Rate-1 Incompressible Encryption from Standard Assumptions

Pedro Branco¹, Nico Döttling²^[0000-0002-5914-7635], and Jesko Dujmović^{2,3}^[0000-0003-1043-5094]

 ¹ Johns Hopkins University, Baltimore, MD 21218, USA pedrodemelobranco@gmail.com
 ² Helmholtz Center for Information Security (CISPA), 66123 Saarbrücken, Germany {doettling,jesko.dujmovic}@cispa.de
 ³ Saarland University, 66123 Saarbrücken, Germany

Abstract. Incompressible encryption, recently proposed by Guan, Wichs and Zhandry (EURO-CRYPT'22), is a novel encryption paradigm geared towards providing strong long-term security guarantees against adversaries with *bounded long-term memory*. Given that the adversary forgets just a small fraction of a ciphertext, this notion provides strong security for the message encrypted therein, even if, at some point in the future, the entire secret key is exposed. This comes at the price of having potentially very large ciphertexts. Thus, an important efficiency measure for incompressible encryption is the message-to-ciphertext ratio (also called the rate). Guan et al. provided a low-rate instantiation of this notion from standard assumptions and a rate-1 instantiation from indistinguishability obfuscation (iO). In this work, we propose a simple framework to build rate-1 incompressible encryption from standard assumptions. Our construction can be realized from, e.g. the DDH and additionally the DCR or the LWE assumptions.

1 Introduction

Incompressible Cryptography [26,19,27,42,34,32] is a flourishing paradigm trying to leverage memory limitations of adversaries to achieve strong security goals. While traditionally, the goal of cryptography in the bounded storage model [40] is to minimize the need for computational assumptions or even obtain information-theoretically secure constructions, incompressible cryptography is geared more toward mitigating the consequences of key exfiltration and key exposure attacks. In this work, we focus on the notion of *incompressible encryption* [26,32] ⁴ recently coined by Guan et al. [32]. An incompressible encryption scheme produces large, incompressible ciphertexts and guarantees that any adversary who forgets even a small fraction of the ciphertext data will learn nothing about the encrypted data, even if he is later given the corresponding secret key!

One motivation for incompressible encryption is to hamper adversaries conducting a mass-surveillance operation by forcing them to store massive amounts of ciphertext data even if they are just interested in a tiny fraction of the encrypted data. In a similar scenario, an adversary trying to exfiltrate information encrypted under an incompressible encryption scheme from a data-center will have to exfiltrate massive amounts of data, even if his exfiltration target is just a small piece of information.

An orthogonal notion to incompressible encryption is encryption in the *bounded-retrieval model* [25,22,4,3,36,9,8,42] where the goal is to make the secret key large and incompressible (to make it hard to exfiltrate) while keeping all other system parameters small, such as the sizes of public keys and ciphertexts, as well as the overhead of encryption and decryption.

Encryption with High Rate. An important efficiency measure of encryption schemes is their ciphertext expansion or rate. The rate of an encryption scheme is the ratio between plaintext size and ciphertext size. The closer the rate is to 1, the more efficient a scheme manages to pack information into a ciphertext. Conversely, the closer the rate is to 0, the less information is encoded in potentially large ciphertexts. For incompressible encryption, achieving a high rate (ideally converging to 1), especially if we think of the data center application above, where a small rate would also put a massive burden on the data center.

Guan et al. [32] provided two constructions of incompressible encryption.

- A construction from the minimal assumption of public-key encryption which has ciphertext-rate approaching 0.

 $^{^4}$ Dziembowski [26] introduced this concept under the name forward-secure storage in the symmetric key setting.

- A construction from indistinguishability obfuscation (iO) [6,28,37] which achieves ciphertext-rate approaching 1.

We remark that their rate-1 construction relies on non-black-box techniques and iO, which gives this result a strong feasibility flavor.

Given this state of affairs, this work is motivated by the following question:

Can we build a rate-1 incompressible encryption scheme based on standard assumptions while only making black-box use of cryptographic primitives?

1.1 Our Results

In this work, we build a rate-1 incompressible encryption scheme from standard assumptions while only using black-box techniques. Our result uses what we call programmable hash proof systems (HPS) (which are a variant of standard HPS [16,17] with some additional properties), plain-model incompressible encodings [42] and a pseudorandom generator (PRG). In particular, we prove the following theorem.

Theorem 1 (Informal). Let S be the storage capacity of the adversary and let n be the size of the encrypted messages. Assuming programmable HPS, incompressible encodings and PRGs exist, there is an incompressible encryption scheme fulfilling the following properties:

- 1. Ciphertexts are of size $n + n^{\varepsilon} \cdot \operatorname{poly}(\lambda)$ for some $\varepsilon > 0$.
- 2. The public key is of size $n^{\varepsilon'} \cdot \operatorname{poly}(\lambda)$ for some $\varepsilon' > 1/2$.
- 3. Moreover, the size of ciphertexts is only slightly larger than the adversary's storage space, that is, $S + poly(\lambda)$.

The ciphertext rate $n/(n + n^{\varepsilon} \cdot \text{poly}(\lambda))$ approaches 1 for large enough messages. Additionally, the public key is sublinear in the size of the encrypted message.

In terms of assumptions, incompressible encodings can be based on either decisional composite residuosity (DCR) or learning with errors (LWE). The PRG can be based on any one-way function. We also show that programmable HPS can be instantiated from the decisional Diffie-Hellman (DDH) assumption by tweaking the famous HPS by Cramer and Shoup [17]. Consequently, our final incompressible encryption scheme can be based solely on standard assumptions.

Streaming encryption. Streaming encryption/decryption is a property of incompressible encryption schemes which allows the honest encryptor/decryptor to perform operations with very low storage capacity. It is easy to see that streaming decryption is an inherently conflicting property with high rate ciphertexts [32]. This is because the honest decryptor needs storage at least as large as the size of the message. Otherwise, an adversary can essentially mimic the decryptor and learn something about the encrypted message (e.g., the most significant bit).

However, we note that our scheme has *stream encryption*, i.e., the honest encryptor does not need much space to perform encryption. This follows from the fact that the incompressible encodings construction of [42] has stream encoding.

Extension to CCA security. In the security experiment for incompressible encryption presented in [32] the adversary is never allowed to query a decryption oracle. In other words, their work only considered IND-CPA incompressible encryption. In this work, we also give the adversary access to an decryption oracle extending incompressible encryption to IND-CCA2 incompressible encryption. We stress that IND-CCA2 security is usually considered the *right* security definition to use in practice. We show that our construction is, in fact, is IND-CCA2 incompressible secure.

1.2 Comparison with Previous Work

[32] presented two incompressible encryption schemes. The first one is based only on the minimal assumption of PKE. However, the ciphertext rate is very far from 1. The second one achieves rate-1 but is based on $i\mathcal{O}$. We compare these schemes in Table 1.

	Ciphertext Rate	Public key Size	Hardness Assumption	Security
[32]	$1/poly(\lambda)$	$poly(\lambda)$	PKE	IND-CPA
[32]	1	$poly(\lambda)$	iO	IND-CPA
Our result	1	$n^{\varepsilon'} \cdot poly(\lambda)$	$\begin{array}{c} \text{DDH} \land \\ (\text{DCR} \lor \text{LWE}) \end{array}$	IND-CCA2

Table 1. Comparison with previous work. Here, n denotes the size of the encrypted messages and ε' is any constant between 1/2 and 1.

Other related work. Some recent works made significant progress in the area of incompressible cryptography. The works of [19,27,42] proposed constructions for incompressible encodings either in the random oracle model or in the CRS model. The work of [34] used the BSM together with computational assumptions to propose constructions of primitives that are not known just from computational assumptions, such as virtual grey-box obfuscation.

Incompressible cryptography is closely related to the bounded storage model (BSM) [39]. However, most works in the BSM (e.g. [13,5,45,33,24]) focus on achieving unconditional security for primitives that are already known from computational assumptions such as public-key encryption and oblivious transfer.

Open Problems. We leave the open problem of developing an incompressible encryption scheme that combines concretely short public keys with small ciphertexts. A possible approach for this would be to find a programmable hash proof system where the size of the public key is essentially independent of the size of the encapsulated key.

Full Version In the full version [11], we justify focusing on the plain model by providing a simple incompressible encryption scheme that is secure in the random oracle model but is broken for any instantiation of the random oracle. This provides another uninstantiability for the random oracle in the vein of [14,21,29,7,41,10,12,31].

The full version also contains a programmable HPS based on isogeny-based assumptions with worse parameters than the DDH programmable HPS and a programmable HPS based on the hardness of LWE with superpolynomial modulus-to-noise ratio with better parameters than the DDH programmable HPS. The LWE construction, however, only results in incompressible IND-CPA security.

Acknowledgement. We would like to thank Stefan Dziembowski, Daniel Wichs, and the anonymous reviewers of TCC for discussions and comments.

Nico Döttling is funded by the European Union. 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 Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. (ERC-2021-STG 101041207 LACONIC).

Part of the work of Pedro Branco was done while at IST University of Lisbon.

2 Technical Overview

In this technical overview, we sketch the main techniques to build an IND-CPA incompressible scheme. We later argue how these techniques can be tweaked to obtain a scheme that is IND-CCA2 incompressible secure.

Security Notion The syntax and correctness notions for incompressible encryption are identical to standard public-key encryption (PKE). The main difference is in the security definition. Since the security notion of incompressible encryption is relatively new, we will briefly detail its security experiment here. Consider the following security game between a challenger C and a 3-stage PPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$.

1. ${\mathcal C}$ creates a pair of public and secret keys $\mathsf{pk},\mathsf{sk}.$

2. Given pk, the first stage \mathcal{A}_1 chooses two messages m_0, m_1 .

- 3. C chooses $b \leftarrow \{0, 1\}$ uniformly at random and encrypts $\mathsf{ct} \leftarrow \mathsf{Enc}(\mathsf{pk}, \mathsf{m}_b)$.
- 4. Given the ciphertext ct and the state of A_1 , the second stage A_2 produces a state st of size S < |ct|. That is, the state st should be somewhat smaller than ct.
- 5. Now, the third stage \mathcal{A}_3 receives as input the state st (produced by \mathcal{A}_2) and the secret key sk. The goal of \mathcal{A}_3 is to guess the bit b.

We say that an incompressible encryption is secure if, for any adversary, \mathcal{A} the advantage of winning the following game is negligible in the security parameter λ .

2.1 The Scheme of GWZ

Before we provide an outline of our construction, we will briefly discuss the underlying ideas of the lowrate incompressible encryption scheme constructed in [32]. At the very core is the following idea: The ciphertext essentially consists of a very long truly random random string R and a short payload part $c = (c_1, c_2)$, where c_1 is an encryption of a seed k for a randomness extractor Ext, and $c_2 = \text{Ext}(k, R) \oplus \text{m}$ is essentially a one-time-pad encryption of the message m under the key Ext(k, R). Clearly, if c_1 was not part of the ciphertext, then security of this scheme follows routinely by the following observations:

- In the view of the third stage A_3 of the adversary R has high min-entropy, as R is uniformly random and the state st is significantly shorter than R.
- Furthermore, as we assume c_1 is not part of the ciphertext, st is independent of k
- Hence by the extraction property of Ext the string Ext(k, R) is uniformly random in the adversary's view, and therefore m_b is statistically hidden.

Now, the main idea of [32] to make this approach work even though c_1 is part of the ciphertext is to encrypt k in such a way that c_1 can be made independent of the extractor seed k. This is achieved by choosing a suitable encryption scheme for which c_1 can be chosen independently of k, and a suitable secret key which decrypts c_1 to k can be chosen after the fact, i.e. after the leakage st has been computed. [32] provide an elegant construction of such a scheme from non-compact single-key functional encryption, which can be built from any public key encryption scheme [30].

2.2 The Big Picture

While our construction departs significantly from the blueprint of [32] we use the same high-level concept of an encryption scheme that allows delaying secret-key generation in the security proof. Rather than constructing incompressible PKE directly, we first tackle the intermediate and simpler task of realizing a rate-1 *incompressible symmetric-key encryption*. In a second step, we will then transform *any* incompressible SKE scheme into an incompressible PKE scheme in a rate-preserving way. It turns out that even constructing a rate-1 incompressible SKE from standard assumptions is a non-trivial task and does not follow, e.g. from the (low-rate) public-key construction of [32].

Since our two steps are independent of one another, improvements of either in future work will lead to better incompressible encryption schemes. For simplicity, in the following outline, we will focus only on CPA security, whereas in the main body, we present a CCA secure construction.



Fig. 1. Overview of the results in this work, bold arrows are contributions of this work.

2.3 Rate-1 Incompressible Symmetric-Key Encryption

In the symmetric-key setting, the syntax and correctness of incompressible SKE are pretty much that of standard symmetric-key encryption, whereas the security notion is similar to that of incompressible PKE, just with the difference that the first stage A_1 of the adversary is not given a public key (as there is none). Thus, the security notion we consider here is the incompressible encryption-analogue of security against an eavesdropper (IND-EAV).

Failed Naive Attempts. As a (failed) very first attempt, one may try "make work" an incompressible SKE construction from the One-Time-Pad (OTP), i.e. the secret key k is a random bit-string as long as the message m and the ciphertext is $c = k \oplus m$. However, the obvious issue with this is that such a ciphertext c decomposes into many one-bit ciphertexts $c_i = k_i \oplus m_i$, and it is enough for A_2 to leak a few bits of c to enable A_3 to partially decrypt c and thus distinguish encryptions of m_0 from encryptions of m_1 . As a next idea, one may try the following: Encryption chooses a (fresh) pseudorandom generator (PRG) seed s, encrypt m into $\hat{m} = m \oplus PRG(s)$, use k to encrypt the seed s into a header ciphertext c, i.e. the encryption of m is (Enc(k, s), $m \oplus PRG(s)$). While this approach does look promising, we observe that it is stuck at either leakage-rate 1/2 or ciphertext-rate 1/2, that is as soon as A_3 learns Enc(k, s) in its entirety and a few bits of $m_b \oplus PRG(s)$, he will be able to distinguish encryptions of m_0 from m_1 .

Introducing Circularity. Clearly, we need some kind of mechanism to glue the two ciphertexts components together, i.e. we want to make it such that if some parts of \hat{m} are missing, then c will be useless (and vice versa). As a first, heuristic "hands-on" approach to achieve this, we can try to use \hat{m} as a source of randomness from which we extract a key to mask the seed s. Thus, let $\mathsf{Ext}(\cdot, \cdot)$ be a *seeded* randomness extractor. We compute a ciphertext (c, \hat{m}) by first computing $\hat{m} = m \oplus \mathsf{PRG}(s)$ for a random seed s as before, but then encrypt s into c via $c = s \oplus \text{Ext}(k, \hat{m})$, i.e. we use k as an extractor seed to extract a onetime-pad key $Ext(k, \hat{m})$ from \hat{m} . Clearly, given k and a ciphertext (c, \hat{m}) , we can decrypt by first computing $s = c \oplus \text{Ext}(k, \hat{m})$ and then $m = \hat{m} \oplus \text{PRG}(s)$. The rationale for why we hope this construction to be secure is that as soon as a significant fraction of the bits of \hat{m} are lost, the output of the extractor $Ext(k, \hat{m})$ should look uniform, and thus $\hat{m} = m \oplus PRG(s)$ should hide m by the pseudorandomness of PRG. However, this circularity backfires when trying to establish security of this construction just from the pseudorandomness of PRG and the randomness-extraction property of Ext: In order to use pseudorandomness of PRG, we first need to *remove* the s from the view of the adversary, but $c = s \oplus \text{Ext}(k, \hat{m})$ is correlated with s given k. On the other hand, in order to use the randomness extraction property of Ext we need that \hat{m} has high entropy given st. But all the entropy of $\hat{m}=m\oplus\mathsf{PRG}(s)$ comes from the seed s, which is very small. Hence $\approx \lambda$ bits of \hat{m} suffice to information-theoretically determine s.

Consequently, while heuristically, this construction seems secure, it seems unlikely the individual security properties of PRG and Ext suffice to *prove* this construction secure.

Breaking Circularity Hence, what we need is a mechanism to break the circularity, which we have just introduced. Looking at where establishing security of the above construction gets stuck, a natural point to start is to make it such that \hat{m} looks like it has a large amount of real entropy once a few bits of \hat{m} are missing, i.e. $L(\hat{m})$ being computationally indistinguishable from $L(\hat{r})$ for a high-entropy distribution \hat{r} for any efficiently computable leakage function $L(\cdot)^{-5}$.

Incompressible encodings. Fortunately, an encoding mechanism achieving this notion called *incompress-ible encodings* was just recently introduced and constructed by Moran and Wichs [42]. As a technical tool, they introduced the notion of *HILL-entropic encoding* in their work, which will be sufficient, if not to say ideally suited for our construction. Such a scheme consists of an encoding algorithm En and a decoding algorithm De, both of which rely on a (large) common random string $crs \leftarrow \{0, 1\}^t$:

- The encoding algorithm $\mathsf{En}_{\mathsf{crs}}(m)$ is a randomized algorithm which takes a message m and produces an encoding \hat{m}
- The decoding algorithm $\mathsf{De}_{\mathsf{crs}}(\hat{m})$ is a deterministic algorithm which takes an encoding \hat{m} and returns a message m.

⁵ In our case the leakage function L is described by the adversary's second stage \mathcal{A}_2

In terms of correctness, one naturally requires that encoding followed by decoding leads to the original message. Security-wise, we require that there exists a simulator Sim which on input a message m produces a pair (crs', \tilde{m}), which is computationally indistinguishable from a real pair of crs and encoding of m, i.e.

$$(\operatorname{crs}, \operatorname{En}_{\operatorname{crs}}(m)) \approx_c \operatorname{Sim}(m),$$

where $\operatorname{crs} \leftarrow \{0,1\}^t$. Additionally, we require that if $(\operatorname{crs}', \tilde{m}) \leftarrow \operatorname{Sim}(m)$, then \tilde{m} has almost full *true* min-entropy given crs' , i.e. $\tilde{H}_{\infty}(\tilde{m}|\operatorname{crs}') \geq (1-\epsilon)n$, where \tilde{H}_{∞} is the average conditional min-entropy. The (very) high level idea how this can be achieved is that in simulation the common random string and the encoded message *switch roles*, in the sense that the simulated common random string crs' encodes the message m, whereas the encoding \tilde{m} now has room to have (very) high entropy.

Moran and Wichs [42] provide two instantiations of their construction, one from DCR and one from LWE. These constructions achieve rate-1, i.e., the encoding is only slightly larger than the encoded message. The conceptual idea behind the construction is rather elegant: The encoding function $En_{crs}(m)$ generates a pair of public key and trapdoor (pk, td) of preimage-sampleable surjective lossy function F (for which we have efficient constructions from DCR or LWE) and sets x to be a randomly sampled preimage of $m \oplus crs$, i.e. $x = F_{td}^{-1}(m \oplus crs)$, and sets $\hat{m} = (pk, x)$. To decode \hat{m} , one computes $m = F_{pk}(x) \oplus crs$. The simulator Sim chooses a highly lossy public key $p\tilde{k}$, chooses x uniformly at random, and sets $crs' = m \oplus F_{p\tilde{k}}(x)$ and $\tilde{m} = (p\tilde{k}, x)$. Given that F_{pk} is regular for surjective keys pk, meaning that if x is uniform then $F_{pk}(x)$ is also (statisticalluy close to) uniform, we can routinely establish that real pairs (crs, \hat{m}) are computationally indistinguishable from simulated (crs', \tilde{m}) using the indistinguishability of surjective public keys pk and highly lossy public keys $p\tilde{k}$. Moreover, for simulated pairs ($crs' = m \oplus F_{p\tilde{k}}(x)$, $\tilde{m} = (\tilde{pk}, x)$) we can easily argue that x (and hence \tilde{m}) has high min-entropy given $crs' = m \oplus F_{p\tilde{k}}(x)$, as $F_{p\tilde{k}}$ is highly lossy and hence x has high min entropy given $F_{p\tilde{k}}(x)$.

Moran and Wichs [42] go on to show that for any incompressible encoding/HILL-entropic encoding, the common random string **crs** must be as long as the message, if one wants to establish security from a *falsifiable assumption* [43] under a black-box reduction.

The Full Construction. We will now provide our complete construction of incompressible SKE and sketch the security proof. For our scheme, the secret key K is a uniformly random bit-string of suitable length which will be parsed as K = (crs, k), where crs is the common random string for a HILL-entropic encoding (En, De), and k is the seed for a randomness extractor Ext. Encryption and decryption work as follows.

- $\mathsf{Enc}(\mathsf{K} = (\mathsf{crs}, \mathsf{k}), \mathsf{m})$: Choose a uniformly random PRG seed $\mathsf{s} \leftarrow \{0, 1\}^{\lambda}$ and compute $\hat{\mathsf{m}} = \mathsf{En}_{\mathsf{crs}}(m \oplus \mathsf{PRG}(\mathsf{s}))$. Compute $c = \mathsf{s} \oplus \mathsf{Ext}(\mathsf{k}, \hat{\mathsf{m}})$ and output the ciphertext $\mathsf{ct} = (c, \hat{\mathsf{m}})$.
- $\mathsf{Dec}(\mathsf{K} = (\mathsf{crs}, \mathsf{k}), \mathsf{ct} = (c, \hat{\mathsf{m}}))$: Compute $\mathsf{s} = c \oplus \mathsf{Ext}(\mathsf{k}, \hat{\mathsf{m}})$ and output $m = \mathsf{De}_{\mathsf{crs}}(\hat{\mathsf{m}}) \oplus \mathsf{PRG}(\mathsf{s})$.

Correctness of this scheme follows routinely.

Security of this scheme is established along the following lines. First we rely on the security of the HILL-entropic encoding to replace (crs, \hat{m}) with a *simulated pair* $(crs', \tilde{m}) = Sim(m \oplus PRG(s))$. By the security of the HILL-entropic encoding, this modification is (computationally) unnoticeable to the adversary. However, now the encoding \tilde{m} has true high min-entropy given crs'. Thus, using a min-entropy chain rule (e.g. by [23]) we can argue that \tilde{m} still has sufficiently high min-entropy given both crs' and a leak $L(\tilde{m})$. Hence, the randomness extraction property guarantees that $Ext(k, \tilde{m})$ will extract uniform randomness (given crs' and $L(\tilde{m})$). To establish this we need a mild extra property of the extractor Ext that given a (uniformly random) extractor output y and \tilde{m} we can sample a key k' after the fact such that $(k', y) \approx (k, Ext(k, \tilde{m}))$. Hence in the next hybrid modification, we can thus replace $c = s \oplus Ext(k, \tilde{m})$ with a uniformly random and independent string c'. Now that c' is independent of s, we can use the pseudorandomness property of PRG to replace $m \oplus PRG(s)$ in $(crs', \tilde{m}) = Sim(m \oplus PRG(s))$ with a uniformly random string u, i.e. $(crs', \tilde{m}) = Sim(u)$. We have finally arrived at an experiment where the ciphertext $ct = (c', \tilde{m})$ is independent of the message m, and hence the adversary's advantage is 0.

Concerning the rate of this scheme, note that a ciphertext $ct = (c, \hat{m})$ has rate 1, as c is just of size $poly(\lambda)$ (independent of the message length n), and the HILL-entropic encoding \hat{m} is rate 1.

2.4 From Symmetric-Key to Public-Key Incompressible Encryption via Hash Proof Systems

Now that we have a construction of incompressible SKE, we need a way to establish a *long* key K between the sender and receiver. This is a job for a *key encapsulation mechanism* (KEM) [18]. A key-encapsulation mechanism consists of:

- A key-encapsulation mechanism consists of a key-generation algorithm KeyGen which produces a pair of public and secret keys (pk, sk).
- An encapsulation algorithm which takes a public key pk and produces a symmetric key K and a ciphertext header c_0 encapsulating K.
- A decapsulation algorithm Dec which takes a secret key sk and a ciphertext header c_0 and outputs a key K.

The correctness requirement is the obvious one, whereas the standard security requirement is that K is pseudorandom given pk and c_0 . A symmetric key K generated via a KEM can now be used to encrypt a message m into a payload ciphertext c_1 using a symmetric key encryption scheme. The full ciphertext is $c = (c_0, c_1)$.

However, to transform an incompressible SKE into an incompressible PKE not just any key encapsulation mechanism will do. The simple reason is that in the incompressible (public key) encryption security game, the adversary gets to see the secret key sk in the end, which will allow him to decapsulate the (short) ciphertext header c_0 into the symmetric key K. But the standard security notion of KEMs discussed above does not require that the encapsulated key K follows a uniform distribution. Indeed, e.g. for simple PRG-based KEMs, the encapsulated key is statistically far from uniform. However, recall that in our construction of incompressible SKE above, we made critical use of the fact that the key K follows a uniform distribution and that the security reduction can program it in a suitable way.

Thus, we need a KEM which we can switch into a mode in which the ciphertext header c_0 encapsulates a truly uniform key K. As we need the ciphertext header c_0 to be substantially shorter than the encapsulated key K, the entropy of K in this mode *must* come from the secret key sk.

Enter Hash proof systems. This is where hash proof systems (HPS) [17] come into play ⁶. Recall that HPS are defined relative to an NP-language $\mathcal{L} \subseteq \{0,1\}^k$. We have a key-generation algorithm KeyGen which generates a public or projected key pk, and a secret or hashing key sk. The hashing or decapsulation algorithm Decap takes the secret key sk and any $x \in \{0,1\}^k$ and produces a hash value K. The restricted hashing or encapsulation algorithm Encap takes a public key pk, an $x \in \mathcal{L}$ and a witness w (with respect to a fixed NP-relation for \mathcal{L}) for membership of x in \mathcal{L} and produces a hash-value K.

In terms of correctness or completeness, we require that Decap and Encap agree on \mathcal{L} , i.e. if $x \in \mathcal{L}$ and w is a valid witness for x, then it holds that Decap(sk, x) = Encap(pk, x, w).

In terms of security, we require *smoothness*, namely given that $x \notin \mathcal{L}$, it holds that $\mathsf{Decap}(\mathsf{sk}, x)$ is statistically close to uniform given pk .

HPS are especially useful for sparse pseudorandom languages \mathcal{L} , such as the decisional Diffie-Hellman (DDH) language) [17]. We define this language with respect to a pair of (randomly chosen) generators $g, h \in \mathbb{G}$, where \mathbb{G} is a cryptographic group of prime order p. A pair x = (g', h') is in \mathcal{L} , if there exists an $r \in \mathbb{Z}_p$ such that $g' = g^r$ and $h' = h^r$. The DDH assumption states that a random element in \mathcal{L} , i.e. a pair (g^r, h^r) is computationally indistinguishable from a pair of uniformly random group elements $(u, v)^7$

In the Cramer-Shoup [17] scheme, the secret key $\mathsf{sk} = (\alpha, \beta)$ consists of two uniformly random values $\alpha, \beta \in \mathbb{Z}_p$, and the public key pk is computed as $\mathsf{pk} = g^{\alpha}h^{\beta}$. Given a public key pk an instance $c_0 = (g^r, h^r)$ with witness r, we compute a key $\mathsf{K} = \mathsf{pk}^r$. Given a secret key $\mathsf{sk} = (\alpha, \beta)$ and an instance $c_0 = (g', h')$ we compute a key $\mathsf{K} = g'^{\alpha}h'^{\beta}$. It follows routinely that encapsulation and decapsulation agree on \mathcal{L} . Moreover, for a $(g^*, h^*) \notin \mathcal{L}$ it holds $\mathsf{K}^* = g^{*\alpha}h^{*\beta}$ is uniformly random given $\mathsf{pk} = g^{\alpha}h^{\alpha}$ by a simple linear algebra argument.

Hash Proof Systems, and in particular the Cramer-Shoup HPS (almost) give us a KEM with the desired properties. Namely, given pk and (g, h), to encapsulate a key k we choose a uniformly random

⁶ HPS have been instrumental in many prior works on leakage resilience cryptography e.g. [3,36]

⁷ Note that such a pair is not in \mathcal{L} , except with negligible probability 1/p.

 $r \in \mathbb{Z}_p$ and compute $c_0 = (g^r, h^r)$ and $\mathsf{K} = \mathsf{pk}^r$. To decapsulate K from $c_0 = (g', h')$ given $\mathsf{sk} = (\alpha, \beta)$, we compute $\mathsf{K} = g'^{\alpha} h'^{\beta}$.

A typical proof-strategy using HPS lets a reduction compute the encapsulated key (on the sender's side) via the decapsulation algorithm using the secret key. Correctness of the HPS ensures that this does not change K. Hence this modification will not be detected by an adversary. Now we don't need the witness r anymore. We can replace c_0 with a uniformly random c'_0 and argue that this modification is computationally undetectable by the adversary, thanks to the DDH assumption. Since now c'_0 is outside of \mathcal{L} w.o.p, it holds that K is uniform even given pk, as desired.

However, this is still not enough to make our security reduction go through. It turns out we not only have to ensure that K is uniform given pk, but also that for any given K and fixed pk and c_0 we can find a secret key sk (compatible with pk) such that $Decap(sk, c_0) = K$. Realizing this property using the Cramer-Shoup HPS directly seems hard, as in order to sample an $sk = (\alpha, \beta)$ with $K = g'^{\alpha} h'^{\beta}$ we would need to compute a discrete logarithm of K.

Programmable Hash Proof Systems For this purpose, we will consider a notion of programmable hash proof systems, which obey a stronger smoothness notion. In short, such an HPS has the following property. Given a public key pk, a (fake) ciphertext header c_0^* (not in \mathcal{L}) and secret auxiliary information aux depending on both pk and c_0 , we can sample a uniformly random secret key sk^{*} such that $\mathsf{Decap}(\mathsf{sk}^*, c_0^*) = \mathsf{K}$, for which it holds that $(\mathsf{pk}, c_0^*, \mathsf{sk}^*) \approx_s (\mathsf{pk}, c_0^*, \mathsf{sk})$ if K is chosen uniformly random.

Our idea to achieve this is simple: We will concatenate Decap (and also Encap) with a balanced small range hash function $HC : \mathbb{G} \to \{0,1\}$, i.e. we have $Decap'(sk, c_0) = HC(Decap(sk, c_0))$ and $Encap'(pk, c_0, r) = HC(Encap(pk, c_0, r))$. Here balanced means that if $h \in \mathbb{G}$ is a uniformly random group element, then HC(h) is statistically close to a uniformly random bit. While there exist deterministic constructions of such extractors for certain groups (e.g. [15]) we can find such an HC for any group via the leftover-hash lemma [35]. For such a hash function, we can efficiently sample a uniformly random pre-image $h \in \mathbb{G}$ of K for which we do know the discrete logarithm (with respect to a generator $g \in \mathbb{G}$). We achieve this via rejection sampling: Given a bit $K \in \{0, 1\}$, choose a uniformly random $z \in \mathbb{Z}_p$ and test whether $HC(g^z) = K$ (which happens with probability 1/2), and reject and resample if the test fails.

Now let $h = g^y$, $\mathsf{pk} = g^t$ and $c_0^* = (g' = g^r, h' = g^s)$ be a public key and (fake) ciphertext, for which the auxiliary information is (y, t, r, s), i.e. the discrete logarithms of pk and c_0^* . Given a key $\mathsf{K} \in \{0, 1\}$, we first sample a uniformly random $z \in \mathbb{Z}_p$ such that $\mathsf{HC}(g^z) = \mathsf{K}$. Now we have 2 linear constraints (over \mathbb{Z}_p) on $\mathsf{sk} = (\alpha, \beta) \in \mathbb{Z}_p^2$, namely

$$t = \alpha + \beta \cdot y$$

from $\mathsf{pk} = q^{\alpha} \cdot h^{\beta}$ and

$$z = \alpha r + \beta s$$

from $\mathsf{HC}(g^z) = \mathsf{HC}(g'^{\alpha} \cdot h'^{\beta})$. Since we now have two equations and two unknowns α and β , we can solve for α and β using basic linear algebra.

We do pay a price to get programmability: Instead of getting $\log(|\mathbb{G}|)$ key bits per public key pk, we only get a single bit. Naturally, this can be improved up to $\log(\lambda)$ key-bits while keeping the above rejection sampling procedure expected polynomial time.

The Full Construction We are now ready to present our fully-fledged construction. This construction will have a large public key. We will later discuss how the size of the public key can be reduced.

Assume thus that (Enc, Dec) is an incompressible SKE scheme, and that (KeyGen, Encap, Decap) is a programmable HPS for a decision-membership-hard language \mathcal{L} , for concreteness assume the DDH language. Our incompressible PKE construction is given by the following algorithms.

- The key-generation algorithm KeyGen' generates random group elements $g, h \in \mathbb{G}$ and n pairs of public and secret keys $(\mathsf{pk}_1, \mathsf{sk}_1), \ldots, (\mathsf{pk}_n, \mathsf{sk}_n)$ using KeyGen (on g, h) and set $\mathsf{PK} = (g, h, \mathsf{pk}_1, \ldots, \mathsf{pk}_n)$ and $\mathsf{SK} = (\mathsf{sk}_1, \ldots, \mathsf{sk}_n)$.
- The encryption algorithm $\mathsf{Enc'}$ proceeds as follows, given a public key $\mathsf{PK} = (g, h, \mathsf{pk}_1, \dots, \mathsf{pk}_n)$ and a message m. First, generate a random DDH instance $c_0 = (g' = g^r, h' = h^r)$ using a random $r \leftarrow \mathbb{Z}_p$. Now compute the key-bits $\mathsf{K}_1 = \mathsf{Encap}(\mathsf{pk}_1, c_0, r), \dots, \mathsf{K}_n = \mathsf{Encap}(\mathsf{pk}_n, c_0, r)$ and set $\mathsf{K} = (\mathsf{K}_1, \dots, \mathsf{K}_n)$. Next, we use K to encrypt m using the incompressible SKE scheme, i.e. we compute $c_1 = \mathsf{Enc}(\mathsf{K}, \mathsf{m})$ and output the ciphertext $c = (c_0, c_1)$.

- The decryption algorithm $\mathsf{Dec'}$ takes a secret key $\mathsf{SK} = (\mathsf{sk}_1, \dots, \mathsf{sk}_n)$ and a ciphertext $c = (c_0, c_1)$, and proceeds as follows. First, it decapsulates the key $\mathsf{K} = (\mathsf{K}_1, \dots, \mathsf{K}_n)$ by computing $\mathsf{K}_1 = \mathsf{Decap}(\mathsf{sk}, c_0)$, $\dots, \mathsf{K}_n = \mathsf{Decap}(\mathsf{sk}_n, c_0)$. Next, it decrypts c_1 to m via $\mathsf{m} = \mathsf{Dec}(\mathsf{K}, c_1)$.

Correctness of this scheme follows routinely from the correctness of its components.

Note that if the incompressible SKE scheme (Enc, Dec) is rate-1, then so is our public-key scheme (KeyGen', Enc', Dec'), as the only additional information in ciphertexts $c = (c_0, c_1)$ is the header c_0 , which consists of just two group elements. On the other hand, note that the size of the public key of this scheme scales with the size n of the symmetric key K, which in our symmetric-key construction scales with the size of the message m.

Security of the Full Construction We will now turn to sketching the security proof for the main construction. In the first hybrid step (somewhat expectedly), we use the HPS Decap algorithm instead of the Encap algorithm to compute the key-bits K_i in the encryption of the challenge ciphertext. That is, in the encryption of the challenge ciphertext we replace $K_i = \text{Encap}(pk_i, c_0, r)$ with $K_i = \text{Decap}(sk_i, c_0)$ for all i = 1, ..., n. Due to the correctness property of the HPS, this modification does not change the distribution of K. Hence this hybrid change goes unnoticed by the adversary. In the second hybrid step, since we don't need r anymore, we replace $c_0 = (g^r, h^r)$ with a uniformly random c'_0 . We can use the DDH assumption to argue that this modification goes unnoticed.

The next hybrid step is the critical one: We choose g, h, the pk_i and c'_0 with auxiliary information, i.e. together with their discrete logarithms with respect to g, choose $\mathsf{K} \leftarrow \{0,1\}^n$ uniformly at random and sample each sk_i such that $\mathsf{K}_i = \mathsf{Decap}'(\mathsf{sk}_i, c'_0)$ using the programming algorithm of the programmable HPS. We can argue statistical indistinguishability using the programmability property of HPS. The crucial observation now is that the public key $\mathsf{PK} = (g, h, \mathsf{pk}_1, \dots, \mathsf{pk}_n)$ and the ciphertext header c_0 are computed *independently* of K and SK, and in fact we choose SK depending on K, i.e. we can choose SK after everything else.

This now allows us to turn an adversary \mathcal{A} with non-negligible advantage in this hybrid experiment into an adversary \mathcal{A}' with the same advantage against the incompressible SKE scheme. \mathcal{A}' first generates PK as in the hybrid experiment and provides PK to the first stage \mathcal{A}_1 of \mathcal{A} , which will output m_0, m_1 . Now the second stage \mathcal{A}'_2 gets to see a symmetric-key encryption c_1 of m_b , and turns this into a public-key encryption by setting $c = (c_0, c_1)$, where c_0 computed as in the hybrid experiment. This ciphertext c is then given \mathcal{A}_2 , which outputs a state/leak st, and \mathcal{A}'_2 outputs the same state st.

Finally, \mathcal{A}'_3 given a symmetric key K and the state st proceeds as follows. Using the auxiliary information aux^8 , and the key K, it samples a secret key $\mathsf{SK} = (\mathsf{sk}_1, \ldots, \mathsf{sk}_n)$ such that for all $i = 1, \ldots, n$ it holds that $\mathsf{K}_i = \mathsf{Decap}'(\mathsf{sk}_i, c_0)$, as in the hybrid experiment. Then, \mathcal{A}'_3 runs \mathcal{A}_3 on SK and st and outputs whatever \mathcal{A}_3 outputs.

It is not hard to see that from the view of \mathcal{A} , \mathcal{A}' simulates the hybrid experiment perfectly. Hence, the advantage of \mathcal{A}' against the incompressible symmetric-key security experiment is the same as that of \mathcal{A} against the hybrid experiment, and we derive the desired contradiction.

Reducing the Public-Key-Size. As mentioned above, the construction we discussed in the last two paragraphs has a near-optimal ciphertext size (i.e. increasing the size of the symmetric-key ciphertext only by two group elements). In contrast, it has a very large public key which scales *linearly* with the size of the encrypted messages/the ciphertexts.

We will now discuss a tradeoff which achieves a better balance between ciphertext size and public key size. Concretely, we will provide a tradeoff which achieves a ciphertext size of $n + n^{\epsilon} \text{poly}(\lambda)$ for an $0 < \epsilon < 1$ and public key size $n^{\epsilon'} \text{poly}(\lambda)$ for an $1/2 < \epsilon' < 1$. I.e. we achieve ciphertext rate $1 - n^{\epsilon-1} \text{poly}(\lambda)$, which approaches 1 for sufficiently large n, while having a key of sublinear size.

⁸ There is a technical subtlety in the security definition of incompressible SKE which we omitted before: We allow the first stage \mathcal{A}'_1 of a symmetric-key adversary \mathcal{A}' to produce a large state (i.e. scaling with the message size), which is provided to *both* \mathcal{A}'_2 and \mathcal{A}'_3 . This is to *communicate* a potentially large public key PK from \mathcal{A}_1 to \mathcal{A}_3 without putting a burden on the leakage-budget of the leaker-stage \mathcal{A}'_2 . One could consider an alternative definition where this communication from \mathcal{A}'_1 to \mathcal{A}'_3 is not allowed. In such a setting we could still prove our construction secure by *compressing* the auxiliary information aux from which PK and c_0 are generated using a PRG

In order to declutter the presentation, we will switch from multiplicative notation of group operations in \mathbb{G} to additive notion in the following discussion. That is we will denote group elements g^x by [x], and write $\alpha \cdot [x]$ instead of $(g^x)^{\alpha}$. Furthermore, we will consider vectors and matrices of group elements, i.e. if $\mathbf{x} \in \mathbb{Z}_p^k$ is a vector, then $[\mathbf{x}]$ is its element-wise encoding in the group \mathbb{G} . Likewise, we write an encoding of a matrix $\mathbf{A} \in \mathbb{Z}_p^{k \times l}$ as $[\mathbf{A}]$.

In our discussion above we considered a HPS for the *two-dimensional* DDH language, i.e. the language consisting of all $r \cdot [\mathbf{v}]$ given two $[\mathbf{v}]$, where $\mathbf{v} \in \mathbb{Z}_p^2$ is a randomly chosen 2-dimensional vector over \mathbb{Z}_p .

Thus let $\mathbf{v} \in \mathbb{Z}_p^k$ be a randomly chosen k-dimensional vector. The goal of the k-dimensional DDH problem is to distinguish $([\mathbf{v}], t \cdot [\mathbf{v}])$ from $([\mathbf{v}], [\mathbf{u}])$, where \mathbf{v} and \mathbf{u} are chosen uniformly random from \mathbb{Z}_p^k and r is chosen uniformly from \mathbb{Z}_p . It follows routinely via a standard rerandomization argument that the k-dimensional DDH problem is hard, given that the 2-dimensional DDH problem is hard.

We can construct an HPS for k-DDH analogously to the 2-dimensional case: Fix a vector $[\mathbf{v}] \in \mathbb{G}^k$. The secret key sk is a random vector $\boldsymbol{\alpha} \in \mathbb{Z}_p^k$, whereas the public key is given by $[\mathbf{pk}] = \boldsymbol{\alpha}^\top[\mathbf{v}]$, i.e. the inner product of $\boldsymbol{\alpha}$ and $[\mathbf{v}]$. Given a vector $[\mathbf{w}] = r \cdot [\mathbf{v}]$ and a witness r, the Encap algorithm computes $[\mathbf{K}] = r \cdot [\mathbf{pk}]$. On the other hand, given any vector $[\mathbf{w}] \in \mathbb{G}^k$ and a secret key sk = $\boldsymbol{\alpha}$, the Decap algorithm computes $\boldsymbol{\alpha}^\top \cdot [\mathbf{w}]$. Arguing correctness and smoothness are again simple exercises in linear algebra. Furthermore, this HPS satisfies a stronger notion of k - 1-smoothness: Given uniformly random $[\mathbf{w}_1], \ldots, [\mathbf{w}_{k-1}]$, it holds that

$$(\mathsf{pk}, \boldsymbol{\alpha}^{\top}[\mathbf{w}_1], \dots, \boldsymbol{\alpha}^{\top}[\mathbf{w}_{k-1}]) \approx_s (\mathsf{pk}, [u_1], \dots, [u_{k-1}]),$$

where the $[u_1], \ldots, [u_{k-1}]$ are uniformly random in \mathbb{G} . Establishing this is again routine linear algebra.

We will first briefly discuss how the HPS can be made programmable. In essence, we follow the same idea as above: We take a balance function $\mathsf{HC} : \mathbb{G} \to \{0, 1\}$ and define the Decap algorithm to compute $\mathsf{HC}(\boldsymbol{\alpha}^{\top}[\mathbf{w}])$. We claim this construction is k - 1-programmable. That is, given $[\mathbf{v}], [\mathsf{pk}] = [t],$ uniformly random $[\mathbf{w}_1], \ldots, [\mathbf{w}_{k-1}]$ together with the witnesses \mathbf{v}, t and $\mathbf{w}_1, \ldots, \mathbf{w}_{k-1}$, and a random $\mathsf{K} = (\mathsf{K}_1, \ldots, \mathsf{K}_{k-1}) \in \{0, 1\}^{k-1}$, we can efficiently sample a uniformly random $\boldsymbol{\alpha} \in \mathbb{Z}_p^k$ such that $t = \boldsymbol{\alpha}^\top \mathbf{v}$ and $\mathsf{K}_i = \mathsf{HC}(\boldsymbol{\alpha}^\top[\mathbf{w}_i])$ for $i = 1, \ldots, n$. We proceed as above: First we choose uniformly random $z_i \in \mathbb{Z}_p$ such that $\mathsf{K}_i = \mathsf{HC}([z_i])$ for all i. Then we get the linear equation system

$$\boldsymbol{\alpha}^{\top} \mathbf{v} = t$$
$$\boldsymbol{\alpha}^{\top} \mathbf{w}_1 = z_1$$
$$\vdots$$
$$\boldsymbol{\alpha}^{\top} \mathbf{w}_{k-1} = z_{k-1}.$$

Since the \mathbf{w}_i are chosen uniformly random, this system has full rank w.o.p., and hence we can find a matching secret key $\boldsymbol{\alpha}$ via simple linear algebra.

Now, plugging this programmable HPS into our construction of incompressible PKE, we obtain the following parameters.

- A single public pk consisting of one group element can be used to encapsulate k key bits. Hence, to encapsulate n key bits we need n/k public keys amounting to n/k group elements.
- The ciphertext header now contains $k \cdot (k-1) \leq k^2$ group elements (in the above notation the vectors $[\mathbf{w}_1], \ldots, [\mathbf{w}_{k-1}]$).

Hence, if we want to strike a balance where the (additive) ciphertext overhead is of the same size as the public key, we obtain the relation

$$\frac{n}{k} = k^2,$$

which yields to $k = n^{1/3}$. Hence, for this choice of parameters the public key consists of a $n^{2/3}$ group elements (which is sublinear), and the size of the ciphertext is $n + n^{2/3} \log(|\mathbb{G}|) = n(1 - n^{-1/3} \log(|\mathbb{G}|))$ bits, which approaches rate 1.

2.5 Extension to CCA security

The scheme described so far achieves IND-CPA incompressible security. This work also considers an IND-CCA2 incompressible security definition where the adversary gets oracle access to a decryption oracle.

To achieve IND-CCA2 security, we follow the framework of [17]. We add a second hash proof system that acts as integrity proof for ciphertexts. The second hash proof system does not need to be programmable but universal₂ [17] or 2-smooth [1]. It allows the decryption oracle to only answer queries to honestly generated ciphertexts. This mechanism ensures that the decryption oracle does not give up entropy of the programmable HPS's secret key.

In the main body of this work, we provide the full construction that achieves this level of security.

3 Preliminaries

The acronym PPT denotes "probabilistic polynomial time". Throughout this work, λ denotes the security parameter. By $\operatorname{negl}(\lambda)$, we denote a negligible function in λ , that is, a function that vanishes faster than any inverse polynomial in λ . Let $n \in \mathbb{N}$. Then, [n] denotes the set $\{1, \ldots, n\}$. If \mathcal{A} is an algorithm, we denote by $y \leftarrow \mathcal{A}(x)$ the output y after running \mathcal{A} on input x. If S is a (finite) set, we denote by $x \leftarrow S$ the experiment of sampling uniformly at random an element x from S. If D is a distribution over S, we denote by $x \leftarrow D$ the element x sampled from S according to D.

For two probability distributions X, Y, we use the notation $X \approx_s Y$ to state that the distributions are statistically indistinguishable and $X \approx_c Y$ to state that the distributions are computationally indistinguishable.

For ease of notation, in any of our constructions we assume public parameters p are known to every algorithm and every secret key sk also contains the corresponding public key pk.

We present some information-theoretical notions and results that will be instrumental throughout this work.

Definition 1 (Average Min-Entropy [23]). For two jointly distributed random variables (X, Y), the average min-entropy of X conditioned on Y is defined as

$$\hat{H}_{\infty}(X|Y) = -log(\mathbb{E}_{y \leftarrow \$Y}[max_x \Pr[X = x|Y = y]])$$

Lemma 1 (Lemma 2.2 b) of [23]). For random variables X, Y, Z where Y is supported over a set of size T, we have

$$\tilde{H}_{\infty}(X|(Y,Z)) \ge \tilde{H}_{\infty}((X,Y)|Z) - \log(T) \ge \tilde{H}_{\infty}(X|Z) - \log(T).$$

Definition 2 (Average-Case Extractor [23]). Let $n, d, m \in \mathbb{N}$. A function $\mathsf{Ext} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ is a (k,ϵ) strong average-case min-entropy extractor if, for all random variables (X,Y) where X takes values in $\{0,1\}^n$ and $\tilde{H}_{\infty}(X|Y) \ge k$, we have that $(U_d,\mathsf{Ext}(X,U_d),Y)$ is ϵ -close to (U_d,U_m,Y) , where U_d and U_m are independent uniformly random strings of length d and m respectively.

Lemma 2 (Generalized Leftover Hash Lemma 2.4 of [23]). Let $n, m \in \mathbb{N}$. Let $\{H_r : \{0,1\}^n \to \{0,1\}^m\}_{r\in R}$ be a family of universal hash functions, then $\mathsf{Ext}(x,r) \mapsto H_r(x)$ is an average-case (k,ϵ) -strong extractor whenever $m \leq k - 2\log(\frac{1}{\epsilon}) + 2$.

Definition 3 (Pseudorandom Generator). Let $n, m = \text{poly}(\lambda)$. A function $G : \{0, 1\}^n \to \{0, 1\}^m$ is a pseudorandom generator if, for uniformly random $s \leftarrow \{0, 1\}^n$ and $r \leftarrow \{0, 1\}^m$, we have

$$G(\mathbf{s}) \approx_c r.$$

Definition 4 (Collision-Resitant Hash Function). Let $n, m, l = \text{poly}(\lambda)$. A collision-resistant hash function is a seeded function $\text{CRHF} : \{0,1\}^n \times \{0,1\}^m \to \{0,1\}^l$ with the property that for all PPT adversaries \mathcal{A} , random seed $\mathbf{s} \in \{0,1\}^n$ we have $\mathcal{A}(\mathbf{s})$ outputs x, x' with $x \neq x'$ and $\text{CRHF}_{\mathbf{s}}(x) = \text{CRHF}_{\mathbf{s}}(x')$ with negligible probability.

3.1 Decisional Diffie-Hellman assumption

In the following, let \mathcal{G} be a (prime-order) group generator, that is, \mathcal{G} is an algorithm that takes as an input a security parameter 1^{λ} and outputs (\mathbb{G}, p, g), where \mathbb{G} is the description of a multiplicative cyclic group, p is the order of the group which is always a prime number unless differently specified, and g is a generator of the group. Sometimes we denote the size of the group by $|\mathbb{G}|$.

We denote by [a] be value g^a . Similarly, if $\mathbf{A} \in \mathbb{Z}_p^{n \times m}$ is a matrix with entries $a_{i,j}$ then [A] denotes the matrix where each (i, j)-entry is the value $g^{a_{i,j}}$. Note that given $\mathbf{x} \in \mathbb{Z}_p^n$, $\mathbf{y} \in \mathbb{Z}_p^m$ and [A], we can compute $\mathbf{x}^T[\mathbf{A}] = [\mathbf{x}^T\mathbf{A}]$ and $[\mathbf{A}]\mathbf{y} = [\mathbf{A}\mathbf{y}]$.

In the following we state the decisional version of the Diffie-Hellman (DDH) assumption.

Definition 5 (Decisional Diffie-Hellman Assumption). Let $(\mathbb{G}, p, g) \leftarrow \mathfrak{G}(1^{\lambda})$. We say that the DDH assumption holds (with respect to \mathcal{G}) if for any PPT adversary \mathcal{A}

$$|\Pr[1 \leftarrow \mathcal{A}((\mathbb{G}, p, g), ([a], [b], [ab]))] - \Pr[1 \leftarrow \mathcal{A}((\mathbb{G}, p, g), ([a], [b], [c]))]| \le \mathsf{negl}(\lambda)$$

where $a, b, c \leftarrow \mathbb{Z}_p$.

3.2 Public-Key Encryption

Definition 6 (Public-Key Encryption). A public-key encryption (PKE) scheme is a triple of PPT algorithms

 $(pk, sk) \leftarrow KeyGen(1^{\lambda})$: Given the security parameter λ the key-generation algorithm outputs a public key pk and a secret key sk.

 $c \leftarrow Enc(pk,m)$: Given a public key pk and a message m encryption outputs a ciphertext c.

 $m \leftarrow Dec(sk, c)$: Given a secret key sk and a ciphertext c decryption outputs a message m.

Correctness. For all $\lambda, S \in \mathbb{N}$, messages m and (pk, sk) in the range of KeyGen we have that m = Dec(sk, Enc(pk, m)).

IND-CPA security. For all $\lambda \in \mathbb{N}$ and all adversaries $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ we have that

$$\Pr\left[b \leftarrow \mathcal{A}_2(\mathsf{st}, c) : \begin{array}{c} (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda) \\ (\mathsf{m}_0, \mathsf{m}_1, \mathsf{st}) \leftarrow \mathcal{A}_1(\mathsf{pk}) \\ b \leftarrow \$ \ \{0, 1\} \\ c \leftarrow \mathsf{Enc}(\mathsf{pk}, \mathsf{m}_b) \end{array} \right] \leq \frac{1}{2} + \mathsf{negl}(\lambda).$$

3.3 HILL-Entropic Encodings

We recall the notion of HILL-entropic encodings from [42].

Definition 7 (HILL-Entropic Encodings [42]). An (α, β) -HILL-entropic encoding scheme with selective security in the CRS setting consists of two PTT algorithms:

- $-c \leftarrow \mathsf{En}_{\mathsf{crs}}(1^{\lambda},\mathsf{m})$: An encoding algorithm that takes a common random string crs and a message m producing an encoding c.
- $-m \leftarrow \mathsf{De}_{\mathsf{crs}}(c)$: A decoding algorithm that takes a common random string crs and an encoding c and produces a message m .

Correctness. There is some negligible μ such that for all $\lambda \in \mathbb{N}$ and all $m \in \{0,1\}^*$ we have

$$\Pr[\mathsf{De}_{\mathsf{crs}}(\mathsf{En}_{\mathsf{crs}}(1^{\lambda},\mathsf{m})) = \mathsf{m}] = 1 - \mu(\lambda).$$

 α -Expansion. For all $\lambda, k \in \mathbb{N}$ and all $\mathbf{m} \in \{0,1\}^k$ we have $|\mathsf{En}_{\mathsf{crs}}(1^\lambda, \mathbf{m})| \leq \alpha(\lambda, k)$.

 β -HILL-Entropy. There exists an algorithm SimEn s.t. for any polynomial $k = k(\lambda)$ and any ensamble of messages $\mathbf{m} = \{\mathbf{m}_{\lambda}\}$ of length $|\mathbf{m}_{\lambda}| = k(\lambda)$, consider the following "real" experiment:

$$-\operatorname{crs} \leftarrow \$ \{0,1\}^{t(\lambda,k)} - c \leftarrow \operatorname{En}_{\operatorname{crs}}(1^{\lambda}, \mathsf{m}_{\lambda})$$

and let CRS, C denote the random variables for the corresponding values in the "real" experiment. Also consider the following "simulated" experiment:

$$-(\mathsf{crs}',c') \leftarrow \mathsf{SimEn}(1^{\lambda},\mathsf{m}_{\lambda})$$

and let CRS', C' denote the random variables for the corresponding values in the "simulated" experiment. We require that $(CRS, C) \approx_c (CRS', C')$ and $\tilde{H}_{\infty}(C'|CRS') \geq \beta(\lambda, k)$.

We call a (α,β) -HILL-entropic encoding good if $\alpha(\lambda,k) = k(1+o(1)) + \text{poly}(\lambda)$ and $\beta(\lambda,k) = k(1-o(1)) - \text{poly}(\lambda)$. Moran and Wichs [42] provide good HILL-entropic encodings from DCR [44,20] or LWE [46] in the CRS model. They also show that the CRS must be as big as the encoded message.

4 Incompressible Symmetric-Key Encryption

In this section, we define incompressible symmetric-key encryption (SKE) and give a construction from entropic encodings.

4.1 Definition

First, we recall the notion of forward-secure storage [26] under the name of incompressible symmetric-key encryption. For our purposes we only need IND-EAV style security but this could be extended similar to what we did with incompressible public-key encryption.

Definition 8 (Incompressible SKE). An incompressible symmetric-key encryption scheme is a tuple of PPT algorithms using uniformly random keys k

 $c \leftarrow Enc(k, m)$: Given a symmetric key k and a message m encryption it outputs a ciphertext c. m $\leftarrow Dec(sk, c)$: Given a symmetric key k and a ciphertext c decryption it outputs a message m.

We require size of message space, size of key space, and size of ciphertext space to be polynomials over the security parameter λ and the space bound S; that is, $n = n(\lambda, S)$, $k = k(\lambda, S)$, and $l = l(\lambda, S)$ respectively.

Correctness For all $\lambda, S \in \mathbb{N}$, messages m and keys $k \in \{0, 1\}^k$ we have that m = Dec(k, Enc(k, m))

Security For security parameter λ and space bound S, a symmetric-key encryption scheme (Enc, Dec) has incompressible SKE security if for all PPT adversaries $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ the probability of winning the following experiment is $\leq \frac{1}{2} + \operatorname{negl}(\lambda)$.

 $\mathsf{Dist}_{\mathcal{A},\Pi}^{\mathsf{IncomSKE}}(\lambda, S)$ **Experiment** :

- Run the adversary $(\mathsf{m}_0, \mathsf{m}_1, \mathsf{st}_1) \leftarrow \mathcal{A}_1(1^{\lambda})$ to receive two messages m_0 and m_1
- Sample a bit $b \leftarrow \{0, 1\}$ uniformly at random
- Sample $\mathsf{k} \leftarrow \{0,1\}^{n(\lambda,S)}$ uniformly at random
- Run $c \leftarrow \mathsf{Enc}(\mathsf{k},\mathsf{m}_b)$ to encrypt m_b
- Run the adversary $\mathsf{st}_2 \leftarrow \mathcal{A}_2(\mathsf{st}_1, c)$ to produce a state st_2 smaller than S
- Run the final adversary $b' \leftarrow \mathcal{A}_3(\mathsf{k},\mathsf{st}_1,\mathsf{st}_2,\mathsf{m}_0,\mathsf{m}_1)$
- The adversary wins if b = b'

4.2Construction

Now we show how to build incompressible symmetric-key encryption using HILL-entropic encodings, extractors, and pseudorandom generators.

Construction 1 Let λ be the security parameter, S be the space bound of the adversary and n be the size of the message space. Let (En, De) be an (α, β) -HILL-entropic encoding, $\mathsf{Ext} : \{0, 1\}^{\alpha(\check{\lambda}, n)} \times \{0, 1\}^{d(\lambda)} \to \mathbb{C}$ $\{0,1\}^{\lambda}$ be a $(\beta(\lambda,n) - S, \mathsf{negl}(\lambda))$ strong average-case min-entropy extractor where $d(\lambda)$ is a polynomial and $G: \{0,1\}^{\lambda} \to \{0,1\}^{n}$ be a PRG.

Enc(k,m):

- Parse $\mathbf{k} = (\mathbf{k}_1, \mathbf{k}_2, \mathbf{crs}).$ - Sample $s \leftarrow \{0,1\}^{\lambda}$ uniformly at random. - Let $c_1 \leftarrow \mathsf{En}_{\mathsf{crs}}(1^{\lambda}, G(\mathsf{s}) \oplus \mathsf{m}).$ - Let $c_2 \leftarrow \mathsf{s} \oplus \mathsf{Ext}(c_1,\mathsf{k}_1) \oplus \mathsf{k}_2$. $- Return \ c = (c_1, c_2).$ $- \textit{ Parse } \mathsf{k} = (\mathsf{k}_1, \mathsf{k}_2, \mathsf{crs}).$

Dec(k, c):

- Parse $c = (c_1, c_2)$.
- Let $\mathsf{s} \leftarrow \mathsf{Ext}(c_1,\mathsf{k}_1) \oplus c_2 \oplus \mathsf{k}_2$.
- Return $\mathsf{De}_{\mathsf{crs}}(c_1) \oplus G(\mathsf{s})$.

Parameters. The ciphertexts are of size $\lambda + \alpha(\lambda, n)$. The keys are of size $d(\lambda) + t(\lambda, n)$, where $t(\lambda, n)$ is the size of the encoding's crs. Notice that the extractor exists if $\beta(\lambda, n) - S - 2\log\left(\frac{1}{\operatorname{negl}(\lambda)} + 2\right) \geq \lambda$

according to Lemma 2. So, the adversary is allowed a leakage of size $S \leq \beta(\lambda, n) - \lambda - 2\log\left(\frac{1}{\mathsf{negl}(\lambda)} + 2\right)$. Therefore, if we choose a "good" entropic encoding we get a rate of $\frac{n}{n(1+o(1))+\text{poly}(\lambda)}$, allowed leakage of $S = n(1 - o(1)) - \text{poly}(\lambda)$, and keysize of $k = n(1 + o(1)) + \text{poly}(\lambda)$.

Correctness. By the correctness of the entropic encoding $\mathsf{De}_{\mathsf{crs}}(\mathsf{En}_{\mathsf{crs}}(\mathsf{1}^{\lambda}, G(\mathsf{s}) \oplus \mathsf{m})) = G(\mathsf{s}) \oplus \mathsf{m}$. Since Ext is deterministic under a fixed key k_1 then $Ext(c_1, k_1) \oplus c_2 \oplus k_2 = Ext(c_1, k_1) \oplus s \oplus Ext(c_1, k_1) = s$. Therefore, $\mathsf{De}_{\mathsf{crs}}(c_1) \oplus G(\mathsf{s}) = \mathsf{m}$.

Theorem 2 (Security). The incompressible SKE presented in Construction 1 has incompressible SKE security if (En,De) is an (α,β) -HILL-entropic encoding, Ext is a $(\beta(\lambda,n)-S, \operatorname{negl}(\lambda))$ strong average-case min-entropy extractor, and G is a pseudorandom generator each with the listed parameters.

Proof. We prove security via hybrids. First we list the hybrid and then argue their indistinguishability. In each hybrid we highlight the changes compared to the previous one.

 H_0 :

- Run the adversary $\mathbf{m}_0, \mathbf{m}_1, \mathbf{st}_1 \leftarrow \mathcal{A}_1(1^{\lambda})$ to receive two messages \mathbf{m}_0 and \mathbf{m}_1 .
- Sample bit $b \leftarrow \{0, 1\}$ uniformly at random.
- Sample $\mathsf{k} \leftarrow \{0,1\}^n$ uniformly at random.
- Run $c \leftarrow \mathsf{Enc}(\mathsf{k},\mathsf{m}_b)$ to encrypt m_b .
- Run the adversary $\mathsf{st}_2 \leftarrow \mathcal{A}_2(\mathsf{st}_1, c)$ to produce a state st_2 smaller than S.
- Run the final adversary $b' \leftarrow \mathcal{A}_3(\mathsf{k}, \mathsf{st}_1, \mathsf{st}_2, \mathsf{m}_0, \mathsf{m}_1)$.
- The adversary wins if b = b'.

In H_1 we explicitly represent what happens in Enc.

 H_1 :

- Run the adversary $\mathbf{m}_0, \mathbf{m}_1, \mathbf{st}_1 \leftarrow \mathcal{A}_1(1^{\lambda})$ to receive two messages \mathbf{m}_0 and m_1 .
- Sample $k_1 \leftarrow \{0,1\}^{d(\lambda,n)}$ uniformly at random.
- Sample $k_2 \leftarrow \{0,1\}^{\lambda}$ uniformly at random.
- Sample crs \leftarrow $\{0,1\}^{t(\lambda,n)}$ uniformly at random.

- Sample $s \leftarrow \{0,1\}^{\lambda}$ uniformly at random.
- Let $c_1 \leftarrow \mathsf{En}_{\mathsf{crs}}(1^\lambda, G(\mathsf{s}) \oplus \mathsf{m}_b)$.
- Let $c_2 \leftarrow \mathsf{s} \oplus \mathsf{Ext}(c_1,\mathsf{k}_1) \oplus \mathsf{k}_2$.
- Let $c \leftarrow (c_1, c_2)$ and $\mathsf{k} \leftarrow (\mathsf{k}_1, \mathsf{k}_2, \mathsf{crs})$.
- Run the adversary $\mathsf{st}_2 \leftarrow \mathcal{A}_2(\mathsf{st}_1, c)$ to produce a state st_2 smaller than S.
- Run the final adversary $b' \leftarrow \mathcal{A}_3(\mathsf{k}, \mathsf{st}_1, \mathsf{st}_2, \mathsf{m}_0, \mathsf{m}_1)$.
- The adversary wins if b = b'

In H_2 we switch the entropic encoding to the simulated code that has a lot of entropy.

H_2 :

- Run the adversary $\mathbf{m}_0, \mathbf{m}_1, \mathbf{st}_1 \leftarrow \mathcal{A}_1(1^{\lambda})$ to receive two messages \mathbf{m}_0 and m_1 .
- Sample bit $b \leftarrow \{0, 1\}$ uniformly at random.
- Sample $k_1 \leftarrow \{0,1\}^{d(\lambda,n)}$ uniformly at random. Sample $k_2 \leftarrow \{0,1\}^{\lambda}$ uniformly at random.

- Sample $\mathbf{s} \leftarrow \{0,1\}^{\lambda}$ uniformly at random.
- Let $(\operatorname{crs}, c_1) \leftarrow \operatorname{SimEn}(1^{\lambda}, G(\mathsf{s}) \oplus \mathsf{m}_b).$
- Let $c_2 \leftarrow \mathsf{s} \oplus \mathsf{Ext}(c_1,\mathsf{k}_1) \oplus \mathsf{k}_2$.
- Let $c \leftarrow (c_1, c_2)$ and $\mathsf{k} \leftarrow (\mathsf{k}_1, \mathsf{k}_2, \mathsf{crs})$.
- Run the adversary $\mathsf{st}_2 \leftarrow \mathcal{A}_2(\mathsf{st}_1, c)$ to produce a state st_2 smaller than S.
- Run the final adversary $b' \leftarrow \mathcal{A}_3(\mathsf{k}, \mathsf{st}_1, \mathsf{st}_2, \mathsf{m}_0, \mathsf{m}_1)$.
- The adversary wins if b = b'.

I H_3 we switch the order in which we sample c_2 and k_2 .

 H_3 :

- Run the adversary $\mathbf{m}_0, \mathbf{m}_1, \mathbf{st}_1 \leftarrow \mathcal{A}_1(1^{\lambda})$ to receive two messages \mathbf{m}_0 and m_1 .
- Sample bit $b \leftarrow \{0, 1\}$ uniformly at random.
- Sample $k_1 \leftarrow \{0,1\}^{d(\lambda,n)}$ uniformly at random.
- Sample $\mathbf{s} \leftarrow \{0,1\}^{\lambda}$ uniformly at random.
- Let $(\operatorname{crs}, c_1) \leftarrow \operatorname{Sim} \operatorname{En}(1^{\lambda}, G(\mathbf{s}) \oplus \mathbf{m}_b).$
- Sample $c_2 \leftarrow \{0, 1\}^{\lambda}$ uniformly at random.
- Let $c \leftarrow (c_1, c_2)$
- Run the adversary $\mathsf{st}_2 \leftarrow \mathcal{A}_2(\mathsf{st}_1, c)$ to produce a state st_2 smaller than S.
- Let $\mathbf{k}_2 \leftarrow c_2 \oplus \mathsf{Ext}(c_1, \mathbf{k}_1) \oplus \mathsf{s}$.
- Let $\mathbf{k} \leftarrow (\mathbf{k}_1, \mathbf{k}_2, \mathbf{crs})$.
- Run the final adversary $b' \leftarrow \mathcal{A}_3(\mathsf{k}, \mathsf{st}_1, \mathsf{st}_2, \mathsf{m}_0, \mathsf{m}_1)$.
- The adversary wins if b = b'.

In H_4 we replace the output of the extractor Ext by a uniformly random value.

H_4 :

- Run the adversary $\mathbf{m}_0, \mathbf{m}_1, \mathbf{st}_1 \leftarrow \mathcal{A}_1(1^{\lambda})$ to receive two messages \mathbf{m}_0 and m_1 .
- Sample bit b ←s {0,1} uniformly at random.
 Sample k₁ ←s {0,1}^{d(λ,n)} uniformly at random.
- Sample $s \leftarrow \{0, 1\}^{\lambda}$ uniformly at random.
- Let $(\operatorname{crs}, c_1) \leftarrow \operatorname{Sim} \operatorname{En}(1^{\lambda}, G(\mathbf{s}) \oplus \mathbf{m}_b).$
- Sample $c_2 \leftarrow \{0,1\}^{\lambda}$ uniformly at random.
- Let $c \leftarrow (c_1, c_2)$.
- Run the adversary $\mathsf{st}_2 \leftarrow \mathcal{A}_2(\mathsf{st}_1, c)$ to produce a state st_2 smaller than S.

- Sample $k_2 \leftarrow \{0,1\}^{\lambda}$ uniformly at random.
- Let $\mathbf{k} \leftarrow (\mathbf{k}_1, \mathbf{k}_2, \mathsf{crs})$.
- Run the final adversary $b' \leftarrow \mathcal{A}_3(\mathsf{k}, \mathsf{st}_1, \mathsf{st}_2, \mathsf{m}_0, \mathsf{m}_1)$.
- The adversary wins if b = b'.

Finally we replace the output of G(s) by a uniformly random value.

H_5 :

- Run the adversary $\mathbf{m}_0, \mathbf{m}_1, \mathbf{st}_1 \leftarrow \mathcal{A}_1(1^{\lambda})$ to receive two messages \mathbf{m}_0 and m_1 . - Sample bit $b \leftarrow \{0, 1\}$ uniformly at random.

- Sample $k_1 \leftarrow \{0,1\}^{d(\lambda,n)}$ uniformly at random.
- Sample $s \leftarrow \{0,1\}^{\lambda}$ uniformly at random.
- Sample $r \leftarrow \{0, 1\}^n$ uniformly at random.
- Let $(\operatorname{crs}, c_1) \leftarrow \operatorname{SimEn}(1^{\lambda}, r)$.
- Sample $c_2 \leftarrow \{0,1\}^{\lambda}$ uniformly at random.
- Let $c \leftarrow (c_1, c_2)$.
- Run the adversary $\mathsf{st}_2 \leftarrow \mathcal{A}_2(\mathsf{st}_1, c)$ to produce a state st_2 smaller than S.
- Sample $k_2 \leftarrow \{0, 1\}^{\lambda}$ uniformly at random.
- Let $\mathbf{k} \leftarrow (\mathbf{k}_1, \mathbf{k}_2, \mathsf{crs})$.
- Run the final adversary $b' \leftarrow \mathcal{A}_3(\mathsf{k}, \mathsf{st}_1, \mathsf{st}_2, \mathsf{m}_0, \mathsf{m}_1)$.
- The adversary wins if b = b'.

$H_0 \approx H_1$:

The differences between H_0 and H_1 are purely syntactical. In H_1 we just show more detail of Enc. $H_1 \approx_c H_2$:

Instead of sampling the common random string for the entropic encoding uniformly at random and then encoding $G(s) \oplus m$ we simulate both steps using SimEn. Assume there exists a PPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ that can distinguish the two hybrids H_1 and H_2 with a non-negligible advantage of ϵ . From this we construct a PPT adversary $\mathcal{A}' = (\mathcal{A}'_1, \mathcal{A}'_2)$ that can break the β -HILL-entropy of (En, De) with advantage ϵ .

$$4'_{1}(1^{\lambda})$$
 :

- Run the adversary $\mathsf{m}_0, \mathsf{m}_1, \mathsf{st}_1 \leftarrow \mathcal{A}_1(1^{\lambda})$ to receive two messages m_0 and m_1
- Sample bit $b \leftarrow \{0, 1\}$ uniformly at random
- Sample $k_1 \leftarrow \{0,1\}^{d(\lambda,n)}$ uniformly at random
- Sample $\mathbf{s} \leftarrow \{0,1\}^{\lambda}$ uniformly at random
- Return $G(s) \oplus m_b$

 $\mathcal{A}_2'(\mathsf{crs}, c_1)$:

- Let $c_2 \leftarrow \mathsf{s} \oplus \mathsf{Ext}(c_1,\mathsf{k}_1)$

- Let $c \leftarrow (c_1, c_2)$
- Run the adversary $\mathsf{st}_2 \leftarrow \mathcal{A}_2(\mathsf{st}_1, c)$ to produce a state st_2 smaller than S
- Run the final adversary $b' \leftarrow \mathcal{A}_3(\mathsf{k}, \mathsf{st}_1, \mathsf{st}_2, \mathsf{m}_0, \mathsf{m}_1)$
- Return b'

If \mathcal{A} can distinguish H_1 from H_2 then \mathcal{A}' can distinguish a uniformly random $\mathsf{crs} \leftarrow \{0,1\}^{t(\lambda,n)}$ and $c_1 \leftarrow \mathsf{En}(1^{\lambda}, G(\mathsf{s}) \oplus \mathsf{m}_b)$ from $(\mathsf{crs}, c_1) \leftarrow \mathsf{Sim}\mathsf{En}(1^{\lambda}, G(\mathsf{s}) \oplus \mathsf{m}_b)$ as it perfectly simulates H_2 in the case that $(crs, c_1) \leftarrow SimEn(1^{\lambda}, G(s) \oplus m_b)$ and perfectly simulates H_1 in the other case.

$$H_2 \approx H_3$$

In H_3 we switch the order in which we sample c_2 and k_2 . From the view of the adversary this is statistically identical.

 $H_3 \approx_s H_4$:

Let C_1 , C_2 , CRS, K_1 , K_2 , and ST_2 denote the random variables for the corresponding values in the experiment and U_{λ} independent uniform randomness of length λ . By the β -HILL entropy of the entropic encoding we know that $\tilde{H}_{\infty}(C_1|CRS) \geq \beta$. Using Lemma 1 we deduce that

$$H_{\infty}(C_1|(CRS, K_2, ST_2, C_2)) \ge \beta - 2\lambda - \log(S)$$

Therefore, the extractor gives us that $(K_1, K_2, CRS, ST_2, U_\lambda)$ and $(K_1, K_2, CRS, ST_2, \mathsf{Ext}(C_1, K_1))$ are statistically close.

 $H_3 \approx_c H_4$:

In H_4 we encode a uniformly random string instead of $G(\mathbf{s}) \oplus \mathbf{m}_b$. Assume there exists a PPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ that can distinguish the two hybrids H_3 and H_4 with a non-negligible advantage of ϵ . From this we construct a PPT adversary \mathcal{A}' that can break the pseudorandomness of G with advantage ϵ .

- $\mathcal{A}'(r')$:
 - Run the adversary $\mathbf{m}_0, \mathbf{m}_1, \mathbf{st}_1 \leftarrow \mathcal{A}_1(1^{\lambda})$ to receive two messages \mathbf{m}_0 and m_1
 - Sample bit $b \leftarrow \{0, 1\}$ uniformly at random
 - Sample $k_1 \leftarrow \{0,1\}^{d(\lambda,n)}$ uniformly at random
 - Sample $\mathbf{s} \leftarrow \{0,1\}^{\lambda}$ uniformly at random
 - Sample $r \leftarrow r' \oplus \mathsf{m}_b$ uniformly at random
 - Let $(\operatorname{crs}, c_1) \leftarrow \operatorname{SimEn}(1^{\lambda}, r)$
 - Sample $c_2 \leftarrow \{0, 1\}^{\lambda}$ uniformly at random
 - Let $c \leftarrow (c_1, c_2)$
 - Run the adversary $\mathsf{st}_2 \leftarrow \mathcal{A}_2(\mathsf{st}_1, c)$ to produce a state st_2 smaller than S
 - Run the final adversary $b' \leftarrow \mathcal{A}_3(\mathsf{k}, \mathsf{st}_1, \mathsf{st}_2, \mathsf{m}_0, \mathsf{m}_1)$
 - Return b'

If \mathcal{A} can distinguish H_3 from H_4 then \mathcal{A}' can distinguish G(s) with uniformly random $s \leftarrow \{0,1\}^{\lambda}$ from uniformly random $r' \leftarrow \{0,1\}^n$ as it perfectly simulates H_3 in the case that $r' \leftarrow G(s)$ and perfectly simulates H_4 in the other case.

H_4 :

In H_4 the winning probability of the adversary is $\frac{1}{2}$ as it gets no information about b at all.

5 Programmable Hash Proof Systems

In this work we think of a hash proof systems as a key encapsulation mechanism where the encapsulated key is independent of the public key and the ciphertext under certain conditions. This allows us to later resample the secret key in the incompressibility experiments.

For our construction we need two different hash proof systems. One that is Y-programmable and one that is 2-smooth both using the same language.

5.1 Definitions

First we define hash proof system that we will use as a mask in our encryption scheme.

Definition 9 (Y-Programmable Hash Proof System [17,38]). A Y-programmable hash proof system is defined over a NP language $\mathcal{L} \subset X$, where each element x in the language \mathcal{L} has a witness w. Additionally there exist a subset $Y \subset X \setminus \mathcal{L}$ and efficient ways to sample a language \mathcal{L} with a corresponding trapdoor $\mathsf{td}_{\mathcal{L}}$, an $x \in \mathcal{L}$ with its witness w and an $x \in Y$ with a corresponding trapdoor td_x

- $-(\mathbf{p}, \mathsf{td}_{\mathcal{L}}) \leftarrow \mathsf{Gen}(1^{\lambda}, 1^{k})$: Given the security parameter λ , the encapsulated key size k the language generation algorithm that outputs public parameters \mathbf{p} defining a language \mathcal{L} and a trapdoor $\mathsf{td}_{\mathcal{L}}$ to that language.
- $-(x \in \mathcal{L}, w) \leftarrow \mathsf{samp}\mathcal{L}(\mathsf{p})$: Given the public parameters, it outputs an element $x \in \mathcal{L}$ with the corresponding witness w.
- $-(x \in Y, \mathsf{td}_x) \leftarrow \mathsf{samp}Y(\mathsf{p}, \mathsf{td}_{\mathcal{L}})$: Given the public parameters and a trapdoor $\mathsf{td}_{\mathcal{L}}$, it outputs $x \in Y$ and the corresponding trapdoor td_x .

The hash proof system itself consists of these algorithms:

- $-(pk, sk) \leftarrow KeyGen(p)$: Given the public parameters, the key generation algorithm outputs a public key pk and a secret key sk.
- $k \leftarrow Encap(pk, x, w)$: Given the public lye pk, en element x and a witness w. the key encapsulation algorithm outputs an encapsulated key k.
- $\mathsf{k} \leftarrow \mathsf{Decap}(\mathsf{sk}, x)$: Given the secret key sk and any $x \in X$, the key decapsulation algorithm outputs an encapsulated key. k . Notice x can be outside \mathcal{L} .
- $-\mathsf{sk}' \leftarrow \mathsf{Program}(\mathsf{td}_{\mathcal{L}}, \mathsf{td}_x, \mathsf{sk}, x, \mathsf{k})$ Given two trapdoors $\mathsf{td}_{\mathcal{L}}, \mathsf{td}_x$, a secret key sk , an element $x \in Y$, and an encapsulated key k , the programming algorithm outputs a new secret key sk' .

Correctness. For all $\lambda, k \in \mathbb{N}$, $(\mathbf{p}, \mathsf{td}_{\mathcal{L}})$ in the range of $\mathsf{Gen}(1^{\lambda}, 1^{k})$, $(\mathbf{pk}, \mathsf{sk})$ in the range of $\mathsf{KeyGen}(\mathbf{p})$, $x \in \mathcal{L}$ and for $\mathsf{k} \leftarrow \mathsf{Encap}(\mathsf{pk}, x, w)$, we have $\mathsf{k} = \mathsf{Decap}(\mathsf{sk}, x)$ with $|\mathsf{k}| = k$.

Language Indistinguishability. For all $\lambda, k \in \mathbb{N}$ if we sample $(\mathsf{p}, \mathsf{td}_{\mathcal{L}}) \leftarrow \mathsf{Gen}(1^{\lambda}, 1^{k}), \mathcal{L} \ni x \leftarrow \mathsf{samp}\mathcal{L}(\mathsf{p}),$ and $(x^* \in Y, \mathsf{td}_{x^*}) \leftarrow \mathsf{samp}Y(\mathsf{p}, \mathsf{td}_{\mathcal{L}})$, we have the computational indistinguishability: $x \approx_c x^*$.

Programmability. For all $\lambda, k \in \mathbb{N}$, $(\mathbf{p}, \mathsf{td}_{\mathcal{L}})$ in the range of samp $\mathcal{L}(1^{\lambda}, 1^{k})$, $(\mathbf{pk}, \mathbf{sk})$ in the range of $\mathsf{KeyGen}(\mathsf{p}), \mathsf{k} \in \{0,1\}^{m}, \text{ and for } (x, \mathsf{td}_{x}) \text{ in the range of } \mathsf{samp}Y(\mathsf{p}, \mathsf{td}_{\mathcal{L}}), \mathsf{sk}' \leftarrow \mathsf{Program}(\mathsf{td}_{\mathcal{L}}, \mathsf{td}_{x}, \mathsf{sk}, x, \mathsf{k}),$ we have $\mathsf{Decap}(\mathsf{sk}', x) = \mathsf{k}$.

Y-Programmable Smoothness. For all $\lambda, k \in \mathbb{N}$, $(\mathsf{p}, \mathsf{td}_{\mathcal{L}})$ in the range of $\mathsf{Gen}(1^{\lambda}, 1^{k})$, $(\mathsf{pk}, \mathsf{sk})$ in the range of KeyGen(p), (x, td_x) in the range of samp $Y(\mathsf{p}, \mathsf{td}_{\mathcal{L}}), \mathsf{k} \in \{0, 1\}^m$, and $\mathsf{sk}' \leftarrow \mathsf{Program}(\mathsf{td}_{\mathcal{L}}, \mathsf{td}_x, \mathsf{sk}, x, \mathsf{k})$ we have statistical indistinguishability $(\mathsf{pk},\mathsf{sk},x) \approx_s (\mathsf{pk},\mathsf{sk}',x)$.

Notice, if $Y = X \setminus \mathcal{L}$ then Y-programmable smoothness implies smoothness.

Next we recall 2-smooth hash proof systems with our adjusted notation.

Definition 10 (2-Smooth Hash Proof System [17,1]). A 2-smooth hash proof system is defined over a NP language $\mathcal{L} \subset X$ as above The hash proof system itself consists of the following algorithms:

- $-(pk, sk) \leftarrow KeyGen(p)$: Given the public parameters, the key generation algorithm that outputs a public key pk and a secret key sk.
- $-k \leftarrow \mathsf{Encap}(\mathsf{pk}, x, w, \tau)$: Given public key pk, an element of the language $x \in \mathcal{L}$, its witness w, and a tag τ , the key encapsulation algorithm outputs an encapsulated key k.
- $-\mathsf{k} \leftarrow \mathsf{Decap}(\mathsf{sk}, x, \tau)$: Given the secret key sk , any $x \in X$, and a tag τ . the key decapsulation algorithm outputs an encapsulated key k. Notice x can be outside \mathcal{L} .

Correctness. For all $\lambda, k \in \mathbb{N}$, $(\mathbf{p}, \mathsf{td}_{\mathcal{L}})$ in the range of $\mathsf{Gen}(1^{\lambda}, 1^{k})$, $(\mathbf{pk}, \mathsf{sk})$ in the range of $\mathsf{KeyGen}(\mathbf{p})$, $x \in \mathcal{L}$, tags τ , and for $\mathsf{k} \leftarrow \mathsf{Encap}(\mathsf{pk}, x, w, \tau)$, we have $\mathsf{k} = \mathsf{Decap}(\mathsf{sk}, x, \tau)$ with $|\mathsf{k}| = k$.

Language Indistinguishability. Exactly as above.

2-Smoothness. For all $\lambda, k \in \mathbb{N}$, $(\mathsf{p}, \mathsf{td}_{\mathcal{L}})$ in the range of $\mathsf{Gen}(1^{\lambda}, 1^{k}), x, x' \in X \setminus \mathcal{L}$, two tags τ, τ' such that $(x,\tau) \neq (x',\tau')$, let $(\mathsf{pk},\mathsf{sk}) \leftarrow \mathsf{KeyGen}(\mathsf{p})$ and sample $\mathsf{k} \leftarrow \{0,1\}^k$ we have computational indistinguishability between (pk, Decap(sk, x, τ), Decap(sk, x', τ')) and (pk, Decap(sk, x, τ), k).

5.2Programmable Hash Proof System from DDH

In our protocols we need programmable HPS with a big encapsulated key space (for classic notation [17] this would be called the hash space).

Some smooth hash proof systems are easily transformed into programmable HPS with big encapsulated keys by generating more public keys and using them on the same $x \in X$. These HPS include the one from weak pseudorandom effective group actions [2]. That transformation causes the public key size to scale linearly with the size of the encapsulated key and leave the size of the ciphertext indepent of the encapsulated key size. We provide more details about this in the full version [11].

We present a variant of the original [17] HPS with an interesting trade off. Here both public key size and ciphertext size scale in the 2/3-power with k, the size of the encapsulated key.

Construction 2 Let $\mathsf{HC} : \mathbb{G} \times \{0,1\}^{\log(|\mathbb{G}|)} \to \{0,1\}$ denote a 1-bit randomness extractor over a group element; if this function is applied over a matrix of group elements, then it means that the function is applied entry-wise with the same randomness. In the following let $\ell, s \in \mathbb{N}$ such that $\ell \cdot s = k$. We get an interesting tradeoff for our application when $\ell = k^{1/3}$ and $s = k^{2/3}$.

 $Gen(1^{\lambda}, 1^{k})$:

- $\begin{array}{l} (\mathbb{G}, p, g) \leftarrow \$ \ \mathcal{G}(1^{\lambda}). \\ Sample \ \mathbf{h} \leftarrow \$ \ \mathbb{Z}_{p}^{\ell} \setminus \{\mathbf{0}\} \ uniformly \ at \ random. \\ \ Return \ \mathbf{p} = (\mathbb{G}, p, g, [\mathbf{h}]) \ and \ \mathsf{td}_{\mathcal{L}} = \mathbf{h}. \end{array}$

samp $\mathcal{L}(p)$: $\begin{array}{l} - \ Parse \ \mathsf{p} = (\mathbb{G}, p, g, [\mathbf{h}]). \\ - \ Sample \ \mathbf{y} \leftarrow \$ \mathbb{Z}_p^{\ell-1} \ uniformly \ at \ random. \end{array}$ - Return $x = [\mathbf{y}]$ and $w = \mathbf{y}$. $\mathsf{samp}Y(\mathsf{p},\mathsf{td}_{\mathcal{L}}):$ - Parse $\mathbf{p} = (\mathbb{G}, p, q, [\mathbf{h}]).$ $- Let td_{\mathcal{L}} = h.$ $- Sample \mathbf{E} \leftarrow \mathbb{Z}_p^{\ell \times (\ell-1)} \text{ such that } (\mathbf{h} \mathbf{E}) \text{ is invertible uniformly at random.}$ - Return $x = [\mathbf{E}]$ and $w = \mathbf{E}$. KeyGen(p): - Parse $\mathbf{p} = (\mathbb{G}, p, g, [\mathbf{h}]).$ $\begin{array}{l} - \ Sample \ r \leftarrow \$ \ \{0,1\}^{log(|\mathbb{G}|)} \ the \ public \ randomness \ for \ a \ extractor. \\ - \ Sample \ \mathbf{A} \leftarrow \$ \ \mathbb{Z}_p^{s \times \ell} \ uniformly \ at \ random. \\ - \ Return \ \mathbf{pk} = (\mathbf{A}[\mathbf{h}], r) \ and \ \mathbf{sk} = \mathbf{A}. \end{array}$ Encap $(\mathbf{pk}, c = [\mathbf{hy}^t], w = \mathbf{y})$: - Parse $\mathbf{p} = (\mathbb{G}, p, q, [\mathbf{h}] \in \mathbb{G}^{\ell}).$ - Parse $\mathsf{pk} = ([\mathbf{f}] \in \mathbb{G}^s, r).$ - Let $\mathbf{K} \leftarrow \mathsf{HC}([\mathbf{f}]\mathbf{y}^t, r)$ the component-wise extractor of the outer product between \mathbf{f} and \mathbf{y} . - Return k = K. Decap $(\mathsf{sk}, x = [\mathbf{E}] \in \mathbb{G}^{\ell \times (\ell-1)})$: - Parse pk = ([f], r) - Parse $\mathsf{sk} = \mathbf{A} \in \mathbb{Z}_p^{s \times \ell}$. - Let $\mathbf{K} \leftarrow \mathsf{HC}(\mathbf{A}[\mathbf{E}], r)$ the component-wise extractor of the product between \mathbf{A} and $[\mathbf{E}]$.

- Return $\mathbf{k} = \mathbf{K}$.

 $Program(td_{\mathcal{L}}, td_x, sk, x, k)$:

- Parse $\mathsf{pk} = ([\mathbf{f}], r), \, \mathsf{td}_{\mathcal{L}} = \mathbf{h} \in \mathbb{Z}_p^{\ell}, \, \mathsf{td}_x = \mathbf{E} \in \mathbb{Z}_p^{\ell \times (\ell-1)}, \, \mathsf{sk} = \mathbf{A}, \, and \, \mathsf{k} = \mathbf{K} \in \{0, 1\}^{s \times (\ell-1)}.$ For each $i \in [\ell-1], j \in [s]$ sample $B_{i,j} \leftarrow \mathbb{Z}_p$ such that $K_{i,j} = \mathsf{HC}([B_{i,j}], r)$ via rejection sampling.
- Set $\mathbf{B} = (B)_{i,j}$. Let $\mathbf{A}' \leftarrow (\mathbf{Ah} \mathbf{B}) (\mathbf{h} \mathbf{E})^{-1}$.
- Return $\mathbf{sk}' = \mathbf{A}'$.

Correctness. For any $(\mathbf{p} = (\mathbb{G}, p, q, [\mathbf{h}]), \mathsf{td}_{\mathcal{L}})$ in the range of Gen, $(\mathbf{pk} = ([\mathbf{Ah}], r), \mathsf{sk} = \mathbf{A})$ in the range of KeyGen, and $[\mathbf{hy}^t] \in \mathcal{L}$ we have $\mathsf{Encap}(\mathsf{pk}, [\mathbf{hy}^t])$ outputs $\mathsf{k} = \mathsf{HC}([\mathbf{Ah}]\mathbf{y}^t, r) = \mathsf{HC}([\mathbf{Ahy}^t], r)$. Decapsulation then outputs $\mathbf{k} = \mathsf{HC}(\mathbf{A}[\mathbf{h}\mathbf{y}^t], r) = \mathsf{HC}([\mathbf{A}\mathbf{h}\mathbf{y}^t], r).$

Programmability. Since we choose \mathbf{h} and \mathbf{E} s.t. ($\mathbf{h} \mathbf{E}$) is invertible Program always outputs a matrix \mathbf{A}' with the property that $\mathbf{A'E} = \mathbf{B}$ and $\mathbf{k} = \mathsf{HC}([\mathbf{B}], r)$.

Programmable Smoothness. If we first sample k uniformly random and then program for the key k $Program(td_{\mathcal{L}}, td_x, sk, x, k)$ the resulting distribution over **B** will be uniformly random. And because (**h E**) is invertible then \mathbf{A}' is a uniformly random under the condition that $\mathbf{A}'\mathbf{h} = \mathbf{A}\mathbf{h}$. The same holds for \mathbf{A} . Therefore, $(\mathsf{pk}, \mathsf{sk} = \mathbf{A}, x)$ and $(\mathsf{pk}, \mathsf{sk}' = \mathbf{A}', x)$ are identically distributed.

Theorem 3 (Language Indistinguishability). If DDH is hard for \mathcal{G} then elements from the language $\mathcal{L} = \{[\mathbf{h}]\mathbf{y}^t | \mathbf{y} \in \mathbb{Z}_p^{\ell-1}\}$ and $Y = \{[\mathbf{E}] | \mathbf{E} \in \mathbb{Z}_p^{\ell \times (\ell-1)} \land (\mathbf{h} \mathbf{E}) \text{ is invertible}\}$ of construction 2 are indistinguishable.

Proof. We prove security via hybrids. First we list the hybrids and then argue their indistinguishability. In each hybrid we highlight the changes compared to the previous one.

 H_0 :

- Let $(\mathsf{p}, \mathsf{td}_{\mathcal{L}}) \leftarrow \mathsf{Gen}(1^{\lambda}, 1^k)$.
- Let $(pk, sk) \leftarrow KeyGen(p)$.
- Let $(x, w) \leftarrow \mathsf{samp}\mathcal{L}(\mathsf{p})$.
- Let $\mathsf{k} \leftarrow \mathsf{Encap}(\mathsf{pk}, x, w)$.
- Run the adversary $\mathcal{A}(\mathsf{pk},\mathsf{sk},x)$.

 H_1 :

- Sample a group $(\mathbb{G}, p, g) \leftarrow \mathfrak{G}(1^{\lambda})$.
- Sample $r \leftarrow \{0,1\}^{\log(|\mathbb{G}|)}$ the randomness for the extractor
- Sample $\mathbf{h} \leftarrow \mathbb{Z}_p^{\ell} \setminus \{\mathbf{0}\}$ uniformly at random
- Sample $\mathbf{A} \leftarrow \mathbb{Z}_p^{s \times \ell}$ uniformly at random
- Let $\mathbf{p} = (\mathbb{G}, p, g, [\mathbf{h}])$.
- Let $\mathsf{pk} = ([\mathbf{Ah}], r)$ and $\mathsf{sk} = \mathbf{A}$.
- Sample $\mathbf{y} \leftarrow \mathbb{Z}_p^{\ell-1}$ uniformly at random
- Let $x = [\mathbf{h}\mathbf{y}^t] = [\mathbf{C}]$.
- Run the adversary $\mathcal{A}(\mathsf{pk},\mathsf{sk},x)$.

$H_{2,i}$:

- Sample a group $(\mathbb{G}, p, g) \leftarrow \mathfrak{G}(1^{\lambda}).$
- Sample $r \leftarrow \{0,1\}^{\log(|\mathbb{G}|)}$ the randomness for the extractor.
- Sample $\mathbf{h} \leftrightarrow \mathbb{Z}_p^{\ell} \setminus \{\mathbf{0}\}$ uniformly at random. Sample $\mathbf{A} \leftarrow \mathbb{Z}_p^{s \times \ell}$ uniformly at random.
- Let $\mathbf{p} = (\mathbb{G}, p, g, [\mathbf{h}]).$
- Let $\mathsf{pk} = ([\mathbf{Ah}], r)$ and $\mathsf{sk} = \mathbf{A}$. Sample $\mathbf{y} \leftarrow \mathbb{Z}_p^{\ell-1}$ uniformly at random. Let $[\mathbf{C}] = [\mathbf{hy}^t]$.
- Sample $\mathbf{E} \leftarrow \mathbb{Z}_p^{l \times (l-1)}$ uniformly at random
- Replace the first *i* entries of $[\mathbf{C}]$ by the first *i* entries in $[\mathbf{E}]$
- Let $x = [\mathbf{C}].$
- Run the adversary $\mathcal{A}(\mathsf{pk},\mathsf{sk},x)$.

H_3 :

- $\begin{array}{l} \text{ Sample a group } (\mathbb{G}, p, g) \xleftarrow{} \mathcal{G}(1^{\lambda}). \\ \text{ Sample } r \xleftarrow{} \{0, 1\}^{\log(|\mathbb{G}|)} \text{ the randomness for the extractor.} \end{array}$
- Sample $\mathbf{h} \leftarrow \mathbb{Z}_p^{\ell} \setminus \{\mathbf{0}\}$ uniformly at random. Sample $\mathbf{A} \leftarrow \mathbb{Z}_p^{s \times \ell}$ uniformly at random.
- Let $\mathbf{p} = (\mathbb{G}, p, g, [\mathbf{h}]).$
- Let $\mathsf{pk} = ([\mathbf{Ah}], r)$ and $\mathsf{sk} = \mathbf{A}$.
- Sample $\mathbf{E} \leftarrow \mathbb{Z}_p^{l \times (l-1)}$ uniformly at random such that $(\mathbf{h} \mathbf{E})$ is invertible
- Let $x = [\mathbf{E}]$.
- Run the adversary $\mathcal{A}(\mathsf{pk},\mathsf{sk},x)$.

$H_0 \approx H_1$:

The differences between H_0 and H_1 are purely syntactical. In H_1 we just show more detail of Gen and Encap.

 $H_1 \approx H_{2,0}$:

The differences between H_1 and $H_{2,0}$ are purely syntactical.

$$H_{2,i} \approx_c H_{2,i+1}$$

In $H_{2,i+1}$ we replace the n+1st element of C by a random one. Assume there exists a PPT adversary \mathcal{A} that can distinguish the two hybrids $H_{2,i}$ and $H_{2,i+1}$ with a non-negligible advantage of ϵ . From this we construct a PPT adversary \mathcal{A}' that can break DDH with advantage ϵ .

- \mathcal{A}' ((\mathbb{G}, p, q), ([a], [b], [ρ])):
 - Let $u \leftarrow i \mod l$
 - Let $v \leftarrow |i/l|$
 - Sample $r \leftarrow \{0,1\}^{\log(|\mathbb{G}|)}$ the randomness for the extractor
 - Sample $\mathbf{h} \leftarrow \mathbb{Z}_p^{\ell} \setminus \{\mathbf{0}\}$ uniformly at random
 - Replace $[x_u]$ by [a]

- Sample $\mathbf{A} \leftarrow \mathbb{Z}_p^{s \times \ell}$ uniformly at random
- Let $\mathbf{p} = (\mathbb{G}, p, g, [\mathbf{h}])$
- Let $\mathsf{pk} = ([\mathbf{Ah}], r)$ and $\mathsf{sk} = \mathbf{A}$
- Sample $\mathbf{y} \leftarrow \mathbb{Z}_p^{\ell-1}$ uniformly at random

- For
$$u' \in [l]$$
 and $v' \in [l-1]$ let $C_{u',v'} \leftarrow \begin{cases} [\rho] & \text{if } u' = u, v' = v \\ [b]x_{u'} & \text{if } u' \neq u, v' = v \\ [x_{u'}]y_{v'} & \text{else} \end{cases}$

- Sample $\mathbf{E} \leftarrow \mathbb{Z}_p^{l \times (l-1)}$
- Replace the first i entries of [C] by the first i entries in [E]
- Let $x = [\mathbf{C}]$
- Run the adversary $b' \leftarrow \mathcal{A}(\mathsf{pk},\mathsf{sk},x)$
- Return b'

If \mathcal{A} distinguishes between $H_{2,i}$ and $H_{2,i+1}$ then \mathcal{A}' distinguishes between $\rho = ab$ and ρ being uniformly random as \mathcal{A}' perfectly simulates $H_{2,i}$ in the case that $\rho = ab$ and $H_{2,i+1}$ if r is uniformly random.

 $H_{2,m} \approx_s H_3$:

 $H_{2,m}$ is statistically close to H_3 because with probability $1 - \operatorname{\mathsf{negl}}(\lambda)$ we have (**h E**) is invertible.

Parameters. For an encapsulated key of size k this scheme roughly gets us public parameters of size $k^{1/3} \cdot \mathsf{poly}(\lambda)$, public key of size $k^{2/3} \cdot \mathsf{poly}(\lambda)$ and elements from X of size $k^{2/3} \cdot \mathsf{poly}(\lambda)$.

5.32-Smooth Hash Proof System from DDH

The above hash proof system only is programmable if $x \in Y$. To make our encryption scheme CCA secure we need a efficient way to check whether $x \in \mathcal{L}$ or $x \in X \setminus \mathcal{L}$. To do this we construct the 2-smooth hash proof system below that is defined over the same language.

Construction 3 We construct a 2-smooth hash proof system with a output size of λ using an extractor $\mathsf{Ext}: \mathbb{G}^{\ell-1} \times \{0,1\}^p \to \{0,1\}^{\lambda}$ and a collision resistant hash function CRHF that maps into \mathbb{Z}_p . As a language description we use the same as in Construction 2.

KeyGen(p):

- Parse $\mathbf{p} = (\mathbb{G}, p, g, [\mathbf{h}]).$ Sample $r \leftarrow \{0, 1\}^{\log(|\mathbb{G}|)}$ uniformly at random.
- Sample $\mathbf{s} \leftarrow \{0,1\}^{\lambda}$ uniformly at random. Sample $\mathbf{a}, \mathbf{b} \leftarrow \{\mathbf{Z}_p^{\ell} \text{ uniformly at random.}$
- Return $\mathsf{pk} = (\mathbf{a}^t[\mathbf{h}], \mathbf{b}^t[\mathbf{h}], r, \mathsf{s})$ and $\mathsf{sk} = (\mathbf{a}, \mathbf{b})$.

 $\mathsf{Encap} \ (\mathsf{pk}, x = [\mathbf{hy}^t] \in \mathbb{G}^{\ell \times (\ell-1)}, w = \mathbf{y} \in \mathbb{Z}_p^{\ell-1}, \tau):$

- $Parse_{\mathbf{p}} = (\mathbb{G}, p, g, [\mathbf{h}] \in \mathbb{G}^{\ell}) \text{ and } \mathbf{pk} = ([f], [f'] \in \mathbb{G}, r, \mathsf{s}).$
- $Let [\mathbf{d}] = ([f]\mathbf{y}) + (\mathsf{CRHF}_{\mathsf{s}}(x,\tau)[f']\mathbf{y}).$

- Return $\mathbf{k} = \mathsf{Ext}([\mathbf{d}^t], r)$. Decap $(\mathsf{sk}, x = [\mathbf{E}] \in \mathbb{G}^{\ell \times (\ell-1)}, \tau)$:

- Parse $\mathbf{p} = (\mathbb{G}, p, g, [\mathbf{h}] \in \mathbb{G}^{\ell})$ and $\mathbf{pk} = ([f], [f'] \in \mathbb{G}, r, s)$.
- $Parse \, \mathsf{sk} = (\mathbf{a} \in \mathbb{Z}_p^{\ell}, \mathbf{b} \in \mathbb{Z}_p^{\ell}).$
- Parse $x = [\mathbf{E}] \in \mathbb{G}^{\ell \times (\ell-1)}$.
- Return $\mathbf{k} = \mathsf{Ext}(\mathbf{a}^t[\mathbf{E}] + \mathsf{CRHF}_{\mathsf{s}}(x,\tau)\mathbf{b}^t[\mathbf{E}], r).$

Correctness. For any $(\mathbf{p} = (\mathbb{G}, p, g, [\mathbf{h}]), \mathsf{td}_{\mathcal{L}})$ in the range of Gen, $(\mathbf{pk} = (\mathbf{a}^t[\mathbf{h}], \mathbf{b}^t[\mathbf{h}], r, \mathbf{s}), \mathsf{sk} = (\mathbf{a}, \mathbf{b}))$ in the range of KeyGen, and $[\mathbf{h}\mathbf{y}^t] \in \mathcal{L}$ we have $\mathsf{Encap}(\mathsf{pk}, [\mathbf{h}\mathbf{y}^t])$ outputs

$$\mathsf{k} = \mathsf{Ext}\left(\left(\left([f]\mathbf{y}\right) + \left([f]\mathsf{CRHF}_{\mathsf{s}}(x,\tau)\mathbf{y}\right)\right)^{t}, r\right) = \mathsf{Ext}\left(\left[(\mathbf{a}^{t}\mathbf{h})\mathbf{y}^{t} + \mathsf{CRHF}_{\mathsf{s}}(x,\tau)(\mathbf{b}^{t}\mathbf{h})\mathbf{y}^{t}\right], r\right)$$

On the other hand, decapsulation outputs

$$\mathsf{k} = \mathsf{Ext}(\mathbf{a}^t[\mathbf{h}\mathbf{y}^t] + \mathsf{CRHF}_\mathsf{s}(x,\tau)\mathbf{b}^t[\mathbf{h}\mathbf{y}^t], r) = \mathsf{Ext}\left([(\mathbf{a}^t\mathbf{h})\mathbf{y}^t + \mathsf{CRHF}_\mathsf{s}(x,\tau)(\mathbf{b}^t\mathbf{h})\mathbf{y}^t], r\right)$$

Language Indistinguishability. Since we use the same language as in construction 2 the language indistinguishability holds by the same argument.

2-Smoothness. For all $\lambda, n \in \mathbb{N}$, $(\mathsf{p}, \mathsf{td}_{\mathcal{L}})$ in the range of $\mathsf{Gen}(\mathsf{p})$, $x, x' \in X \setminus \mathcal{L}$, two tags τ, τ' such that $(x, \tau) \neq (x', \tau')$, let $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(\mathsf{p})$ and sample $\mathsf{k} \leftarrow \$ \{0, 1\}^m$. Let $\gamma \leftarrow \mathsf{CRHF}_{\mathsf{s}}(x, \tau)$ and $\gamma' \leftarrow \mathsf{CRHF}_{\mathsf{s}}(x', \tau')$. Using $(x, \tau) \neq (x', \tau')$ and the collision resistance of CRHF we can assume that $\gamma \neq \gamma'$.

In the following let $\mathbf{d} = \mathbf{a}^t[\mathbf{E}] + \gamma \mathbf{b}^t[\mathbf{E}]$ (computed in $\mathsf{Decap}(\mathsf{sk}, x, \tau)$) and $\mathbf{d}' = \mathbf{a}^t[\mathbf{E}'] + \gamma' \mathbf{b}^t[\mathbf{E}']$ (computed in $\mathsf{Decap}(\mathsf{sk}, x', \tau')$). Then the following equation holds:

$$\left(f f' \mathbf{d}^{t} \mathbf{d}^{\prime t}\right) = \left(\mathbf{a}^{t} \mathbf{b}^{t}\right) \left(\begin{array}{cc} \mathbf{h} \mathbf{0} \mathbf{E} \mathbf{E}^{\prime} \\ \mathbf{0} \mathbf{h} \gamma \mathbf{E} \gamma^{\prime} \mathbf{E}^{\prime} \end{array}\right)$$

If $x, x' \in X \setminus \mathcal{L}$ then there exists a column \mathbf{z} with index i in \mathbf{E} s.t. \mathbf{z} is linearly independent of \mathbf{h} and \mathbf{z}' with index i' in \mathbf{E}' s.t. \mathbf{z}' is l.i. of \mathbf{h} . Then the following equation also holds:

$$\left(f f' d_i d'_{i'}\right) = \left(\mathbf{a}^t \mathbf{b}^t\right) \begin{pmatrix} \mathbf{h} \ \mathbf{0} \ \mathbf{z} \ \mathbf{z}' \\ \mathbf{0} \ \mathbf{h} \ \gamma \mathbf{z} \ \gamma' \mathbf{z}' \end{pmatrix}$$

Now, we argue that the matrix on the right side has rank 4. We have that $\begin{pmatrix} \mathbf{h} \\ \mathbf{0} \end{pmatrix}$ and $\begin{pmatrix} \mathbf{0} \\ \mathbf{h} \end{pmatrix}$ are linearly independent. Moreover, $\begin{pmatrix} \mathbf{z} \\ \gamma \mathbf{z} \end{pmatrix}$ is outside the span of $\begin{pmatrix} \mathbf{h} \\ \mathbf{0} \end{pmatrix}$ and $\begin{pmatrix} \mathbf{0} \\ \mathbf{h} \end{pmatrix}$ because \mathbf{h} and \mathbf{z} are linearly independent. Finally, $\begin{pmatrix} \mathbf{z}' \\ \gamma' \mathbf{z}' \end{pmatrix}$ is outside the span of $\begin{pmatrix} \mathbf{z} \\ \gamma \mathbf{z} \end{pmatrix}$, $\begin{pmatrix} \mathbf{h} \\ \mathbf{0} \end{pmatrix}$, and $\begin{pmatrix} \mathbf{0} \\ \mathbf{h} \end{pmatrix}$. To see this, assume that this is not the case, i.e., that there exists a linear combination

$$\begin{pmatrix} \mathbf{z}' \\ \gamma' \mathbf{z}' \end{pmatrix} = c_1 \begin{pmatrix} \mathbf{z} \\ \gamma \mathbf{z} \end{pmatrix} + c_2 \begin{pmatrix} \mathbf{h} \\ \mathbf{0} \end{pmatrix} + c_3 \begin{pmatrix} \mathbf{0} \\ \mathbf{h} \end{pmatrix}$$

Assume there exist $c_1, c_2, c_3 \in \mathbb{N}^+$ such that $\mathbf{z}' = c_1 \mathbf{z} + c_2 \mathbf{h}$ and $\gamma' \mathbf{z}' = c_1 \gamma \mathbf{z} + c_3 \mathbf{h}$. Then we replace \mathbf{z}' in the second equation

$$\gamma'(c_1\mathbf{z} + c_2\mathbf{h}) = c_1\gamma\mathbf{z} + c_3\mathbf{h}$$

$$\Leftrightarrow (\gamma' - \gamma)c_1\mathbf{z} = (c_3 - \gamma'c_2)\mathbf{h}$$

This however can only be true if $\gamma' - \gamma = 0$ because **z** is linearly independent of **h**.

Since **a** and **b** are chosen uniformly at random then so are f, f', d_i , and d'_i . If d_i and d'_i are uniformly random then $\mathsf{Decap}(\mathsf{sk}, x, \tau) = \mathsf{Ext}(\mathbf{d}^t, r)$ and $\mathsf{Decap}(\mathsf{sk}, x', \tau') = \mathsf{Ext}(\mathbf{d}'^t, r)$ are statistically close to uniformly random by the extractor property.

Parameters. For the same language as in Construction 2 with public parameters of size $k^{1/3} \cdot \operatorname{poly}(\lambda)$ and elements of $k^{2/3} \cdot \operatorname{poly}(\lambda)$ construction 3 roughly results in public keys of size $2k^{1/3} \cdot \operatorname{poly}(\lambda)$ and an encapsulated key of size λ .

6 Incompressible PKE from Incompressible SKE and HPS

First we extend the incompressible encryption security notion [32] to the chosen ciphertext scenario and then we show a new construction paradigm using hash proof systems and incompressible symmetric-key encryption.

6.1 CCA Incompressible Encryption

We use the definition of incompressible encryption by Guan et al.[32]. It defines a public-key encryption scheme where the adversary has to know most of the ciphertext to decrypt it even with access to the secret key.

Definition 11 (Incompressible PKE). An incompressible public-key encryption scheme is a triple of PPT algorithms

 $(\mathsf{pk},\mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^{\lambda},1^{S})$: Given the security parameter λ and a space bound S the key-generation algorithm outputs a public key pk and a secret key sk.

 $c \leftarrow \mathsf{Enc}(\mathsf{pk},\mathsf{m})$: Given a public key pk and a message m the encryption algorithm outputs a ciphertext c.

 $\mathsf{m} \leftarrow \mathsf{Dec}(\mathsf{sk}, c)$: Given a secret key sk and a ciphertext c the decryption algorithm outputs a message m .

Both size of message space and size of ciphertext space are polynomials over security parameter λ and space bound S, that is, $n = n(\lambda, S)$ and $l = l(\lambda, S)$ respectively.

Correctness. For all $\lambda, S \in \mathbb{N}$, messages m and (pk, sk) in the range of KeyGen we have that m = Dec(sk, Enc(pk, m)).

CCA Incompressible Security. Similar to standard IND-CCA (sometimes referred to as IND-CCA2) security we extend incompressible encryption such that the adversary has access to an encryption oracle.

For security parameter λ and space bound S, a public key encryption scheme (KeyGen, Enc, Dec) has incompressible CCA PKE security if for all PPT adversaries $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ wins the following experiment with probability $\leq \frac{1}{2} + \operatorname{negl}(\lambda)$.

$\mathsf{Dist}_{\mathcal{A},\Pi}^{\mathsf{CCAIncomPKE}}(\lambda,S)$ **Experiment** :

- Run key generation algorithm $\mathsf{KeyGen}(1^{\lambda}, 1^{S})$ to obtain $(\mathsf{pk}, \mathsf{sk})$.
- Run the adversary $\mathbf{m}_0, \mathbf{m}_1, \mathbf{st}_1 \leftarrow \mathcal{A}_1^{\mathsf{Dec}_{\mathsf{sk}}}(\mathsf{pk})$ on public key pk with oracle access to $\mathsf{Dec}(\mathsf{sk}, \cdot)$ to receive two messages m_0 , m_1 and state st_1 .
- Sample bit $b \leftarrow \{0, 1\}$ uniformly at random.
- Run $c \leftarrow \mathsf{Enc}(\mathsf{pk},\mathsf{m}_b)$ to encrypt m_b .
- Run the adversary $\mathsf{st}_2 \leftarrow \mathcal{A}_2^{\mathsf{Dec}_{\mathsf{sk}}}(\mathsf{pk}, c, \mathsf{st}_1)$ with oracle access to $\mathsf{Dec}(\mathsf{sk}, \cdot)$ for all inputs but c to produce a state st_2 smaller than S.
- Run the final adversary $b' \leftarrow \mathcal{A}_3(\mathsf{sk}, \mathsf{st}_1, \mathsf{st}_2, \mathsf{m}_0, \mathsf{m}_1)$.
- The adversary wins if b = b'.

Rate We define the rate by $\frac{|\mathbf{m}|}{|\mathsf{Enc}(\mathsf{pk},\mathsf{m})|}$ the size of a message divided by a ciphertext encrypting the message. We say a scheme has rate-1 when the rate is 1 - o(1).

6.2 Construction

We construct a encryption scheme that very much resembles the classic Cramer-Shoup [17] scheme. Instead of masking the ciphertext with the randomness that comes out of the hash proof system we use it as a key for an incompressible symmetric-key encryption scheme.

Construction 4 (Incompressible PKE) Given security parameter λ , space bound S, and message length n let (KeyGen', Encap', Decap', Program') be a Y-programmable hash proof system for a language $\mathcal{L} \subset X$ (where you can sample x with according witness from \mathcal{L} and sample x with according trapdoor from Y) where the representation size of X is $p(\lambda, S, n)$ and encapsulated keys of size $k(\lambda, S_{sym}, n)$, (KeyGen", Encap", Decap") is a 2-smooth hash proof system for the same language with encapsulation key size of λ and public key size $p'(\lambda, S, n)$, and (Enc_{sym}, Dec_{sym}) be an incompressible SKE with messages of size n, keys of size $k(\lambda, S_{sym}, n)$ and ciphertexts of size $l(\lambda, S_{sym}, n)$ with incompressible SKE adversary being allowed to leak a state of size $S_{sym} = S + p(\lambda, S, n) + p'(\lambda, S, n)$.

KeyGen $(1^{\lambda}, 1^{S})$:

- Generate language and corresponding trapdoor $(\mathbf{p}, \mathsf{td}_{\mathcal{L}}) \leftarrow \mathsf{Gen}(1^{\lambda}, 1^n)$.
- $\begin{array}{l} \ \ Let \ (pk', sk') \leftarrow KeyGen'(p). \\ \ \ Let \ (pk'', sk'') \leftarrow KeyGen''(p). \end{array}$
- Return pk = (pk', pk'') and sk = (sk', sk'').

Enc(pk, m):

- Parse pk = (pk', pk'')

 $\begin{array}{l} - \ Let \ (x,w) \leftarrow \mathsf{samp}\mathcal{L}(\mathsf{p}). \\ - \ Let \ \mathsf{k} \leftarrow \mathsf{Encap}'(\mathsf{pk}',x,w). \\ - \ Let \ \mathsf{c}_{\mathsf{sym}} \leftarrow \mathsf{Enc}_{\mathsf{sym}}(\mathsf{k},\mathsf{m}). \\ - \ Let \ \pi \leftarrow \mathsf{Encap}''(\mathsf{pk}'',x,w,c_{\mathsf{sym}}). \\ - \ Return \ c = (x,c_{\mathsf{sym}},\pi). \\ \\ \mathsf{Dec}(\mathsf{sk},c): \\ - \ Parse \ \mathsf{sk} = (\mathsf{sk}',\mathsf{sk}''). \\ - \ Parse \ \mathsf{sk} = (\mathsf{sk}',\mathsf{sk}''). \\ - \ If \ \pi = \mathsf{Decap}''(\mathsf{sk}'',x,c_{\mathsf{sym}}) \\ \bullet \ Let \ \mathsf{k} \leftarrow \mathsf{Decap}'(\mathsf{sk}',x) \end{array}$

• Return $m = Dec_{sym}(k, c_{sym})$.

- Return
$$\perp$$
.

Parameters. (KeyGen, Enc, Dec) is an incompressible PKE with messages of size n, ciphertexts of size $l(\lambda, S_{sym}, n) + p(\lambda, S, n) + p'(\lambda, S, n)$, the adversary is allowed a leak of size $S = S_{sym} - p(\lambda, S, n) - p'(\lambda, S, n)$, and the public key is of size $p(\lambda, S, n) + p'(\lambda, S, n)$.

When instantiating the two hash proof systems with constructions 2,3 and the incompressible SKE with construction 1 then (KeyGen, Enc, Dec) is an incompressible PKE with messages of size n, ciphertexts of size $(n + n^{2/3} \text{poly}(\lambda))(1 + o(1))$, the adversary is allowed a leak of size $S = n(1 - o(1)) - \text{poly}(\lambda)(n(1 + o(1)))^{2/3}$, the public key is of size $n^{2/3}(1 + o(1))\text{poly}(\lambda)$, and the secret key is of size $n(1 + o(1))\text{poly}(\lambda)$.

Correctness. Follows from the correctness of (KeyGen', Encap', Decap', Program'), (KeyGen'', Encap'', Decap''), and (Enc_{sym}, Dec_{sym}).

Theorem 4 (Security). The PKE construction 4 has incompressible CCA PKE security if (KeyGen', Encap', Decap', Program') is a programmable hash proof system with the listed parameters, (KeyGen'', Encap'', Decap'') is a 2-smooth hash proof system with the listed parameters, and (Enc_{sym}, Dec_{sym}) is an incompressible secure SKE with the listed parameters.

The proof of this theorem can be found in the full version [11] of this paper.

References

- Abdalla, M., Benhamouda, F., Pointcheval, D.: Disjunctions for hash proof systems: New constructions and applications. In: Oswald, E., Fischlin, M. (eds.) Advances in Cryptology – EUROCRYPT 2015, Part II. Lecture Notes in Computer Science, vol. 9057, pp. 69–100. Springer, Heidelberg, Germany, Sofia, Bulgaria (Apr 26–30, 2015). https://doi.org/10.1007/978-3-662-46803-6_3
- Alamati, N., De Feo, L., Montgomery, H., Patranabis, S.: Cryptographic group actions and applications. In: Moriai, S., Wang, H. (eds.) Advances in Cryptology – ASIACRYPT 2020, Part II. Lecture Notes in Computer Science, vol. 12492, pp. 411–439. Springer, Heidelberg, Germany, Daejeon, South Korea (Dec 7–11, 2020). https://doi.org/10.1007/978-3-030-64834-3_14
- Alwen, J., Dodis, Y., Naor, M., Segev, G., Walfish, S., Wichs, D.: Public-key encryption in the boundedretrieval model. In: Gilbert, H. (ed.) Advances in Cryptology – EUROCRYPT 2010. Lecture Notes in Computer Science, vol. 6110, pp. 113–134. Springer, Heidelberg, Germany, French Riviera (May 30 – Jun 3, 2010). https://doi.org/10.1007/978-3-642-13190-5_6
- Alwen, J., Dodis, Y., Wichs, D.: Leakage-resilient public-key cryptography in the bounded-retrieval model. In: Halevi, S. (ed.) Advances in Cryptology – CRYPTO 2009. Lecture Notes in Computer Science, vol. 5677, pp. 36–54. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 16–20, 2009). https://doi.org/ 10.1007/978-3-642-03356-8_3
- Aumann, Y., Rabin, M.O.: Information theoretically secure communication in the limited storage space model. In: Wiener, M.J. (ed.) Advances in Cryptology – CRYPTO'99. Lecture Notes in Computer Science, vol. 1666, pp. 65–79. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 15–19, 1999). https: //doi.org/10.1007/3-540-48405-1_5
- Barak, B., Goldreich, O., Impagliazzo, R., Rudich, S., Sahai, A., Vadhan, S.P., Yang, K.: On the (im)possibility of obfuscating programs. In: Kilian, J. (ed.) Advances in Cryptology – CRYPTO 2001. Lecture Notes in Computer Science, vol. 2139, pp. 1–18. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 19–23, 2001). https://doi.org/10.1007/3-540-44647-8_1

- Bellare, M., Boldyreva, A., Palacio, A.: An uninstantiable random-oracle-model scheme for a hybridencryption problem. In: Cachin, C., Camenisch, J. (eds.) Advances in Cryptology – EUROCRYPT 2004. Lecture Notes in Computer Science, vol. 3027, pp. 171–188. Springer, Heidelberg, Germany, Interlaken, Switzerland (May 2–6, 2004). https://doi.org/10.1007/978-3-540-24676-3_11
- Bellare, M., Dai, W.: Defending against key exfiltration: Efficiency improvements for big-key cryptography via large-alphabet subkey prediction. In: Thuraisingham, B.M., Evans, D., Malkin, T., Xu, D. (eds.) ACM CCS 2017: 24th Conference on Computer and Communications Security. pp. 923–940. ACM Press, Dallas, TX, USA (Oct 31 – Nov 2, 2017). https://doi.org/10.1145/3133956.3133965
- Bellare, M., Kane, D., Rogaway, P.: Big-key symmetric encryption: Resisting key exfiltration. In: Robshaw, M., Katz, J. (eds.) Advances in Cryptology – CRYPTO 2016, Part I. Lecture Notes in Computer Science, vol. 9814, pp. 373–402. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 14–18, 2016). https: //doi.org/10.1007/978-3-662-53018-4_14
- Black, J.: The ideal-cipher model, revisited: An uninstantiable blockcipher-based hash function. In: Robshaw, M.J.B. (ed.) Fast Software Encryption – FSE 2006. Lecture Notes in Computer Science, vol. 4047, pp. 328– 340. Springer, Heidelberg, Germany, Graz, Austria (Mar 15–17, 2006). https://doi.org/10.1007/11799313_ 21
- Branco, P., Döttling, N., Dujmovic, J.: Rate-1 incompressible encryption from standard assumptions. IACR Cryptol. ePrint Arch. p. 697 (2022), https://eprint.iacr.org/2022/697
- Brzuska, C., Farshim, P., Mittelbach, A.: Random-oracle uninstantiability from indistinguishability obfuscation. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015: 12th Theory of Cryptography Conference, Part II. Lecture Notes in Computer Science, vol. 9015, pp. 428–455. Springer, Heidelberg, Germany, Warsaw, Poland (Mar 23–25, 2015). https://doi.org/10.1007/978-3-662-46497-7_17
- Cachin, C., Maurer, U.M.: Unconditional security against memory-bounded adversaries. In: Kaliski Jr., B.S. (ed.) Advances in Cryptology – CRYPTO'97. Lecture Notes in Computer Science, vol. 1294, pp. 292–306. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 17–21, 1997). https://doi.org/10.1007/ BFb0052243
- Canetti, R., Goldreich, O., Halevi, S.: The random oracle methodology, revisited. J. ACM 51(4), 557–594 (jul 2004). https://doi.org/10.1145/1008731.1008734, https://doi.org/10.1145/1008731.1008734
- Chevalier, C., Fouque, P.A., Pointcheval, D., Zimmer, S.: Optimal randomness extraction from a Diffie-Hellman element. In: Joux, A. (ed.) Advances in Cryptology – EUROCRYPT 2009. Lecture Notes in Computer Science, vol. 5479, pp. 572–589. Springer, Heidelberg, Germany, Cologne, Germany (Apr 26–30, 2009). https://doi.org/10.1007/978-3-642-01001-9_33
- Cramer, R., Shoup, V.: A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In: Krawczyk, H. (ed.) Advances in Cryptology CRYPTO'98. Lecture Notes in Computer Science, vol. 1462, pp. 13–25. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 23–27, 1998). https://doi.org/10.1007/BFb0055717
- Cramer, R., Shoup, V.: Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In: Knudsen, L.R. (ed.) Advances in Cryptology – EUROCRYPT 2002. Lecture Notes in Computer Science, vol. 2332, pp. 45–64. Springer, Heidelberg, Germany, Amsterdam, The Netherlands (Apr 28 – May 2, 2002). https://doi.org/10.1007/3-540-46035-7_4
- Cramer, R., Shoup, V.: Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. SIAM J. Comput. 33(1), 167–226 (2003)
- Damgård, I., Ganesh, C., Orlandi, C.: Proofs of replicated storage without timing assumptions. In: Boldyreva, A., Micciancio, D. (eds.) Advances in Cryptology – CRYPTO 2019, Part I. Lecture Notes in Computer Science, vol. 11692, pp. 355–380. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 18–22, 2019). https://doi.org/10.1007/978-3-030-26948-7_13
- 20. Damgård, I., Jurik, M.: A generalisation, a simplification and some applications of Paillier's probabilistic public-key system. In: Kim, K. (ed.) PKC 2001: 4th International Workshop on Theory and Practice in Public Key Cryptography. Lecture Notes in Computer Science, vol. 1992, pp. 119–136. Springer, Heidelberg, Germany, Cheju Island, South Korea (Feb 13–15, 2001). https://doi.org/10.1007/3-540-44586-2_9
- Dent, A.W.: Adapting the weaknesses of the random oracle model to the generic group model. In: Zheng, Y. (ed.) Advances in Cryptology – ASIACRYPT 2002. Lecture Notes in Computer Science, vol. 2501, pp. 100–109. Springer, Heidelberg, Germany, Queenstown, New Zealand (Dec 1–5, 2002). https://doi.org/10. 1007/3-540-36178-2_6
- 22. Di Crescenzo, G., Lipton, R.J., Walfish, S.: Perfectly secure password protocols in the bounded retrieval model. In: Halevi, S., Rabin, T. (eds.) TCC 2006: 3rd Theory of Cryptography Conference. Lecture Notes in Computer Science, vol. 3876, pp. 225–244. Springer, Heidelberg, Germany, New York, NY, USA (Mar 4–7, 2006). https://doi.org/10.1007/11681878_12
- Dodis, Y., Ostrovsky, R., Reyzin, L., Smith, A.: Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. SIAM Journal on Computing 38(1), 97–139 (2008). https://doi.org/10.1137/060651380, https://doi.org/10.1137/060651380

- 24. Dodis, Y., Quach, W., Wichs, D.: Speak much, remember little: Cryptography in the bounded storage model, revisited. Cryptology ePrint Archive, Report 2021/1270 (2021), https://eprint.iacr.org/2021/1270
- 25. Dziembowski, S.: Intrusion-resilience via the bounded-storage model. In: Halevi, S., Rabin, T. (eds.) TCC 2006: 3rd Theory of Cryptography Conference. Lecture Notes in Computer Science, vol. 3876, pp. 207–224. Springer, Heidelberg, Germany, New York, NY, USA (Mar 4–7, 2006). https://doi.org/10.1007/ 11681878_11
- Dziembowski, S.: On forward-secure storage (extended abstract). In: Dwork, C. (ed.) Advances in Cryptology – CRYPTO 2006. Lecture Notes in Computer Science, vol. 4117, pp. 251–270. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 20–24, 2006). https://doi.org/10.1007/11818175_15
- 27. Garg, R., Lu, G., Waters, B.: New techniques in replica encodings with client setup. In: Pass, R., Pietrzak, K. (eds.) TCC 2020: 18th Theory of Cryptography Conference, Part III. Lecture Notes in Computer Science, vol. 12552, pp. 550–583. Springer, Heidelberg, Germany, Durham, NC, USA (Nov 16–19, 2020). https://doi.org/10.1007/978-3-030-64381-2_20
- Garg, S., Gentry, C., Halevi, S., Raykova, M., Sahai, A., Waters, B.: Candidate indistinguishability obfuscation and functional encryption for all circuits. In: 54th Annual Symposium on Foundations of Computer Science. pp. 40–49. IEEE Computer Society Press, Berkeley, CA, USA (Oct 26–29, 2013). https: //doi.org/10.1109/F0CS.2013.13
- Goldwasser, S., Kalai, Y.T.: On the (in)security of the Fiat-Shamir paradigm. In: 44th Annual Symposium on Foundations of Computer Science. pp. 102–115. IEEE Computer Society Press, Cambridge, MA, USA (Oct 11–14, 2003). https://doi.org/10.1109/SFCS.2003.1238185
- Gorbunov, S., Vaikuntanathan, V., Wee, H.: Functional encryption with bounded collusions via multi-party computation. In: Safavi-Naini, R., Canetti, R. (eds.) Advances in Cryptology – CRYPTO 2012. Lecture Notes in Computer Science, vol. 7417, pp. 162–179. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 19–23, 2012). https://doi.org/10.1007/978-3-642-32009-5_11
- Goyal, R., Koppula, V., Waters, B.: Lockable obfuscation. In: Umans, C. (ed.) 58th Annual Symposium on Foundations of Computer Science. pp. 612–621. IEEE Computer Society Press, Berkeley, CA, USA (Oct 15– 17, 2017). https://doi.org/10.1109/F0CS.2017.62
- Guan, J., Wichs, D., Zhandry, M.: Incompressible cryptography. In: Dunkelman, O., Dziembowski, S. (eds.) Advances in Cryptology – EUROCRYPT 2022, Part I. Lecture Notes in Computer Science, vol. 13275, pp. 700–730. Springer, Heidelberg, Germany, Trondheim, Norway (May 30 – Jun 3, 2022). https://doi.org/ 10.1007/978-3-031-06944-4_24
- Guan, J., Zhandry, M.: Simple schemes in the bounded storage model. In: Ishai, Y., Rijmen, V. (eds.) Advances in Cryptology – EUROCRYPT 2019, Part III. Lecture Notes in Computer Science, vol. 11478, pp. 500–524. Springer, Heidelberg, Germany, Darmstadt, Germany (May 19–23, 2019). https://doi.org/10. 1007/978-3-030-17659-4_17
- 34. Guan, J., Zhandry, M.: Disappearing cryptography in the bounded storage model. In: Nissim, K., Waters, B. (eds.) TCC 2021: 19th Theory of Cryptography Conference, Part II. Lecture Notes in Computer Science, vol. 13043, pp. 365–396. Springer, Heidelberg, Germany, Raleigh, NC, USA (Nov 8–11, 2021). https://doi.org/10.1007/978-3-030-90453-1_13
- Håstad, J., Impagliazzo, R., Levin, L.A., Luby, M.: A pseudorandom generator from any one-way function. SIAM Journal on Computing 28(4), 1364–1396 (1999)
- Hazay, C., López-Alt, A., Wee, H., Wichs, D.: Leakage-resilient cryptography from minimal assumptions. In: Johansson, T., Nguyen, P.Q. (eds.) Advances in Cryptology – EUROCRYPT 2013. Lecture Notes in Computer Science, vol. 7881, pp. 160–176. Springer, Heidelberg, Germany, Athens, Greece (May 26–30, 2013). https://doi.org/10.1007/978-3-642-38348-9_10
- Jain, A., Lin, H., Sahai, A.: Indistinguishability obfuscation from well-founded assumptions. In: STOC. pp. 60–73. ACM (2021)
- Kalai, Y.T.: Smooth projective hashing and two-message oblivious transfer. In: Cramer, R. (ed.) Advances in Cryptology – EUROCRYPT 2005. Lecture Notes in Computer Science, vol. 3494, pp. 78–95. Springer, Heidelberg, Germany, Aarhus, Denmark (May 22–26, 2005). https://doi.org/10.1007/11426639_5
- Maurer, U.: Conditionally-perfect secrecy and a provably-secure randomized cipher. Jornal of Cryptology 5(1), 53–66 (Jan 1992)
- Maurer, U.M.: Protocols for secret key agreement by public discussion based on common information. In: Brickell, E.F. (ed.) Advances in Cryptology – CRYPTO'92. Lecture Notes in Computer Science, vol. 740, pp. 461–470. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 16–20, 1993). https://doi.org/ 10.1007/3-540-48071-4_32
- Maurer, U.M., Renner, R., Holenstein, C.: Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. In: Naor, M. (ed.) TCC 2004: 1st Theory of Cryptography Conference. Lecture Notes in Computer Science, vol. 2951, pp. 21–39. Springer, Heidelberg, Germany, Cambridge, MA, USA (Feb 19–21, 2004). https://doi.org/10.1007/978-3-540-24638-1_2

- Moran, T., Wichs, D.: Incompressible encodings. In: Micciancio, D., Ristenpart, T. (eds.) Advances in Cryptology – CRYPTO 2020, Part I. Lecture Notes in Computer Science, vol. 12170, pp. 494–523. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 17–21, 2020). https://doi.org/10.1007/ 978-3-030-56784-2_17
- 43. Naor, M.: On cryptographic assumptions and challenges (invited talk). In: Boneh, D. (ed.) Advances in Cryptology – CRYPTO 2003. Lecture Notes in Computer Science, vol. 2729, pp. 96–109. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 17–21, 2003). https://doi.org/10.1007/978-3-540-45146-4_6
- Paillier, P.: Public-key cryptosystems based on composite degree residuosity classes. In: Stern, J. (ed.) Advances in Cryptology – EUROCRYPT'99. Lecture Notes in Computer Science, vol. 1592, pp. 223– 238. Springer, Heidelberg, Germany, Prague, Czech Republic (May 2–6, 1999). https://doi.org/10.1007/ 3-540-48910-X_16
- Raz, R.: A time-space lower bound for a large class of learning problems. In: Umans, C. (ed.) 58th Annual Symposium on Foundations of Computer Science. pp. 732-742. IEEE Computer Society Press, Berkeley, CA, USA (Oct 15-17, 2017). https://doi.org/10.1109/F0CS.2017.73
- 46. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: Gabow, H.N., Fagin, R. (eds.) 37th Annual ACM Symposium on Theory of Computing. pp. 84–93. ACM Press, Baltimore, MA, USA (May 22–24, 2005). https://doi.org/10.1145/1060590.1060603