallel \text{ds.vk}_i$ or $\text{Open}(\text{com.pp}, 0, \text{com}_i, \text{op}_i, \text{ds.vk}_i) = \perp$ whereas in **CheckPKE**, we check that $\text{pke.ct}_{i,0}$ is an encryption of $m_i = 1 \parallel \text{op}_i \parallel \text{ds.vk}_i$ and $\text{Open}(\text{com.pp}, 0, \text{com}_i, \text{op}_i, \text{ds.vk}_i) = \top$. This concludes the proof of the second case.

Having considered the above two cases, it remains to argue the last part of the lemma. From $\text{Dec}(\text{sk}, \text{ct}) \neq \perp$, we obtain that **PRC.CheckNoise** succeeded without any errors which, in turn, means that $\Delta((c_{i,b}), (c'_{i,b})) \leq \delta \ell_{\text{ct}}$ (follows from Step 2 of **CheckNoise**). This implies that $\Delta(c_{i,b}, c'_{i,b}) \geq \alpha \ell_{\text{cw}}$ for at most γn indices. Therefore, for at least $(n - \gamma n)$ indices, we have that $\Delta(c_{i,b}, c'_{i,b}) \leq \alpha \ell_{\text{cw}}$ for both $b \in \{0, 1\}$. Applying the first half of Lemma 9, we get that $\tilde{t}_i = t_i$ for at least $(n - \gamma n)$ indices. This implies that $\Delta(\tilde{t}, t) \leq n - (n - \gamma n) = \gamma n$. $\square$

The same analysis follow analogously for the **Find**[†]. Therefore, we obtain the following corollary.

Corollary 2. If $\text{Dec}^\dagger(\text{sk}, \text{ct}) \neq \perp$, then it must hold that $\Delta(t^\dagger, \tilde{t}^\dagger) \leq \gamma n$ (here, $t^\dagger$ is the output of **Find**[†], and $\tilde{t}^\dagger$ is the output of Step 1 of **Find**[†]).

Proof. The proof follows as above with the following helping lemma.

Lemma 10. If $\text{Dec}^\dagger(\text{sk}, \text{ct}) \neq \perp$, then it follows that for all $i \in [n]$ with $\Delta(c_{i,b}, c'_{i,b}) \leq \alpha \ell_{\text{cw}}$ it holds that $\tilde{t}^\dagger_i = t^\dagger_i$.

¹⁶Recall $\tilde{t}$ is produced in Step 1 of **Find** and t is the output of **Find**.

□

Next, we proceed with the proof of the transition from $\mathcal{H}_2$ to $\mathcal{H}_3$. The difference between $\mathcal{H}_2$ and $\mathcal{H}_3$ is that the challenger uses Find and CheckNoise in $\mathcal{H}_2$, and uses Find[†] and CheckNoise[†] along path t^* in $\mathcal{H}_3$. The proof of this Lemma does not rely on the secrecy of t^* . Below, we state a stronger lemma which implies Lemma 8. For any path $\mathbf{p} \in \{0, 1\}^n$, consider a game $\mathcal{G}_{\mathbf{p}}$ between a PPT adversary $\mathcal{A}$ and a challenger. The challenger picks a bit $b \in \{0, 1\}$. If $b = 0$, the challenger responds to all decryption queries using Find and CheckNoise, else it answers all decryption queries using Find[†] and CheckNoise[†] along path $\mathbf{p}$. The adversary wins if it can correctly guess b .

Lemma 11. *Assuming perfect correctness of PKE and ECC, and adaptive robustness of PRC, then for any path $\mathbf{p} \in \{0, 1\}^n$, for any PPT adversary $\mathcal{A}$, there exists a negligible function negl such that for all λ ,*

$$\Pr[\mathcal{A} \text{ wins } \mathcal{G}_{\mathbf{p}}(\lambda)] \leq \frac{1}{2} + \text{negl}(\lambda).$$

Notations. First we introduce a few notations that will be useful for our proof.

- *differing decryption query* - we call ct a differing decryption query if the decryption of ct using Dec (along path 0^n) is different from decryption of ct using Find[†] and CheckNoise[†] (along path $\mathbf{p}$).
- decryption performed using Find[†] followed by CheckNoise[†] will be referred to as Dec[†]. We adopt the convention that all variables defined via Dec[†] will be written with a dagger superscript. For example, Find[†] first computes $\tilde{t}^\dagger$, then decodes it to $s^\dagger$ and outputs $t^\dagger$.

Proof. Suppose there exists a PPT adversary $\mathcal{A}$ that makes Q decryption queries, and can win in game $\mathcal{G}_{\mathbf{p}}$ with non-negligible advantage, then it must make a *differing decryption query* ct with non-negligible probability. Let $\mathcal{E}$ denote the event that $\mathcal{A}$ makes a differing decryption query, and $\mathcal{E}_q$ the event that the q^{th} decryption query is the first differing decryption query. Therefore, it suffices to bound $\Pr[\mathcal{E}_q]$ for each q . Fix any $q \in [Q]$, we will show that $\Pr[\mathcal{E}_q] \leq \text{negl}(\lambda)$. We will further split $\mathcal{E}_q$ into the following sub-events:

- $\mathcal{E}_{q0}$: decryption using Find and CheckNoise outputs $\text{msg} \neq \perp$, Find(sk, ct) outputs t , Find[†]($\text{sk}^\dagger, \text{ct}$) = $t^\dagger$, and $t \neq t^\dagger$
- $\mathcal{E}_{q1}$: decryption using Find[†] and CheckNoise[†] outputs $\text{msg}^\dagger \neq \perp$, Find(sk, ct) outputs t , Find[†]($\text{sk}^\dagger, \text{ct}$) = $t^\dagger$, and $t \neq t^\dagger$
- $\mathcal{E}_{q2}$: Find(sk, ct) = Find[†]($\text{sk}^\dagger, \text{ct}$) = t , the output of CheckNoise is not $\perp$, but CheckNoise[†] outputs $\perp$.
- $\mathcal{E}_{q3}$: Find(sk, ct) = Find[†]($\text{sk}^\dagger, \text{ct}$) = t , the output of CheckNoise is $\perp$, but CheckNoise[†] does not output $\perp$.

We can also consider the event $\mathcal{E}_{q4}$ where Find(sk, ct) = Find[†]($\text{sk}^\dagger, \text{ct}$) = t and both CheckNoise and CheckNoise[†] produce different non- $\perp$ outputs. However, this event never occurs as both CheckNoise and CheckNoise[†] yield the identical tuple $(\text{msg}, \text{K}, (\text{pke.ct}_{i,b}, \text{com}_{i,b}))$ in Step (1), and as both procedures do not abort in Step (3), they necessarily output the same value.

Also, note that if Find(sk, ct) = Find[†]($\text{sk}^\dagger, \text{ct}$) and CheckNoise = CheckNoise[†], then the decryption in both scenarios outputs the same value. Therefore, we have $\mathcal{E}_q = \mathcal{E}_{q0} \cup \mathcal{E}_{q1} \cup \mathcal{E}_{q2} \cup \mathcal{E}_{q3}$. Therefore, it suffices to bound $\Pr[\mathcal{E}_{q0}]$, $\Pr[\mathcal{E}_{q1}]$, $\Pr[\mathcal{E}_{q2}]$ and $\Pr[\mathcal{E}_{q3}]$ separately. The events $\mathcal{E}_{q0}$ and $\mathcal{E}_{q1}$ are similar, and likewise $\mathcal{E}_{q2}$ and $\mathcal{E}_{q3}$ are similar. Below we show $\Pr[\mathcal{E}_{q0}]$ and $\Pr[\mathcal{E}_{q2}]$ are negligible.

Bounding the probability of event $\mathcal{E}_{q0}$. In this event, decryption using Find and CheckNoise outputs $\text{msg} \neq \perp$, Find(sk, ct) = t , Find[†]($\text{sk}^\dagger, \text{ct}$) = $t^\dagger$, and $t \neq t^\dagger$. We will first show that there exists an index i such that $\tilde{t}_i$ and $\tilde{t}_i^\dagger$ are different,¹⁷ $\mathbf{p}_i = 1$ and the two PRC codewords $c_{i,0}, c_{i,1}$ have low noise. This is formally captured in the following lemma.

¹⁷Recall, $\tilde{t}$ (resp. $\tilde{t}^\dagger$) is the noisy versions of t (resp. $t^\dagger$) computed in the first step of Find (resp. Find[†]).

Lemma 12. *With overwhelming probability over choice of the public/secret key, if ct is any ciphertext such that the following conditions hold:*

- decryption using Find and CheckNoise outputs $\text{msg} \neq \perp$
- $\text{Find}(\text{sk}, \text{ct}) = t$, and $\tilde{t}$ is the string produced in Step 1 of Find
- $\text{Find}^\dagger(\text{sk}, \text{ct}) = t^\dagger$ where $t^\dagger \neq t$ (note that $t^\dagger$ can be $\perp$, but t cannot be $\perp$). The string $\tilde{t}^\dagger$ is produced in Step 1 of $\text{Find}^\dagger$

then there exists an index i such that $\mathbf{p}_i = 1$, $\tilde{t}_i \neq \tilde{t}^\dagger_i$, and for both $b \in \{0, 1\}$, $\Delta(c_{i,b}, c'_{i,b}) \leq \alpha \ell_{\text{cw}}$. Here, $c'_{i,b}$ is the PRC ciphertext computed in Step 1(e) of CheckNoise.

Proof. We will now use the Lemma 9 for proving Lemma 12. Recall, for an overwhelming choice of (pk, sk) , for every $i \in [n]$ and every ct , if $\Delta(c_{i,b}, c'_{i,b}) \leq \alpha \ell_{\text{cw}}$, then $t_i = \tilde{t}_i$. Fix any such ‘good’ (pk, sk) . We are given that decryption using Find and CheckNoise outputs $\text{msg} \neq \perp$, $\text{Find}(\text{sk}, \text{ct}) = t$, $\text{Find}^\dagger(\text{sk}, \text{ct}) = \tilde{t}$ and $t \neq \tilde{t}$. Since CheckNoise does not output $\perp$ in Step 2(a), it follows that $\Delta\left(\left(c, (c_{i,b})_{i,b}, \sigma\right), \left(c', (c'_{i,b})_{i,b}, \sigma'\right)\right) \leq \delta \ell_{\text{ct}}$ which, in turn, implies $\Delta\{(c_{i,b})_{i,b}, (c'_{i,b})_{i,b}\} \leq \delta \ell_{\text{ct}}$. As a result, if

$$\mathcal{I} = \{i : \Delta(c_{i,b}, c'_{i,b}) \leq \alpha \ell_{\text{cw}} \text{ for both } b \in \{0, 1\}\}$$

then $|\mathcal{I}| \geq n(1 - \gamma)$, and for all $i \in \mathcal{I}$, $t_i = \tilde{t}_i$. We show that there exists an $i \in \mathcal{I}$ such that $\mathbf{p}_i = 1$ and $\tilde{t}_i \neq t^\dagger_i$. Assume the contrary, i.e., $\forall i \in \mathcal{I}, \tilde{t}_i = t^\dagger_i$. (We know that if $\mathbf{p}_i = 0$, the checks performed in Find and $\text{Find}^\dagger$ for this index are identical.) This implies that $\forall i \in \mathcal{I}, t_i = \tilde{t}_i$. Due to this, we have $\Delta(t, \tilde{t}) \leq \gamma n$ and due to perfect correctness of the ECC, $\tilde{t}^\dagger$ will be decoded to t , i.e., $t^\dagger = t$ which is a contradiction to the assumption. $\square$

Using Lemma 12, there exists an i such that $\mathbf{p}_i = 1$, $\tilde{t}_i \neq \tilde{t}^\dagger_i$, and for both $b \in \{0, 1\}$, $\Delta(c_{i,b}, c'_{i,b}) \leq \alpha \ell_{\text{cw}}$. We now consider two cases, first, $\tilde{t}_i = 0$ and $\tilde{t}^\dagger_i = 1$ (we denote this event by $\mathcal{E}_{q0i0}$) and second, $\tilde{t}_i = 1$ and $\tilde{t}^\dagger_i = 0$ (we denote this event by $\mathcal{E}_{q0i1}$). Note that

$$\mathcal{E}_{q0} = \bigcup_{i \in [n]} (\mathcal{E}_{q0i0} \cup \mathcal{E}_{q0i1}).$$

We start by analyzing the probability of $\mathcal{E}_{q0i0}$ ($\tilde{t}_i = 0$, $\tilde{t}^\dagger_i = 1$, decryption using Find outputs non- $\perp$, and both $c_{i,0}, c_{i,1}$ have low noise). Let us first make a few observations based on the fact that Find outputs t where $t_i = 0$, $\text{Find}^\dagger$ outputs $t^\dagger$ where $t^\dagger_i = 1$, and CheckNoise did not output $\perp$.

1. Step 1 of Find sets $\tilde{t}_i = 0$, therefore when $c_{i,0}$ was decoded using $\text{prc.sk}_{i,0}$, it produces $(\text{pke.ct}_{i,0}, \text{com}_i)$ where $\text{pke.ct}_{i,0}$ decrypts to $\text{m}_i = 1 \parallel \text{op}_i \parallel \text{ds.vk}_i$ and $\text{Open}(\text{com.pp}, 0, \text{com}_i, \text{op}_i, \text{ds.vk}_i) = \top$.
2. Likewise, since $\tilde{t}^\dagger_i$ is set to 1 in Step 1 of the decryption using $\text{Find}^\dagger$, and since $\mathbf{p}_i = 1$, we get that when $c_{i,1}$ is decoded using $\text{prc.sk}_{i,1}$, it produces $(\text{pke.ct}_{i,1}^\dagger, \text{com}_i^\dagger)$ where $\text{pke.ct}_{i,1}^\dagger$ decrypts to $\text{m}_i^\dagger = 1 \parallel \text{op}_i^\dagger \parallel \text{ds.vk}_i^\dagger$ and $\text{Open}(\text{com.pp}, 1, \text{com}_i^\dagger, \text{op}_i^\dagger, \text{ds.vk}_i^\dagger) = \top$.
3. Step 3 of CheckNoise checks that if $\Delta(c_{i,b}, c'_{i,b}) \leq \alpha \ell_{\text{cw}}$ for both $b \in \{0, 1\}$, then $\text{com}_i = \text{com}_i$ and $\text{pke.ct}_{i,0} = \text{pke.ct}_{i,0}$, where $\text{pke.ct}_{i,b}$ and com_i are obtained from the payload c .

We will now use the above observations to show that if the above case happens with non-negligible probability, then the base PRC scheme does not satisfy adaptive robustness. First, we observe that with overwhelming

probability over the choice of public parameters, $\text{com}_i^\dagger \neq \text{com}_i$. This is due to the statistical binding property of the commitment scheme.¹⁸

Therefore, it only remains to handle the case where $\text{com}_i^\dagger \neq \text{com}_i$. Note that $\text{com}_i = \text{com}_i$, and thus $\text{com}_i^\dagger \neq \text{com}_i$. Note that

1. $\Delta(c_{i,1}, c'_{i,1}) \leq \alpha \ell_{\text{cw}}$ (from the description of the first case)
2. $c'_{i,1} = \text{PRC.Encode}(\text{prc.pk}_{i,1}, (\text{pke.ct}_{i,1}, \text{com}_i); \text{PRF.Eval}(\text{K}, (i, 1)))$ (from Step 1e of CheckNoise)
3. $\text{PRC.Decode}(\text{prc.sk}_{i,1}, c_{i,1}) = (\text{pke.ct}_{i,1}^\dagger, \text{com}_i^\dagger) \neq (\text{pke.ct}_{i,1}, \text{com}_i)$.

From here, we get an attack on adaptive robustness. If an adversary $\mathcal{A}$, with non-negligible probability, is able to trigger the event $\mathcal{E}_{q0i0}$, then we can construct an adversary that breaks adaptive robustness of PRC as follows.

1. The PRC challenger sends pk^* to the reduction $\mathcal{B}$.
2. It simulates $\mathcal{H}_2$ by setting $\text{prc.pk}_{i,1} = \text{pk}^*$ and chooses the remaining public/secret keys by itself.
3. On receiving the q -th query $(c^{(q)}, \{c_{i,b}^{(q)}\}, \sigma^{(q)})$ from $\mathcal{A}$, it runs PRC.Find and PRC.CheckNoise to obtain $(\text{msg}^{(q)}, \text{K}^{(q)}, (\text{pke.ct}_{i,b}^{(q)}, \text{com}_i^{(q)})_{i,b}, \text{ds}.\sigma^{(q)})$ and returns $c_{i,1}^{(q)}$ as the codeword, $(\text{pke.ct}_{i,1}^{(q)}, \text{com}_i^{(q)})$ as the message, and $\text{PRF.Eval}(\text{K}^{(q)}, (i, 1))$ as the randomness to the challenger.

We observe that if an adversary $\mathcal{A}$ triggers the mentioned event in the q -th query for position i , then the adversary $\mathcal{B}$ forwards a valid adaptive robustness attack to the underlying challenger. Therefore, if the event $\mathcal{E}_{q0i0}$ happens with non-negligible probability, then the reduction also wins the adaptive robustness game with the same probability.

Next, we analyze the event $\mathcal{E}_{q0i1}$ ($\tilde{t}_i = 1, \tilde{t}_i^\dagger = 0$, decryption using Find outputs non- $\perp$, and both $c_{i,0}, c_{i,1}$ have low noise). Again, we first make a few observations based on the fact that Find outputs t where $t_i = 1$, Find † outputs $t^\dagger$ where $t_i^\dagger = 0$, and CheckNoise did not output $\perp$.

1. Step 1 of Find sets $\tilde{t}_i = 1$, therefore when $c_{i,0}$ is decoded using $\text{prc.sk}_{i,0}$, it produces $\text{pke.ct}_{i,0}$ and com_i , and when $\text{pke.ct}_{i,0}$ is decrypted using $\text{pke.sk}_{i,0}$, either it does not produce a decryption of the form $m_i = 1 \parallel \text{op}_i \parallel \text{ds.vk}_i$, or $\text{Open}(\text{com.pp}, 0, \text{com}_i, \text{op}_i, \text{ds.vk}_i) = \perp$.
2. Step 3 of CheckNoise checks that if $\Delta(c_{i,b}, c'_{i,b}) \leq \alpha \ell_{\text{cw}}$ for both $b \in \{0, 1\}$, then $\text{com}_i = \text{com}_i$ and $\text{pke.ct}_{i,1} = \text{pke.ct}_{i,1}$, where $\text{pke.ct}_{i,b}$ and com_i are obtained from the payload c .
3. In Step 2(b) of CheckPKE, it is checked that $m_i = 1 \parallel \text{op}_i \parallel \text{ds.vk}_i$ and $\text{Open}(\text{com.pp}, 1, \text{com}_i, \text{op}_i, \text{ds.vk}_i) = \top$ where $m_i = \text{PKE.Recover}(\text{pke.pk}_{i,1}, \text{pke.ct}_{i,1}, r_i)$ and also checked that $\text{pke.ct}_{i,1} = \text{PKE.Enc}(\text{pke.pk}_{i,1}, m_i; r_i)$.
4. Next, Find † sets $\tilde{t}_i = 0$ and $\text{p}_i = 1$, either $\underline{m}_i^\dagger$ is not of the form $1 \parallel \text{op}_i^\dagger \parallel \text{ds.vk}_i^\dagger$ or $\text{Open}(\text{com.pp}, 1, \text{com}_i^\dagger, \text{op}_i^\dagger, \text{ds.vk}_i^\dagger) = \perp$ (here $\text{pke.ct}_{i,1}^\dagger, \text{com}_i^\dagger$ is obtained from $\text{PRC.Decode}(\text{prc.sk}_{i,1}, c_{i,1})$ and $\underline{m}_i^\dagger$ from $\text{PKE.Dec}(\text{pke.sk}_{i,1}, \text{pke.ct}_{i,1}^\dagger)$).
5. If the pair $(\text{pke.ct}_{i,1}^\dagger, \text{com}_i^\dagger)$ (which is obtained by decrypting $c_{i,1}$ using $\text{prc.sk}_{i,1}$) is equal to the pair $(\text{pke.ct}_{i,1}, \text{com}_i)$ (obtained from the payload c in the CheckNoise procedure) then this violates the perfect correctness. This follows from the observations listed above.

Therefore, it only remains to handle the case where $(\text{pke.ct}_{i,1}^\dagger, \text{com}_i^\dagger) \neq (\text{pke.ct}_{i,1}, \text{com}_i)$. Since, $\Delta(c_{i,1}, c'_{i,1}) \leq \alpha \ell_{\text{cw}}$ and the fact that $c'_{i,1}$ is an encoding of $(\text{pke.ct}_{i,1}, \text{com}_i)$ computed in Step 1 of CheckNoise, then we use adaptive robustness to show that this event happens with negligible probability. More precisely, if an adversary $\mathcal{A}$, with non-negligible probability, is able to trigger this event, then we can construct an adversary that breaks adaptive robustness of PRC as follows.

¹⁸Note that the commitments can potentially be for different tags. However, since the commitment to 0 is independent of the tag, the statistical binding property is violated.

1. The PRC challenger outputs pk^* to the reduction $\mathcal{B}$.
2. It simulates $\mathcal{H}_2$ by setting $\text{prc.pk}_{i,1} = \text{pk}^*$ (and samples the remaining public/secret keys by itself).
3. On receiving the q -th query $(c^{(q)}, \{c_{i,b}^{(q)}\}, \sigma^{(q)})$ from $\mathcal{A}$, it runs PRC.Find and PRC.CheckNoise to obtain $(\text{msg}^{(q)}, K^{(q)}, (\text{pke.ct}_{i,b}^{(q)}, \text{com}_i^{(q)})_{i,b}, \text{ds}.\sigma^{(q)})$ and returns $c_{i,1}^{(q)}$ as the codeword, $(\text{pke.ct}_{i,1}^{(q)}, \text{com}_i^{(q)})$ as the message, and $\text{PRF.Eval}(K^{(q)}, (i, 1))$ as the randomness to the challenger.

We observe that if an adversary $\mathcal{A}$ triggers the mentioned event in the q -th query for position i , then the adversary $\mathcal{B}$ forwards a valid adaptive robustness attack to the underlying challenger. Given that this only happens if $\mathcal{B}$ guesses the correct index i and query q , this introduces an additional security loss of $n \cdot Q$.

Bounding the probability of event $\mathcal{E}_{q1}$. In this event, decryption using $\text{Find}^\dagger$ and $\text{CheckNoise}^\dagger$ outputs $\text{msg}^\dagger \neq \perp$, $\text{Find}(\text{sk}, \text{ct}) = t$, and $t \neq t^\dagger$. This case is analogous to $\mathcal{E}_{q0}$ and the proof is almost identical; the lemmas and claims will be with respect to decryption using $\text{Find}^\dagger$ and $\text{CheckNoise}^\dagger$.

Bounding the probability of event $\mathcal{E}_{q2}$. In this event, both Find and $\text{Find}^\dagger$ output the same string t , CheckNoise does not output $\perp$, but $\text{CheckNoise}^\dagger$ outputs $\perp$. Since Find and $\text{Find}^\dagger$ output the same t , the computation in Step 1 of CheckNoise and $\text{CheckNoise}^\dagger$ is identical, and therefore,

$$\begin{aligned}
& - (\text{com}_i)_i = (\text{com}_i^\dagger)_i, \\
& - (\text{pke.ct}_{i,b})_{i,b} = (\text{pke.ct}_{i,b}^\dagger)_{i,b}, \\
& - (c'_{i,b})_{i,b} = (c_{i,b}^\dagger)_{i,b}
\end{aligned}$$

However, the difference in CheckNoise and $\text{CheckNoise}^\dagger$ can arise from Step 3. CheckNoise decrypts $(c_{i,0})_i$ using the secret keys $(\text{prc.sk}_{i,0})_i$ and obtains $(\text{pke.ct}_{i,0}, \text{com}_i)_i$ while $\text{CheckNoise}^\dagger$ decrypts $(c_{i,\mathbf{p}_i})_i$ using the secret keys $(\text{prc.sk}_{i,\mathbf{p}_i})_i$ and obtains $(\text{pke.ct}_{i,\mathbf{p}_i}^\dagger, \text{com}_i^\dagger)_i$. Since CheckNoise does not output $\perp$ but $\text{CheckNoise}^\dagger$ outputs $\perp$, there exists an index i such that

$$\begin{aligned}
& - \mathbf{p}_i = 1 \\
& - \Delta(c_{i,b}, c'_{i,b}) \leq \delta \ell_{\text{ct}} \text{ for both } b \in \{0, 1\} \\
& - \Delta(c_{i,b}, c_{i,b}^\dagger) \leq \delta \ell_{\text{ct}} \text{ for both } b \in \{0, 1\} \text{ (since } c'_{i,b} = c_{i,b}^\dagger) \\
& - \text{PRC.Decode}(\text{prc.sk}_{i,0}, c_{i,0}) = (\text{pke.ct}_{i,0}, \text{com}_i)_i = (\text{pke.ct}_{i,0}, \text{com}_i) \text{ (since } \text{CheckNoise} \text{ does not output } \perp \text{ in Step 3b)} \\
& - \text{PRC.Decode}(\text{prc.sk}_{i,\mathbf{p}_i}, c_{i,\mathbf{p}_i}) = (\text{pke.ct}_{i,\mathbf{p}_i}^\dagger, \text{com}_i^\dagger)_i \neq (\text{pke.ct}_{i,\mathbf{p}_i}, \text{com}_i) \text{ (since } \text{CheckNoise}^\dagger \text{ outputs } \perp \text{ in Step 3b)}
\end{aligned}$$

Using the first, third and fifth lines above, we get that

$$\begin{aligned}
& - \Delta(c_{i,1}, c_{i,1}^\dagger) \leq \alpha \ell_{\text{cw}}, \\
& - \text{PRC.Decode}(\text{prc.sk}_{i,1}, c_{i,1}) \neq \text{PRC.Decode}(\text{prc.sk}_{i,1}, c_{i,1}^\dagger).
\end{aligned}$$

Therefore, we obtain an attack on adaptive robustness if event $\mathcal{E}_{q2}$ happens with non-negligible probability. The reduction algorithm receives the PRC public key prc.pk^* from the challenger. It samples $2n - 1$ PRC public/secret keys, and sets $\text{prc.pk}_{i,\mathbf{p}_i} = \text{prc.pk}^*$. For the first $q - 1$ decryption queries, it uses $(\text{prc.sk}_{i,0}, \text{pke.sk}_{i,0})$ to decrypt (note that the first $q - 1$ decryptions can be performed along either path $\mathbf{p}$ or the path 0^n). On receiving the q^{th} decryption query $(c, (c_{i,0}, c_{i,1}))$, the reduction first computes $t = \text{Find}((\text{sk}_{i,0})_i, (c_{i,0})_i)$. Next, it computes the various components in Step 1 of CheckNoise (which are same as the components computed

in Step 1 of CheckNoise[†]). The reduction finds an i that satisfies the five conditions listed above. It sends $c_{i,1}$ as the noisy codeword, $(\text{pke.ct}_{i,p_i}, \text{com}_i)$ as the message and $\text{PRF.Eval}(K, (i, 1))$ as the randomness to the adaptive robustness challenger.

□
□

Lemma 13. *Assuming PKE is IND-CPA secure, then for all PPT adversaries $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all λ , $|\Pr[\mathcal{H}_3(\lambda, \mathcal{A}) = 1] - \Pr[\mathcal{H}_4(\lambda, \mathcal{A}) = 1]| \leq \text{negl}(\lambda)$.*

Proof. To prove the transition from $\mathcal{H}_3$ to $\mathcal{H}_4$, we introduce a set of intermediate hybrids $\mathcal{H}_{3,0}, \dots, \mathcal{H}_{3,n}$. Here, $\mathcal{H}_{3,j}$ is defined as the hybrid where the first j ciphertexts $\text{pke.ct}_{1,t_1^*}^*, \dots, \text{pke.ct}_{j,t_j^*}^*$ are generated as $\text{pke.ct}_{i,t_i^*}^* := \text{PKE.Enc}(\text{pke.pk}_{i,t_i^*}, 1 \parallel \text{op}_{i,t_i^*}^* \parallel \text{ds.vk}^*; r_i^*)$ for all $i \in [j]$ and the remaining ciphertexts $\text{pke.ct}_{j+1,t_{j+1}^*}^*, \dots, \text{pke.ct}_{n,t_n^*}^*$ are generated as $\text{pke.ct}_{i,t_i^*}^* := \text{PKE.Enc}(\text{pke.pk}_{i,t_i^*}, 0^{1+\ell_{\text{vk}}+\ell_{\text{com}}}; r_i^*)$ for all $i \in [n] \setminus [j]$. To conclude the transition, it remains to show that $\mathcal{H}_{3,i-1} \approx_c \mathcal{H}_{3,i}$ for all $i \in [n]$, which we prove for a general index i^* in the next step. We prove the transition from $\mathcal{H}_{3,i^*-1}$ to $\mathcal{H}_{3,i^*}$ by contradiction. In more detail, we assume that there exists an adversary $\mathcal{A}$ that is able to distinguish between the hybrids $\mathcal{H}_{3,i^*-1}$ and $\mathcal{H}_{3,i^*}$ and turn it into an adversary $\mathcal{B}$ that violates the IND-CPA security of the PKE scheme. The reduction $\mathcal{B}$ proceeds as follows, it proceeds as described in the initialization of $\mathcal{H}_3$ until the generation of the ciphertexts, with the exception that it does not generate the public key $\text{pke.pk}_{i^*,t_{i^*}^*}$ itself but receives it from the PKE challenger. Now, we distinguish between three cases: first, the ciphertexts $\{\text{pke.ct}_{i,t_i^*}^*\}_{i \in [i^*-1]}$, second, $\text{pke.ct}_{i^*,t_{i^*}^*}^*$, and, third, $\{\text{pke.ct}_{i,t_i^*}^*\}_{i \in [n] \setminus [i^*]}$. The ciphertexts $\{\text{pke.ct}_{i,t_i^*}^*\}_{i \in [i^*-1]}$ are generated as $\{\text{PKE.Enc}(\text{pke.pk}_{i,t_i^*}, 1 \parallel \text{op}_{i,t_i^*}^* \parallel \text{ds.vk}^*; r_i^*)\}_{i \in [i^*-1]}$ by $\mathcal{B}$ and the ciphertexts $\{\text{pke.ct}_{i,t_i^*}^*\}_{i \in [n] \setminus [i^*]}$ are generated as $\{\text{PKE.Enc}(\text{pke.pk}_{i,t_i^*}, 0^{1+\ell_{\text{vk}}+\ell_{\text{com}}}; r_i^*)\}_{i \in [n] \setminus [i^*]}$ by $\mathcal{B}$. To generate the ciphertext $\text{pke.ct}_{i^*,t_{i^*}^*}^*$, $\mathcal{B}$ submits $((0^{1+\ell_{\text{vk}}+\ell_{\text{com}}}), (1 \parallel \text{op}_{i^*,t_{i^*}^*}^* \parallel \text{ds.vk}^*))$ to its underlying challenger and receives $\text{pke.ct}_{i^*,t_{i^*}^*}^*$ in reply. Afterwards, $\mathcal{B}$ sends all those ciphertexts to the adversary $\mathcal{A}$ and proceeds as described in $\mathcal{H}_3$.

We observe that if the underlying challenger encrypts $0^{1+\ell_{\text{vk}}+\ell_{\text{com}}}$ as the ciphertext $\text{pke.ct}_{i^*,t_{i^*}^*}^*$, then $\mathcal{B}$ simulates hybrid $\mathcal{H}_{3,i^*-1}$ towards $\mathcal{A}$. In the case that the underlying challenger encrypts $1 \parallel \text{op}_{i^*,t_{i^*}^*}^* \parallel \text{ds.vk}^*$ as the ciphertext $\text{pke.ct}_{i^*,t_{i^*}^*}^*$, then $\mathcal{B}$ simulates hybrid $\mathcal{H}_{3,i^*}$ towards $\mathcal{A}$. This results in a perfect reduction and, therefore, the success probability of $\mathcal{A}$ is being bounded by the security advantage of IND-CPA of PKE. Since the above reduction works for all $i^* \in [n]$, we obtain the indistinguishability between the hybrids $\mathcal{H}_3$ and $\mathcal{H}_4$, bounded by n times the security advantage of IND-CPA. □

Lemma 14. *Assuming the security of the hinting PRG and the PRG security, then for all PPT adversaries $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all λ , $|\Pr[\mathcal{H}_5(\lambda, \mathcal{A}) = 1] - \Pr[\mathcal{H}_6(\lambda, \mathcal{A}) = 1]| \leq \text{negl}(\lambda)$.*

Proof. To prove this lemma, we introduce an intermediate hybrid $\mathcal{H}_{5,5}$. This hybrid is identical to $\mathcal{H}_5$, except that the values $\tilde{r}_i^*$ are sampled uniformly at random instead of being generated by the hinting PRG.

We will first show that $\mathcal{H}_5$ and $\mathcal{H}_{5,5}$ are indistinguishable by contradiction. In more detail, we assume that there exists an adversary $\mathcal{A}$ that is able to distinguish between the hybrids $\mathcal{H}_5$ and $\mathcal{H}_{5,5}$ and turn it into an adversary $\mathcal{B}$ that violates the security of the hinting PRG. The reduction $\mathcal{B}$ proceeds as follows, in the first step, it receives hprg.pp and $\tilde{r}_0^*, \{\tilde{r}_{i,b}^*\}_{i \in [n], b \in \{0,1\}}$ from the challenger. It then proceeds as described in the description of $\mathcal{H}_5$ until the adversary $\mathcal{A}$ asks its challenge encryption query using a message m . To generate the corresponding PKE ciphertexts, the adversary $\mathcal{B}$ proceeds as described in $\mathcal{H}_5$ while it uses the values $\tilde{r}_0^*, \{\tilde{r}_{i,b}^*\}_{i \in [n], b \in \{0,1\}}$ obtained from the challenger instead of sampling them itself. Afterwards, $\mathcal{B}$ proceeds again as in the description of $\mathcal{H}_5$.

We observe that if the underlying challenger generates the values $\tilde{r}_0^*$ and $\{\tilde{r}_{i,b}^*\}_{i \in [n], b \in \{0,1\}}$ using the hinting PRG HPRG, then $\mathcal{B}$ simulates hybrid $\mathcal{H}_5$ towards $\mathcal{A}$. In the case that the underlying challenger samples the values $\tilde{r}_0^*$ and $\{\tilde{r}_{i,b}^*\}_{i \in [n], b \in \{0,1\}}$ at random, $\mathcal{B}$ simulates $\mathcal{H}_{5,5}$ towards $\mathcal{A}$. This results in a perfect reduction and, therefore, the success probability of $\mathcal{A}$ is being bounded by the security advantage of HPRG.

To show $\mathcal{H}_6$ and $\mathcal{H}_{5.5}$ are indistinguishable, we rely on the security of the PRG. The only distinction lies in the way c^* and σ^* are generated. In $\mathcal{H}_{5.5}$, c^* is set to $\text{PRG}(\tilde{r}_0^*[0]) + \text{ECC}(m^*, K^*, \{\text{pke.ct}_{i,b}^*, c_{i,2}^*\}_{i \in [n], b \in \{0,1\}}) \cdot B_{\text{pl}}$ and σ^* is set to $\text{PRG}(\tilde{r}_0^*[1]) \oplus \text{ECC.Encode}(\text{ecc.pp}, \text{ds}.\sigma^*)$. Whereas, in $\mathcal{H}_6$, they are sampled uniformly at random. Since, the value $\tilde{r}_0^*$ is not used elsewhere in either experiment, the security of PRG guarantees that no PPT adversary can distinguish between $\mathcal{H}_{5.5}$ and $\mathcal{H}_6$ with more than negligible advantage. $\square$

Lemma 15. *Assuming the security of the PRF, then for all PPT adversaries $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all λ , $|\Pr[\mathcal{H}_6(\lambda, \mathcal{A}) = 1] - \Pr[\mathcal{H}_7(\lambda, \mathcal{A}) = 1]| \leq \text{negl}(\lambda)$.*

Proof. We prove this lemma by contradiction. In more detail, we assume that there exists an adversary $\mathcal{A}$ that is able to distinguish between the hybrids $\mathcal{H}_6$ and $\mathcal{H}_7$ and turn it into an adversary $\mathcal{B}$ that violates the security of PRF. The reduction $\mathcal{B}$ proceeds as follows - in the initialization step, it proceeds as described in the description of $\mathcal{H}_6$ until the generation of the PRC ciphertexts. The adversary $\mathcal{B}$ sends (i, b) for all $i \in [n], b \in \{0, 1\}$ and obtains from the PRF challenger, $u_{i,b}$ instead of sampling them itself to generate $c_{i,b}^*$. Afterwards, $\mathcal{B}$ proceeds again as in the description of $\mathcal{H}_6$. It is important to note that the PRF key is not required because c^* is sampled uniformly at random.

We observe that if the underlying challenger generates the values $u_{i,b}$ using the PRF, then $\mathcal{B}$ simulates hybrid $\mathcal{H}_6$ towards $\mathcal{A}$. In the case that the underlying challenger samples the values at random, $\mathcal{B}$ simulates $\mathcal{H}_7$ towards $\mathcal{A}$. This results in a perfect reduction and, therefore, the success probability of $\mathcal{A}$ is being bounded by the security advantage of PRF. $\square$

Lemma 16. *Assuming PRC is IND-CPA secure, then for all PPT adversaries $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all λ , $|\Pr[\mathcal{H}_7(\lambda, \mathcal{A}) = 1] - \Pr[\mathcal{H}_8(\lambda, \mathcal{A}) = 1]| \leq \text{negl}(\lambda)$.*

Proof. To prove the transition from $\mathcal{H}_7$ to $\mathcal{H}_8$, we introduce a set of intermediate hybrids $\mathcal{H}_{7,0}, \dots, \mathcal{H}_{7,n}$. Here, $\mathcal{H}_{7,j}$ is defined as the hybrid where the first j ciphertexts $c_{1,1}^*, \dots, c_{j,1}^*$ are generated as random values, i.e., $c_{i,1}^* \leftarrow \{0, 1\}^{\ell_{\text{ct}}}$ for all $i \in [j]$, and the remaining ciphertexts $c_{j+1,1}^*, \dots, c_{n,1}^*$ are generated as $c_{i,1}^* := \text{PRC.Encode}(\text{prc.pk}_{i,1}, (\text{pke.ct}_{i,1}, \text{com}_i^*); u_{i,1})$ for all $i \in [n] \setminus [j]$. To conclude the transition, it remains to show that $\mathcal{H}_{7,i-1} \approx_c \mathcal{H}_{7,i}$ for all $i \in [n]$, which we prove for an index i^* in the next step.

We prove the transition from $\mathcal{H}_{7,i^*-1}$ to $\mathcal{H}_{7,i^*}$ by contradiction. In more detail, we assume that there exists an adversary $\mathcal{A}$ that is able to distinguish between the hybrids $\mathcal{H}_{7,i^*-1}$ and $\mathcal{H}_{7,i^*}$ and turn it into an adversary $\mathcal{B}$ that violates the IND-CPA security of the PRC scheme. The reduction $\mathcal{B}$ proceeds as follows, it proceeds as described in the initialization of $\mathcal{H}_7$ until the generation of the ciphertexts, with the exception that it does not generate the public key $\text{prc.pk}_{i^*,1}$ itself but receives it from the underlying challenger. Now, we distinguish between three cases: first, the ciphertexts $\{c_{i,1}^*\}_{i \in [i^*-1]}$, second, $c_{i^*,1}^*$, and, third, $\{c_{i,1}^*\}_{i \in [n] \setminus [i^*-1]}$. The ciphertexts $\{c_{i,1}^*\}_{i \in [i^*-1]}$ are sampled at random, i.e., $c_{i,1}^* \leftarrow \{0, 1\}^{\ell_{\text{ct}}}$ for all $i \in [i^*-1]$, by $\mathcal{B}$ and the ciphertexts $\{c_{i,1}^*\}_{i \in [n] \setminus [i^*-1]}$ are generated as $\{\text{PRC.Encode}(\text{prc.pk}_{i,1}, (\text{pke.ct}_{i,1}, \text{com}_i^*); u_{i,1})\}_{i \in [n] \setminus [i^*-1]}$ by $\mathcal{B}$. To generate the ciphertext $c_{i^*,1}^*$, $\mathcal{B}$ submits $(\text{pke.ct}_{i^*,1}, \text{com}_{i^*}^*)$ to the PRC challenger and receives $c_{i^*,1}^*$ in reply. Afterwards, $\mathcal{B}$ sends all those ciphertexts to the adversary $\mathcal{A}$ and proceeds as described in $\mathcal{H}_7$.

We observe that if the underlying challenger encodes $(\text{pke.ct}_{i^*,1}, \text{com}_{i^*}^*)$ as the ciphertext $c_{i^*,1}^*$, then $\mathcal{B}$ simulates hybrid $\mathcal{H}_{7,i^*-1}$ towards $\mathcal{A}$. In the case that the underlying challenger sends a random value $c_{i^*,1}^* \leftarrow \{0, 1\}^{\ell_{\text{ct}}}$ as the ciphertext, then $\mathcal{B}$ simulates hybrid $\mathcal{H}_{7,i^*}$ towards $\mathcal{A}$. This results in a perfect reduction and, therefore, the success probability of $\mathcal{A}$ is being bounded by the security advantage of IND-CPA of PRC. Since the above reduction works for all $i^* \in [n]$, we obtain the indistinguishability between the hybrids $\mathcal{H}_7$ and $\mathcal{H}_8$, bounded by n times the security advantage of IND-CPA. $\square$

Acknowledgements

We thank the anonymous Crypto 2026 reviewers for their helpful feedback, and in particular one reviewer for suggesting Raw Reed–Solomon codes as a useful cryptanalytic target for noisy McEliece-type assumptions.

Nico Döttling and Mahesh Sreekumar Rajasree: Funded by the European Union (ERC, LACONIC, 101041207). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. Neither the European Union nor the granting authority can be held responsible for them. This work has been partly done while the third author was at IIT Delhi Abu Dhabi and the CISPA Helmholtz Center for Information Security.

References

- AAC⁺25. Omar Alrabiah, Prabhanjan Ananth, Miranda Christ, Yevgeniy Dodis, and Sam Gunn. Ideal pseudorandom codes. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, STOC '25, page 1638–1647, New York, NY, USA, 2025. Association for Computing Machinery.
- ACPS09. Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In Shai Halevi, editor, *Advances in Cryptology – CRYPTO 2009*, volume 5677 of *Lecture Notes in Computer Science*, pages 595–618, Santa Barbara, CA, USA, August 16–20, 2009. Springer Berlin Heidelberg, Germany.
- Ale03. Michael Alekhnovich. More on average case vs approximation complexity. In *44th Annual Symposium on Foundations of Computer Science*, pages 298–307, Cambridge, MA, USA, October 11–14, 2003. IEEE Computer Society Press.
- AP22. Navid Alamati and Sikhar Patranabis. Cryptographic primitives with hinting property. In Shweta Agrawal and Dongdai Lin, editors, *Advances in Cryptology – ASIACRYPT 2022, Part I*, volume 13791 of *Lecture Notes in Computer Science*, pages 33–62, Taipei, Taiwan, December 5–9, 2022. Springer, Cham, Switzerland.
- BBDP01. Mihir Bellare, Alexandra Boldyreva, Anand Desai, and David Pointcheval. Key-privacy in public-key encryption. In Colin Boyd, editor, *Advances in Cryptology – ASIACRYPT 2001*, volume 2248 of *Lecture Notes in Computer Science*, pages 566–582, Gold Coast, Australia, December 9–13, 2001. Springer Berlin Heidelberg, Germany.
- BBP04. Mihir Bellare, Alexandra Boldyreva, and Adriana Palacio. An uninstantiable random-oracle-model scheme for a hybrid-encryption problem. In Christian Cachin and Jan Camenisch, editors, *Advances in Cryptology – EUROCRYPT 2004*, volume 3027 of *Lecture Notes in Computer Science*, pages 171–188, Interlaken, Switzerland, May 2–6, 2004. Springer Berlin Heidelberg, Germany.
- BC24. Daniel J. Bernstein and Tung Chou. CryptAttackTester: high-assurance attack analysis. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology – CRYPTO 2024, Part VI*, volume 14925 of *Lecture Notes in Computer Science*, pages 141–182, Santa Barbara, CA, USA, August 18–22, 2024. Springer, Cham, Switzerland.
- BCHK07. Dan Boneh, Ran Canetti, Shai Halevi, and Jonathan Katz. Chosen-ciphertext security from identity-based encryption. *SIAM J. Comput.*, 36(5):1301–1328, 2007.
- BDD22. Pedro Branco, Nico Döttling, and Jesko Dujmovic. Rate-1 incompressible encryption from standard assumptions. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022: 20th Theory of Cryptography Conference, Part II*, volume 13748 of *Lecture Notes in Computer Science*, pages 33–69, Chicago, IL, USA, November 7–10, 2022. Springer, Cham, Switzerland.
- BG10. Zvika Brakerski and Shafi Goldwasser. Circular and leakage resilient public-key encryption under subgroup indistinguishability - (or: Quadratic residuosity strikes back). In Tal Rabin, editor, *Advances in Cryptology – CRYPTO 2010*, volume 6223 of *Lecture Notes in Computer Science*, pages 1–20, Santa Barbara, CA, USA, August 15–19, 2010. Springer Berlin Heidelberg, Germany.
- BHHO08. Dan Boneh, Shai Halevi, Michael Hamburg, and Rafail Ostrovsky. Circular-secure encryption from decision Diffie-Hellman. In David Wagner, editor, *Advances in Cryptology – CRYPTO 2008*, volume 5157 of *Lecture Notes in Computer Science*, pages 108–125, Santa Barbara, CA, USA, August 17–21, 2008. Springer Berlin Heidelberg, Germany.
- BK05. Dan Boneh and Jonathan Katz. Improved efficiency for CCA-secure cryptosystems built using identity-based encryption. In Alfred Menezes, editor, *Topics in Cryptology – CT-RSA 2005*, volume 3376 of *Lecture Notes in Computer Science*, pages 87–103, San Francisco, CA, USA, February 14–18, 2005. Springer Berlin Heidelberg, Germany.
- BKR23. Andrej Bogdanov, Pravesh K Kothari, and Alon Rosen. Public-key encryption, local pseudorandom generators, and the low-degree method. In *Theory of Cryptography Conference*, pages 268–285. Springer, 2023.

- BLP08. Daniel J Bernstein, Tanja Lange, and Christiane Peters. Attacking and defending the McEliece cryptosystem. In *International Workshop on Post-Quantum Cryptography*, pages 31–46. Springer, 2008.
- BLSV18. Zvika Brakerski, Alex Lombardi, Gil Segev, and Vinod Vaikuntanathan. Anonymous IBE, leakage resilience and circular security from new assumptions. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018, Part I*, volume 10820 of *Lecture Notes in Computer Science*, pages 535–564, Tel Aviv, Israel, April 29 – May 3, 2018. Springer, Cham, Switzerland.
- BR93. Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *ACM CCS 93: 1st Conference on Computer and Communications Security*, pages 62–73, Fairfax, Virginia, USA, November 3–5, 1993. ACM Press.
- CC98. Anne Canteaut and Florent Chabaud. A new algorithm for finding minimum-weight words in a linear code: Application to McEliece’s cryptosystem and to narrow-sense BCH codes of length 511. *IEEE Transactions on Information Theory*, 44(1):367–378, 1998.
- CG24. Miranda Christ and Sam Gunn. Pseudorandom error-correcting codes. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology – CRYPTO 2024, Part VI*, volume 14925 of *Lecture Notes in Computer Science*, pages 325–347, Santa Barbara, CA, USA, August 18–22, 2024. Springer, Cham, Switzerland.
- CGG⁺25. Miranda Christ, Noah Golowich, Sam Gunn, Ankur Moitra, and Daniel Wichs. Improved Pseudorandom Codes from Permuted Puzzles. Cryptology ePrint Archive, Paper 2025/2222, 2025. To appear in STOC 2026.
- CGGU⁺14. Alain Couvreur, Philippe Gaborit, Valérie Gauthier-Umaña, Ayoub Otmani, and Jean-Pierre Tillich. Distinguisher-based attacks on public-key cryptosystems using Reed–Solomon codes. *Designs, Codes and Cryptography*, 73(2):641–666, 2014.
- CHK03. Ran Canetti, Shai Halevi, and Jonathan Katz. A forward-secure public-key encryption scheme. In Eli Biham, editor, *Advances in Cryptology – EUROCRYPT 2003*, volume 2656 of *Lecture Notes in Computer Science*, pages 255–271, Warsaw, Poland, May 4–8, 2003. Springer Berlin Heidelberg, Germany.
- CHS25. Aloni Cohen, Alexander Hoover, and Gabe Schoenbach. Watermarking language models for many adaptive users. In Marina Blanton, William Enck, and Cristina Nita-Rotaru, editors, *IEEE Symposium on Security and Privacy, SP 2025, San Francisco, CA, USA, May 12–15, 2025*, pages 2583–2601. IEEE, 2025.
- CMCP17. Alain Couvreur, Irene Márquez-Corbella, and Ruud Pellikaan. Cryptanalysis of McEliece cryptosystem based on algebraic geometry codes and their subcodes. *IEEE Transactions on Information Theory*, 63(8):5404–5418, 2017.
- CMT23. Alain Couvreur, Rocco Mora, and Jean-Pierre Tillich. A new approach based on quadratic forms to attack the McEliece cryptosystem. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 3–38. Springer, 2023.
- COT16. Alain Couvreur, Ayoub Otmani, and Jean-Pierre Tillich. Polynomial time attack on wild McEliece over quadratic extensions. *IEEE Transactions on Information Theory*, 63(1):404–427, 2016.
- COTGU15. Alain Couvreur, Ayoub Otmani, Jean-Pierre Tillich, and Valérie Gauthier-Umana. A polynomial-time attack on the BBCRS scheme. In *IACR International Workshop on Public Key Cryptography*, pages 175–193. Springer, 2015.
- CS98. Ronald Cramer and Victor Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In Hugo Krawczyk, editor, *Advances in Cryptology – CRYPTO’98*, volume 1462 of *Lecture Notes in Computer Science*, pages 13–25, Santa Barbara, CA, USA, August 23–27, 1998. Springer Berlin Heidelberg, Germany.
- CS02. Ronald Cramer and Victor Shoup. Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In Lars R. Knudsen, editor, *Advances in Cryptology – EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 45–64, Amsterdam, The Netherlands, April 28 – May 2, 2002. Springer Berlin Heidelberg, Germany.
- DDMN12. Nico Döttling, Rafael Dowsley, Jörn Müller-Quade, and Anderson C. A. Nascimento. A CCA2 secure variant of the McEliece cryptosystem. *IEEE Trans. Inf. Theory*, 58(10):6672–6680, 2012.
- DDN91. Danny Dolev, Cynthia Dwork, and Moni Naor. Non-malleable cryptography (extended abstract). In *23rd Annual ACM Symposium on Theory of Computing*, pages 542–552, New Orleans, LA, USA, May 6–8, 1991. ACM Press.
- DGHM18. Nico Döttling, Sanjam Garg, Mohammad Hajiabadi, and Daniel Masny. New constructions of identity-based and key-dependent message secure encryption schemes. In Michel Abdalla and Ricardo Dahab, editors, *PKC 2018: 21st International Conference on Theory and Practice of Public Key Cryptography*,

- Part I*, volume 10769 of *Lecture Notes in Computer Science*, pages 3–31, Rio de Janeiro, Brazil, March 25–29, 2018. Springer, Cham, Switzerland.
- DMN12. Nico Döttling, Jörn Müller-Quade, and Anderson C. A. Nascimento. IND-CCA secure cryptography based on a variant of the LPN problem. In Xiaoyun Wang and Kazue Sako, editors, *Advances in Cryptology – ASIACRYPT 2012*, volume 7658 of *Lecture Notes in Computer Science*, pages 485–503, Beijing, China, December 2–6, 2012. Springer Berlin Heidelberg, Germany.
- DMR25. Nico Döttling, Anne Müller, and Mahesh Sreeksar Rajasree. Separating pseudorandom codes from local oracles. *Cryptology ePrint Archive*, Paper 2025/1020, 2025.
- DNR04a. Cynthia Dwork, Moni Naor, and Omer Reingold. Immunizing encryption schemes from decryption errors. In Christian Cachin and Jan Camenisch, editors, *Advances in Cryptology – EUROCRYPT 2004*, volume 3027 of *Lecture Notes in Computer Science*, pages 342–360, Interlaken, Switzerland, May 2–6, 2004. Springer Berlin Heidelberg, Germany.
- DNR04b. Cynthia Dwork, Moni Naor, and Omer Reingold. Immunizing encryption schemes from decryption errors. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 342–360. Springer, 2004.
- Döt15. Nico Döttling. Low noise LPN: KDM secure public key encryption and sample amplification. In Jonathan Katz, editor, *PKC 2015: 18th International Conference on Theory and Practice of Public Key Cryptography*, volume 9020 of *Lecture Notes in Computer Science*, pages 604–626, Gaithersburg, MD, USA, March 30 – April 1, 2015. Springer Berlin Heidelberg, Germany.
- FO99. Eiichiro Fujisaki and Tatsuki Okamoto. How to enhance the security of public-key encryption at minimum cost. In Hideki Imai and Yuliang Zheng, editors, *PKC’99: 2nd International Workshop on Theory and Practice in Public Key Cryptography*, volume 1560 of *Lecture Notes in Computer Science*, pages 53–68, Kamakura, Japan, March 1–3, 1999. Springer Berlin Heidelberg, Germany.
- FOP⁺16. Jean-Charles Faugere, Ayoub Otmani, Ludovic Perret, Frédéric De Portzamparc, and Jean-Pierre Tillich. Structural cryptanalysis of McEliece schemes with compact keys. *Designs, Codes and Cryptography*, 79(1):87–112, 2016.
- FOPT10. Jean-Charles Faugere, Ayoub Otmani, Ludovic Perret, and Jean-Pierre Tillich. Algebraic cryptanalysis of McEliece variants with compact keys. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 279–298. Springer, 2010.
- FS90. Uriel Feige and Adi Shamir. Witness indistinguishable and witness hiding protocols. In *22nd Annual ACM Symposium on Theory of Computing*, pages 416–426, Baltimore, MD, USA, May 14–16, 1990. ACM Press.
- GG24. Surendra Ghentiyala and Venkatesan Guruswami. New constructions of pseudorandom codes. *arXiv preprint arXiv:2409.07580*, 2024.
- GGW25. Sanjam Garg, Sam Gunn, and Mingyuan Wang. Black-box crypto is useless for pseudorandom codes. *Cryptology ePrint Archive*, Paper 2025/1021, 2025.
- GK03. Shafi Goldwasser and Yael Tauman Kalai. On the (in)security of the Fiat-Shamir paradigm. In *44th Annual Symposium on Foundations of Computer Science*, pages 102–115, Cambridge, MA, USA, October 11–14, 2003. IEEE Computer Society Press.
- GM84. Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270–299, 1984.
- GM24. Noah Golowich and Ankur Moitra. Edit distance robust watermarks for language models. *arXiv preprint arXiv:2406.02633*, 2024.
- GOT12. Valérie Gauthier, Ayoub Otmani, and Jean-Pierre Tillich. A distinguisher-based attack on a variant of McEliece’s cryptosystem based on Reed-Solomon codes. *arXiv preprint arXiv:1204.6459*, 2012.
- GS95. Arnaldo Garcia and Henning Stichtenoth. A tower of Artin-Schreier extensions of function fields attaining the Drinfeld-Vladut bound. *Inventiones Mathematicae*, 121(1):211–222, 1995.
- HKW20. Susan Hohenberger, Venkata Koppula, and Brent Waters. Chosen ciphertext security from injective trapdoor functions. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology – CRYPTO 2020, Part I*, volume 12170 of *Lecture Notes in Computer Science*, pages 836–866, Santa Barbara, CA, USA, August 17–21, 2020. Springer, Cham, Switzerland.
- JST21. Fernando Granha Jeronimo, Shashank Srivastava, and Madhur Tulsiani. Near-linear time decoding of Ta-Shma’s codes via splittable regularity. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1527–1536, 2021.
- JSV24. Nikola Jovanovic, Robin Staab, and Martin T. Vechev. Watermark stealing in large language models. In Ruslan Salakhutdinov, Zico Kolter, Katherine A. Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett,

- and Felix Berkenkamp, editors, *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*, volume 235 of *Proceedings of Machine Learning Research*, pages 22570–22593. PMLR / OpenReview.net, 2024.
- Kil06. Eike Kiltz. Chosen-ciphertext security from tag-based encryption. In Shai Halevi and Tal Rabin, editors, *TCC 2006: 3rd Theory of Cryptography Conference*, volume 3876 of *Lecture Notes in Computer Science*, pages 581–600, New York, NY, USA, March 4–7, 2006. Springer Berlin Heidelberg, Germany.
- KM23. Elena Kirshanova and Alexander May. Breaking Goppa-based McEliece with hints. *Information and Computation*, 293:105045, 2023.
- KM25. Fuyuki Kitagawa and Takahiro Matsuda. Adaptive TDF from any TDF via pseudorandom-ciphertext PKE. *IACR Commun. Cryptol.*, 2(3):34, 2025.
- KMP14. Eike Kiltz, Daniel Masny, and Krzysztof Pietrzak. Simple chosen-ciphertext security from low-noise LPN. In Hugo Krawczyk, editor, *PKC 2014: 17th International Conference on Theory and Practice of Public Key Cryptography*, volume 8383 of *Lecture Notes in Computer Science*, pages 1–18, Buenos Aires, Argentina, March 26–28, 2014. Springer Berlin Heidelberg, Germany.
- KMT19. Fuyuki Kitagawa, Takahiro Matsuda, and Keisuke Tanaka. CCA security and trapdoor functions via key-dependent-message security. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part III*, volume 11694 of *Lecture Notes in Computer Science*, pages 33–64, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Cham, Switzerland.
- KW19. Venkata Koppula and Brent Waters. Realizing chosen ciphertext security generically in attribute-based encryption and predicate encryption. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part II*, volume 11693 of *Lecture Notes in Computer Science*, pages 671–700, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Cham, Switzerland.
- Lam79. Leslie Lamport. Constructing Digital Signatures from a One Way Function. October 1979.
- LN09. Gaëtan Leurent and Phong Q. Nguyen. How risky is the random-oracle model? In Shai Halevi, editor, *Advances in Cryptology – CRYPTO 2009*, volume 5677 of *Lecture Notes in Computer Science*, pages 445–464, Santa Barbara, CA, USA, August 16–20, 2009. Springer Berlin Heidelberg, Germany.
- LQR⁺19. Alex Lombardi, Willy Quach, Ron D. Rothblum, Daniel Wichs, and David J. Wu. New constructions of reusable designated-verifier NIZKs. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part III*, volume 11694 of *Lecture Notes in Computer Science*, pages 670–700, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Cham, Switzerland.
- LS02. Pierre Loidreau and Nicolas Sendrier. Weak keys in the McEliece public-key cryptosystem. *IEEE Transactions on Information Theory*, 47(3):1207–1211, 2002.
- McE78. Robert J McEliece. A public-key cryptosystem based on algebraic. *Coding Thv*, 4244(1978):114–116, 1978.
- Mce26. Classic McEliece. <https://classic.mceliece.org/papers.html>, 2026. [Accessed 12-02-2026].
- MP12. Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology – EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 700–718, Cambridge, UK, April 15–19, 2012. Springer Berlin Heidelberg, Germany.
- Nao90. Moni Naor. Bit commitment using pseudo-randomness. In Gilles Brassard, editor, *Advances in Cryptology – CRYPTO’89*, volume 435 of *Lecture Notes in Computer Science*, pages 128–136, Santa Barbara, CA, USA, August 20–24, 1990. Springer, New York, USA.
- Nie86. Harald Niederreiter. Knapsack-type cryptosystems and algebraic coding theory. *Prob. Contr. Inform. Theory*, 15(2):157–166, 1986.
- Nie02. Jesper Buus Nielsen. Separating random oracle proofs from complexity theoretic proofs: The non-committing encryption case. In Moti Yung, editor, *Advances in Cryptology – CRYPTO 2002*, volume 2442 of *Lecture Notes in Computer Science*, pages 111–126, Santa Barbara, CA, USA, August 18–22, 2002. Springer Berlin Heidelberg, Germany.
- NY90. Moni Naor and Moti Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *22nd Annual ACM Symposium on Theory of Computing*, pages 427–437, Baltimore, MD, USA, May 14–16, 1990. ACM Press.
- PHZS24. Qi Pang, Shengyuan Hu, Wenting Zheng, and Virginia Smith. No free lunch in llm watermarking: Trade-offs in watermarking design choices, 2024.
- PW08. Chris Peikert and Brent Waters. Lossy trapdoor functions and their applications. In Richard E. Ladner and Cynthia Dwork, editors, *40th Annual ACM Symposium on Theory of Computing*, pages 187–196, Victoria, BC, Canada, May 17–20, 2008. ACM Press.

- QRW19. Willy Quach, Ron D. Rothblum, and Daniel Wichs. Reusable designated-verifier NIZKs for all NP from CDH. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2019, Part II*, volume 11477 of *Lecture Notes in Computer Science*, pages 593–621, Darmstadt, Germany, May 19–23, 2019. Springer, Cham, Switzerland.
- Ran25. Hugues Randriambololona. The syzygy distinguisher. In Serge Fehr and Pierre-Alain Fouque, editors, *Advances in Cryptology - EUROCRYPT 2025 - 44th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Madrid, Spain, May 4-8, 2025, Proceedings, Part VI*, volume 15606 of *Lecture Notes in Computer Science*, pages 324–354. Springer, 2025.
- Sah99. Amit Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In *40th Annual Symposium on Foundations of Computer Science*, pages 543–553, New York, NY, USA, October 17–19, 1999. IEEE Computer Society Press.
- SKS19. Jad Silbak, Swastik Kopparty, and Ronen Shaltiel. Quasilinear time list-decodable codes for space bounded channels. In David Zuckerman, editor, *60th Annual Symposium on Foundations of Computer Science*, pages 302–333, Baltimore, MD, USA, November 9–12, 2019. IEEE Computer Society Press.
- SS92. Vladimir Michilovich Sidelnikov and Sergey O Shestakov. On insecurity of cryptosystems based on generalized Reed-Solomon codes. 1992.
- Sti09. Henning Stichtenoth. *Algebraic Function Fields and Codes*, volume 254 of *Graduate Texts in Mathematics*. Springer, 2 edition, 2009.
- vT94. Johan van Tilburg. Security-analysis of a class of cryptosystems based on linear error-correcting codes. 1994.
- Xag22. Keita Xagawa. The boneh-katz transformation, revisited: Pseudorandom/obliviously-samplable PKE from lattices and codes and its application. In Riham AlTawy and Andreas Hülsing, editors, *SAC 2021: 28th Annual International Workshop on Selected Areas in Cryptography*, volume 13203 of *Lecture Notes in Computer Science*, pages 47–67, Virtual Event, September 29 – October 1, 2022. Springer, Cham, Switzerland.
- Zha22. Mark Zhandry. Augmented random oracles. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022, Part III*, volume 13509 of *Lecture Notes in Computer Science*, pages 35–65, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Cham, Switzerland.