






Multi-instance Randomness Extraction and Security Against Bounded-Storage Mass Surveillance

Jiaxin Guan¹ , Daniel Wichs^{2,3} , and Mark Zhandry³ 

¹ Princeton University, Princeton, NJ 08544, USA

² Northeastern University, Boston, MA 02115, USA

³ NTT Research, Inc., Sunnyvale, CA 94085, USA

mzhandry@gmail.com

Abstract. Consider a state-level adversary who observes and stores large amounts of encrypted data from all users on the Internet, but does not have the capacity to store it all. Later, it may target certain “persons of interest” in order to obtain their decryption keys. We would like to guarantee that, if the adversary’s storage capacity is only (say) 1% of the total encrypted data size, then even if it can later obtain the decryption keys of arbitrary users, it can only learn something about the contents of (roughly) 1% of the ciphertexts, while the rest will maintain full security. This can be seen as an extension of *incompressible cryptography* (Dziembowski CRYPTO’06, Guan, Wichs and Zhandry EUROCRYPT’22) to the *multi-user* setting. We provide solutions in both the symmetric key and public key setting with various trade-offs in terms of computational assumptions and efficiency.

As the core technical tool, we study an information-theoretic problem which we refer to as “multi-instance randomness extraction”. Suppose X_1, \dots, X_t are correlated random variables whose total joint min-entropy rate is α , but we know nothing else about their individual entropies. We choose t random and independent seeds S_1, \dots, S_t and attempt to individually extract some small amount of randomness $Y_i = \text{Ext}(X_i; S_i)$ from each X_i . We’d like to say that roughly an α -fraction of the extracted outputs Y_i should be indistinguishable from uniform even given all the remaining extracted outputs and all the seeds. We show that this indeed holds for specific extractors based on Hadamard and Reed-Muller codes.

1 Introduction

Bounded-Storage Mass Surveillance. We consider a scenario where a powerful (e.g., state-level) adversary wants to perform mass surveillance of the population. Even if the population uses encryption to secure all communication, the adversary can collect large amounts of encrypted data from the users (e.g., by monitoring encrypted traffic on the Internet). The data is encrypted and hence

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/978-3-031-48621-0_4.

the adversary does not learn anything about its contents when it is collected. However, the adversary may store this data for the future. Later, it may identify various “persons of interest” and perform expensive targeted attacks to get their secret keys (e.g., by remote hacking or by physically compromising their devices). We will assume the adversary is capable of eventually getting any secret key of any user of its choosing. Can we still achieve any meaningful notion of security against such mass-surveillance?

One option is to rely on cryptosystems having *forward secrecy* [19], which exactly addresses the problem of maintaining security even if the secret key is later compromised. Unfortunately, forward-secure encryption schemes inherently require either multi-round interaction between the sender and receiver or for the receiver to perform key updates, both of which can be impractical or impossible in many natural scenarios. Without these, it may seem that no reasonable security is possible – if the adversary collects all the ciphertexts and later can get any secret key, clearly it can also get any plaintext!

In this work, we restrict the adversary to have *bounded storage*, which is much smaller than the total of size of all the encrypted data it can observe. This is a reasonable assumption since the total communication of an entire population is likely huge.¹ As a running example throughout the introduction, we will assume that the adversary’s storage capacity is 1% of the total encrypted data size. We allow the adversary to observe all the encrypted data simultaneously and then compress it in some arbitrary way to fit within its storage budget. Later, the adversary can get any secret key of any user of its choosing, and eventually it may even get all the keys of all the users. What kind of security guarantees can we provide in this setting?

Clearly, the adversary can simply store 1% of the ciphertexts and discard the remaining 99%, which will allow it to later compromise the security of 1% of the users by getting their secret keys. While one may pessimistically see this as a significant privacy violation already, we optimistically regard this as a potentially reasonable privacy outcome that’s vastly preferable to the adversary being able to compromise all the users. For example, if the adversary later chooses a random user and wants to learn something about their data, it will only be able to do so with 1% probability, even if it can get their secret key. But can we argue that this is the best that the adversary can do? In particular, we’d like to say that, no matter what compression strategy the adversary employs, it will be unable to learn anything about the contents of 99% of the ciphertexts, even if it later gets all the secret keys. Unfortunately, this is not generically true. For example, the adversary could store the first 1% of the bits of every ciphertext. If the encryption scheme is (e.g.,) the one-time pad, then an adversary who later learns the secret keys would later be able to learn the first 1% of every encrypted message of every user, which may provide a pretty good idea of the overall message contents. In fact, it can get even worse than this. If the encryption scheme is fully homomorphic, the adversary can individually compress each ciphertext

¹ Global annual Internet traffic has long surpassed 1 zettabyte (10^{21} bytes) [4], while total world-wide datacenter storage is only a couple zettabytes in 2022 [11].

into a small evaluated ciphertext encrypting some arbitrary predicate of the data (e.g., was the message insulting of the supreme leader), and therefore learn the outcome of this predicate about the encrypted data of every user. Even worse, if the encryption scheme is multi-key fully homomorphic, the adversary can derive a compressed ciphertext that encrypts the output of a joint computation over all the data of all the users, as long as the output is sufficiently small. Thus, in general, an adversary whose storage capacity is only 1%, may still be able to learn some partial information about the encrypted messages of a 100% of the users. The question is then, whether or not it is indeed possible to guarantee only 1% of users are compromised, and if so to actually design such a scheme.

Connection to Incompressible Cryptography. Encryption schemes that offer protection against bounded-storage mass surveillance can be seen as a generalization of *incompressible encryption* [6, 15, 17] to the setting of multiple ciphertexts. To clarify the distinction, we refer to the earlier notion of incompressible encryption as *individually incompressible* and our new notion as *multi-incompressible*.

In an *individually incompressible encryption* scheme, we can make the size of a ciphertext flexibly large, and potentially huge (e.g., many gigabytes). An adversary observes a single ciphertext, but cannot store it in its entirety and can instead only store some compressed version of it. Security dictates that even if the adversary later gets the user’s secret key, it cannot learn anything about the encrypted message. The work of [15] gave a construction of one-time symmetric-key encryption with information-theoretic security in this setting, and the work of [17] showed how to achieve public-key encryption in this setting, under the minimal assumption that standard public-key encryption exists. The works of [6, 17] also constructed such public-key encryption schemes having rate 1, meaning that the size of the message can be almost as large as the ciphertext size, and the latter work even showed how to do so under specific but standard public-key assumptions.

In our new notion of *multi-incompressible encryption*, we also have the flexibility to make the ciphertext size arbitrarily large. But now the adversary observes a large number of ciphertexts from many users and compresses them down to something that’s roughly an α -fraction of the size of all the original ciphertexts, for some α . In particular, the adversary’s storage may be much larger than a single ciphertext. Later the adversary gets all the secret keys, and we want to say that the adversary can only learn something about a (roughly) α -fraction of the messages, but cannot learn anything about the rest.

Our new notion of multi-incompressibility implies individual incompressibility. In particular, in the case of a single ciphertext, unless the adversary stores essentially all of it (i.e., $\alpha \approx 1$), it cannot learn anything about the encrypted message (= 100% of the messages). But our notion is significantly more general. For example, individual incompressibility does not even offer any guarantees if an adversary can take even 2 ciphertexts and compress them down to the size of 1, while multi-incompressibility ensures that one of the messages stays secure.

Formalizing multi-incompressibility is tricky: the natural indistinguishability-based approach would be to insist that the encryptions of two lists of messages are indistinguishable. But unlike individually incompressible encryption, in our

setting the adversary can always learn *something*, namely the messages contained in ciphertexts it chose to store. We therefore need a fine-grained notion which captures that some messages to be learned, but other messages remain completely hidden. We give details on our solution below.

Extracting Randomness Against Correlated Sources. Before getting to our results, we discuss randomness extraction, which is a central tool in all existing constructions of incompressible encryption. A randomness extractor Ext takes as input a source of imperfect randomness X and uses it to distill out some (nearly) uniformly random string Y . Here, we consider seeded extractors, which use a public uniformly random seed S as a catalyst to extract $Y = \text{Ext}(X; S)$, such that Y should be (nearly) uniform even conditioned on the seed S .

While randomness extraction is very well studied, it is most often in the *single-use* case, where a single string $Y = \text{Ext}(X; S)$ is extracted from a single source X having sufficient entropy. Here we ask: what if many strings $Y_i = \text{Ext}(X_i; S_i)$ are extracted from multiple sources X_i respectively (using independent random seeds S_i), but where the sources X_i may be arbitrarily correlated? What guarantees can be made? We consider the case where we only know that the total joint entropy of all the sources is high, but we know nothing else about their individual entropies; indeed some of the sources may have no entropy at all! In this case, clearly not all of the extracted values Y_i can be uniform, and some may even be entirely deterministic. One may nevertheless hope that *some* of the extracted values remain uniform, where the fraction of uniform values roughly correlates to combined total entropy rate of all the sources. To the best of our knowledge, randomness extraction in this setting has not been studied before.

1.1 Our Results

Formalizing Multi-user Incompressible Encryption. We first provide definitions for multi-user incompressible encryption. We depart from the indistinguishability-based definitions of the prior work on incompressible cryptography [6, 15, 17], and instead give a *simulation*-based definition. Essentially, it says that anything that an adversary can learn by taking many ciphertexts of different users, compressing them down sufficiently, and later getting all the secret keys, can be simulated by a simulator that can only ask to see some small fraction of the plaintexts but learns nothing about the remaining ones. In the single-instance case, this definition implies indistinguishability-based security, but appears stronger. Nevertheless, existing constructions and proofs are readily adapted to satisfy simulation security. The distinction becomes more important in the multi-user setting, however, where simulation security allows us to naturally define what it means for some messages to be revealed and some to remain hidden.

Multi-instance Randomness Extractors. As our main technical tool, we explore a new kind of extractor that we call a multi-instance randomness extractor, which aims to solve the extraction problem outlined above. Syntactically, this is a standard extractor $Y = \text{Ext}(X; S)$ that takes as input a source X and a seed S and outputs some short randomness Y . However, we now imagine that the extractor

is applied separately to t correlated sources X_i , with each invocation using an independent seed S_i , to derive extracted values $Y_i = \text{Ext}(X_i; S_i)$. The only guarantee on the sources is that the total joint min-entropy of $X = (X_1, \dots, X_t)$ is sufficiently high. Any individual source X_i , however, may actually be deterministic (have 0 entropy), in which case the corresponding extracted value Y_i is of course not random. However, provided the total min-entropy rate of X is high, it is guaranteed that *many* of the t extracted values are statistically-close uniform. Ideally, if the joint min-entropy rate of X is α , we would hope that roughly αt of the extracted values are uniform.

Formalizing the above requires some care. For example, it may be the case that X is chosen by selecting a random index $i^* \leftarrow [t]$, setting X_{i^*} to be all 0's, and choosing the remaining block X_j for $j \neq i^*$ uniformly at random. In that case X has a very high entropy rate, but for any fixed index i , the min-entropy of X_i is small (at most $\log t$ since with polynomial probability $1/t$ the value of X_i is all 0's), and not enough to extract even 1 bit with negligible bias. Therefore, we cannot argue that $Y_i = \text{Ext}(X_i; S_i)$ is close to uniform for any particular index i ! Instead, we allow the set of indices i , for which Y_i is close to uniform, itself be a random variable correlated with X . (See Definition 3.)

We show constructions of multi-instance randomness extractors nearing the optimal number of uniform extracted values. In particular, we show that if the joint min-entropy rate of $X = (X_1, \dots, X_t)$ is α then there exists some random variable I_X denoting a subset of $\approx \alpha \cdot t$ indices in $[t]$ such that nobody can distinguish between seeing all the extracted values $Y_i = \text{Ext}(X_i; S_i)$ versus replacing all the Y_i for $i \in I_X$ by uniform, even given all the seeds S_i . (See Corollary 1.) Our constructions are based on Hadamard codes (long seed) and Reed-Muller codes (short seed). While the constructions themselves are standard, our analysis is novel, leveraging the list-decodability of the codes, plus a property we identify called *hinting*. Hinting roughly means that the values of $\{\text{Ext}(x; S_i)\}_i$ on some particular exponentially large set of pairwise independent seeds S_i can be compressed into a single small hint, of size much smaller than x . This hinting property is a crucial feature in the *local* list-decoding algorithms for these codes, but appears not to have been separately formalized/utilized as a design goal for an extractor.²

Applications. We then show that multi-instance randomness extraction can be used essentially as a drop-in replacement for standard randomness extractors in prior constructions of individual incompressible encryption, lifting them to multi-incompressible encryption. As concrete applications, we obtain multi-incompressible encryption in a variety of settings:

² The work of [1] studied a notion of extractors for “Somewhere Honest Entropic Look Ahead” (SHELA) sources. The notions are largely different and unrelated. In particular: (i) in our work X is an arbitrary source of sufficient entropy while [1] places additional restrictions, (ii) we use a seeded extractor while [1] wants a deterministic extractor, (iii) we apply the seeded extractor separately on each X_i while [1] applies it jointly on the entire X , (iv) we guarantee that a large fraction of extracted outputs is uniform even if the adversary sees the rest, while in [1] the adversary cannot see the rest.

- A symmetric key scheme with information-theoretic security, by replacing the extractor in [15].
- A “rate-1” symmetric key scheme, meaning the ciphertext is only slightly larger than the message. Here, we assume either decisional composite residuosity (DCR) or learning with errors (LWE), matching [6]³.
- A public key scheme, assuming any ordinary public key encryption scheme, matching [17].
- A rate-1 public key scheme, under the same assumptions as [6]⁴. The scheme has large public keys.
- A rate-1 public key scheme that additionally has succinct public keys, assuming general functional encryption, matching [17].

In all cases, we guarantee that if the adversary’s storage is an α fraction of the total size of all the ciphertexts, then it can only learn something about a $\beta \approx \alpha$ fraction of the encrypted messages. We can make $\beta = \alpha - 1/p(\lambda)$ for any polynomial p in the security parameter λ , by choosing a sufficiently large ciphertext size.

Multiple Ciphertexts Per User. Prior work, in addition to only considering a single user, also only considers a single ciphertext per user. Perhaps surprisingly, security does not compose, and indeed for any fixed secret key size, we explain that simulation security for unbounded messages is impossible.

We therefore develop schemes for achieving a bounded number of ciphertexts per user. We show how to modify each of the constructions above to achieve multi-ciphertext security under the same assumptions.

The Random Oracle Model. In the full version [18] of the paper, we also show how to construct symmetric key multi-user incompressible encryption with an unbounded number of ciphertexts per user and also essentially optimal secret key and ciphertext sizes, from random oracles. This shows that public key tools are potentially not inherent to rate-1 symmetric incompressible encryption.

1.2 Concurrent Work

A concurrent and independent work of Dinur et al. [12] (Sect. 6.2) considers an extraction problem that turns out to be equivalent to our notion of *Multi-Instance Randomness Extractor*. They study this problem in a completely different context of differential-privacy lower bounds. They show that (in our language) universal hash functions are “multi-instance randomness extractors” with good parameters, similar to the ones in our work. While conceptually similar, the results are technically incomparable since we show our result for hinting

³ One subtlety is that, for all of our rate-1 constructions, we need a PRG secure against *non-uniform* adversaries, whereas the prior work could have used a PRG against uniform adversaries.

⁴ [6] explores CCA security, but in this work for simplicity we focus only on CPA security.

extractors while they show it for universal hash functions. One advantage of our result is that we show how to construct hinting extractors with short seeds, while universal hash functions inherently require a long seed. Their proof is completely different from the one in our paper.

The fact that multi-instance randomness extractors have applications in different contexts, as demonstrated in our work and Dinur et al. [12], further justifies them as a fundamental primitive of independent interest. We believe that having two completely different techniques/approaches to this problem is both interesting and valuable.

1.3 Our Techniques: Multi-instance Randomness Extraction

We discuss how to construct a multi-instance randomness extractor Ext . Recall, we want to show that, if the joint min-entropy rate of $X = (X_1, \dots, X_t)$ is α then there exists some random variable I_X denoting a subset of $\approx \alpha \cdot t$ indices in $[t]$ such that the distribution $(S_i, Y_i = \text{Ext}(X_i; S_i))_{i \in [t]}$ is statistically indistinguishable from $(S_i, Z_i)_{i \in [t]}$ where Z_i is uniformly random for $i \in I_X$ and $Z_i = Y_i$ otherwise.

A Failed Approach. A natural approach would be to try to show that every standard seeded extractor is also a “multi-instance randomness extractor”. As a first step, we would show that there is some random variable I_X denoting a large subset of $[t]$ such that the values X_i for $i \in I_X$ have large min-entropy conditioned on $i \in I_X$. Indeed, such results are known; see for example the “block-entropy lemma” of [13] (also [9, 16]). In fact, one can even show a slightly stronger statement that the random variables X_i for $i \in I_X$ have high min-entropy even conditioned on all past blocks X_1, \dots, X_{i-1} . However, it cannot be true that X_i has high min-entropy conditioned on *all* other blocks past and future (for example, think of X being uniform subject to $\bigoplus_{i=1}^t X_i = 0$). Unfortunately, this prevents us for using the block-entropy lemma to analyze multi-instance extraction, where the adversary sees some extracted outputs from all the blocks.⁵ It remains as a fascinating open problem whether every standard seeded extractor is also a multi-instance randomness extractor or if there is some counterexample.⁶

Our Approach. We are able to show that particular seeded extractors Ext based on Hadamard or Reed-Muller codes are good multi-instance randomness

⁵ This strategy would allow us to only prove a very weak version of multi-instance extraction when the number of blocks t is sufficiently small. In this case we can afford to lose the t extracted output bits from the entropy of *each* block. However, in our setting, we think of the number of blocks t as huge, much larger than the size/entropy of each individual block.

⁶ We were initially convinced that the general result does hold and invested much effort trying to prove it via some variant of the above approach without success. We also mentioned the problem to several experts in the field who had a similar initial reaction, but were not able to come up with a proof.

extractors. For concreteness, let us consider the Hadamard extractor $\text{Ext}(x; s) = \langle x, s \rangle$.⁷ Our proof proceeds in 3 steps:

Step 1: Switch quantifiers. We need to show that there *exists* some random variable I_X such that *every* statistical distinguisher fails to distinguish between the two distributions $(S_i, Y_i)_{i \in [t]}$ and $(S_i, Z_i)_{i \in [t]}$. We can use von Neumann's minimax theorem to switch the order quantifiers.⁸ Therefore, it suffices to show that for every (randomized) statistical distinguisher D there is some random variable I_X such that D fails to distinguish the above distributions.

Step 2: Define I_X . For any fixed $x = (x_1, \dots, x_t)$ we define the set I_x to consist of indices $i \in [t]$ such that D fails to distinguish between the hybrid distributions $(\{S_j\}_{j \in [t]}, Z_1, \dots, Z_{i-1}, Y_i, \dots, Y_t)$ versus $(\{S_j\}_{j \in [t]}, Z_1, \dots, Z_i, Y_{i+1}, \dots, Y_t)$, where in both distributions we condition on $X = x$. In other words, these are the indices where we can replace the next extracted output by random and fool the distinguisher. We then define the random variable I_X that chooses the correct set I_x according to X . It is easy to show via a simple hybrid argument that with this definition of I_X it is indeed true that D fails to distinguish $(S_i, Y_i)_{i \in [t]}$ and $(S_i, Z_i)_{i \in [t]}$.

Step 3: Argue that I_X is large. We still need to show that I_X is a large set, containing $\approx \alpha \cdot t$ indices. To do so, we show that if I_X were small (with non negligible probability) then we could “guess” X with sufficiently high probability that would contradict X having high min-entropy. In particular, we provide a guessing strategy such that for any x for which I_x is small, our strategy has a sufficiently high chance of guessing x . First, we guess the small set $I_x \subseteq [t]$ as well as all of the blocks x_i for $i \in I_x$ uniformly at random. For the rest of the blocks $i \notin I_x$, we come up with a guessing strategy that does significantly better than guessing randomly. We rely on the fact that distinguishing implies predicting, to convert the distinguisher D into a predictor P such that for all $i \notin I_x$ we have: $P(S_i, \{S_j, \text{Ext}(x_j; S_j)\}_{j \in [t] \setminus \{i\}}) = \text{Ext}(x_i; S_i)$ with probability significantly better than $1/2$. Now we would like to use the fact that the Hadamard code $(\text{Ext}(x; s) = \langle x, s \rangle)_s$ is list-decodable to argue that we can use such predictor P to derive a small list of possibilities for x . Unfortunately, there is a problem with this argument. To call the predictor, the predictor requires an auxiliary input, namely $\text{aux}_i = \{S_j, \text{Ext}(x_j; S_j)\}_{j \in [t] \setminus \{i\}}$. Supplying the aux_i in turn requires knowing at least t bits about x . We could hope to guess a good choice of aux_i , but there may be a different good choice for each $i \in [t]$, and therefore we would need to guess a fresh t bits of information about x just to recover each block x_i , which when $|x_i| < t$ is worse than the trivial approach of guessing x_i directly! Instead, we use a trick inspired by the proof of the Goldreich-Levin theorem. For each

⁷ For the sake of exposition, here we only show the case where the extractor output is a single bit. In Sect. 3, we construct extractors with multiple-bit outputs.

⁸ Think of the above as a 2 player game where one player chooses I_X , the other chooses the distinguisher and the payout is the distinguishing advantage; the minimax theorem says that the value of the game is the same no matter which order the players go in.

block $j \in [t]$, we guess the values of $b^{(k)} := \langle x_j, S_j^{(k)} \rangle$ for a very small “base set” of h random seeds $S_j^{(1)}, \dots, S_j^{(h)}$. We can then expand this small “base set” of seeds into an exponentially larger “expanded set” of $Q = 2^h - 1$ seeds $S_j^{(K)} := \sum_{k \in K} S_j^{(k)}$ for $K \subseteq [h] \setminus \emptyset$, and derive guesses for $b^{(K)} := \langle x_j, S_j^{(K)} \rangle$ by setting $b^{(K)} = \sum_{k \in K} b^{(k)}$. By linearity, the expanded set of guesses is correct if the base set is correct, and moreover the expanded sets of seeds $(S_j^{(K)})_K$ are pairwise independent for different sets K . Therefore, for each set K , we can derive the corresponding $\text{aux}_i^{(K)}$. We can now apply Chebyshev’s bound to argue that if for each i we take the majority value for $P(S_i, \text{aux}_i^{(K)})$ across all Q sets K , it is likely equal to $\text{Ext}(x_i; S_i)$ with probability significantly better than $1/2$. Notice that we got our saving by only guessing ht bits about $x = (x_1, \dots, x_t)$ for some small value h (roughly $\log(1/\varepsilon)$ if we want indistinguishability ε) and were able to use these guesses to recover all the blocks x_i for $i \notin I_x$.

Generalizing. We generalize the above analysis for the Hadamard extractor to any extractor that is list-decodable and has a “hinting” property as discussed above. In particular, this also allows us to use a Reed-Muller based extractor construction with a much smaller seed and longer output length.

1.4 Our Techniques: Multi-incompressible Encryption

We then move to considering incompressible encryption in the multi-user setting.

Definition. We propose a simulation-based security definition for multi-instance incompressible encryption. Roughly, the simulator first needs to simulate all the ciphertexts for all the instances *without* seeing any of the message queries, corresponding to the fact that at this point the adversary can’t learn anything about any of the messages. To model the adversary then learning the secret keys, we add a second phase where the simulator can query for a *subset* of the messages, and then must simulate *all* the private keys. We require that no *space-bounded* distinguisher can distinguish between the receiving real encryptions/real private keys vs receiving simulated encryptions/keys. The number of messages the simulator can query will be related to the storage bound of the distinguisher.

Upgrading to Multi-incompressible Encryption Using Multi-instance Randomness Extraction. All prior standard-model constructions of individual incompressible encryption [6, 15, 17] utilize a randomness extractor. For example, Dziembowski [15] gives the following simple construction of a symmetric key incompressible encryption scheme:

- The secret key k is parsed as (s, k') where s is a seed for a randomness extractor, and k' is another random key.
- To encrypt a message m , choose a large random string R , and output $c = (R, d = \text{Ext}(R; s) \oplus k' \oplus m)$.

The intuition for (individual) incompressible security is that an adversary that cannot store essentially all of c can in particular not store all of R , meaning

R has min-entropy conditioned on the adversary’s state. The extraction guarantee then shows that $\text{Ext}(R; s)$ can be replaced with a random string, thus masking the message m .

We demonstrate that our multi-instance randomness extractors can be used as a drop-in replacement for ordinary random extractors in all prior constructions of individual incompressible encryption, upgrading them to multi-incompressible encryption. In the case of [15], this is almost an immediate consequence of our multi-instance randomness extractor definition. Our simulator works by first choosing random s for each user, and sets the ciphertexts of each user to be random strings. Then it obtains from the multi-instance randomness extractor guarantee the set of indices i where Y_i is close to uniform. For these indices, it sets k' to be a uniform random string. This correctly simulates the secret keys for these i .

For i where Y_i is not uniform, the simulator then queries for messages for these i . It programs k' as $k' = d \oplus \text{Ext}(R; s) \oplus m$; decryption under such k' will correctly yield m . Thus, we correctly simulate the view of the adversary, demonstrating multi-incompressible security.

Remark 1. The set of indices where Y_i is uniform will in general not be efficiently computable, and multi-instance randomness extraction only implies that the set of indices exist. Since our simulator must know these indices, our simulator is therefore inefficient. In general, an inefficient simulator seems inherent in the standard model, since the adversary’s state could be scrambled in a way that hides which ciphertexts it is storing.

We proceed to show that various constructions from [6, 17] are also secure in the multi-user setting, when plugging in multi-instance randomness extractors. In all cases, the proof is essentially identical to the original single-user counterpart, except that the crucial step involving extraction is replaced with the multi-instance randomness extraction guarantee. We thus obtain a variety of parameter size/security assumption trade-offs, essentially matching what is known for the single-user setting.

One small issue that comes up is that, once we have invoked the multi-instance randomness extractor, the simulation is inefficient. This presents a problem in some of the security proofs, specifically in the “rate-1” setting where messages can be almost as large as ciphertexts. In the existing proofs in this setting, there is a computational hybrid step that comes *after* applying the extractor. Naively, this hybrid step would seem to be invalid since the reduction now has to be inefficient. We show, however, that the reduction can be made efficient as long as it is *non-uniform*, essentially having the choice of indices (and maybe some other quantities) provided as non-uniform advice. As long as the underlying primitive for these post-extraction hybrids has non-uniform security, the security proof follows.

Multiple Ciphertexts Per User. We also consider the setting where there may be multiple ciphertexts per user, which has not been considered previously.

It is not hard to see that having an *unbounded* number of ciphertexts per user is impossible in the standard model. This is because the simulator has to simulate everything but the secret key without knowing the message. Then, for the ciphertexts stored by the adversary, the simulator queries for the underlying messages and must generate the secret key so that those ciphertexts decrypt to the given messages. By incompressibility, this means the secret key length must be at least as large as the number of messages.

We instead consider the case of bounded ciphertexts per user. For a stateful encryption scheme, it is trivial to upgrade a scheme supporting one ciphertext per user into one supporting many: simply have the secret key be a list of one-time secret keys. In the symmetric key setting, this can be made stateless by utilizing k -wise independent hash functions.

In the public key setting, achieving a stateless construction requires more work, and we do not believe there is a simple generic construction. We show instead how to modify all the existing constructions to achieve multiple ciphertexts per user. Along the way, we show an interesting combinatorial approach to generically lifting non-committing encryption to the many-time setting without sacrificing ciphertext rate.

2 Preliminaries

Notation-wise, for $n \in \mathbb{N}$, we let $[n]$ denote the ordered set $\{1, 2, \dots, n\}$. We use capital bold letters to denote a matrix \mathbf{M} . Lowercase bold letters denote vectors \mathbf{v} . Let $\mathbf{M}_{i,j}$ denote the element on the i -th row, and j -th column of \mathbf{M} , and \mathbf{v}_i denote the i -th element of \mathbf{v} .

Lemma 1 (Johnson Bound, Theorem 3.1 of [20]). *Let $\mathcal{C} \subseteq \Sigma^n$ with $|\Sigma| = q$ be any q -ary error-correcting code with relative distance $p_0 = 1 - (1 + \rho)^{\frac{1}{q}}$ for $\rho > 0$, meaning that for any two distinct values $x, y \in \mathcal{C}$, the Hamming distance between x, y is at least $p_0 \cdot n$. Then for any $\delta > \sqrt{\rho(q-1)}$ there exists some $L \leq \frac{(q-1)^2}{\delta^2 - \rho(q-1)}$ such that the code is $(p_1 = (1 - (1 + \delta)^{\frac{1}{q}}), L)$ -list decodable, meaning that for any $y \in \Sigma_q^n$ there exist at most L codewords $x \in \mathcal{C}$ that are within Hamming distance $p_1 n$ of y .*

Lemma 2 (Distinguishing implies Predicting). *For any randomized function $D : \{0, 1\}^n \times \{0, 1\}^m \rightarrow \{0, 1\}$ there exists some randomized function $P : \{0, 1\}^n \rightarrow \{0, 1\}^m$ such that for any jointly distributed random variables (A, B) over $\{0, 1\}^n \times \{0, 1\}^m$:*

if $\Pr[D(A, B) = 1] - \Pr[D(A, U_m) = 1] \geq \varepsilon$ then $\Pr[P(A) = B] \geq \frac{1}{2^m}(1 + \varepsilon)$.

Proof. Define $P(a)$ as follows. Sample a random $b_0 \leftarrow \{0, 1\}^m$, if $D(a, b_0) = 1$ output b_0 else sample a fresh $b_1 \leftarrow \{0, 1\}^m$ and output b_1 .

Define $p = \Pr[D(A, U_m) = 1]$. Let B_0, B_1 be independent random variables that are uniform over $\{0, 1\}^m$ corresponding to the strings b_0, b_1 . Then we have

$$\begin{aligned} \Pr[P(A) = B] &= \Pr[D(A, B_0) = 1 \wedge B_0 = B] + \Pr[D(A, B_0) = 0 \wedge B_1 = B] \\ &= \Pr[B_0 = B] \Pr[D(A, B) = 1] + \Pr[D(A, B_0) = 0] \Pr[B_1 = B] \\ &= \frac{1}{2^m}(\varepsilon + p) + (1 - p) \frac{1}{2^m} = \frac{1}{2^m}(1 + \varepsilon). \end{aligned}$$

□

Min-Entropy Extractor. Recall the definition for average min-entropy:

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

$$H_\infty(X|Y) = -\log \mathbf{E}_{y \leftarrow Y} [\max_x \Pr[X = x|Y = y]].$$

Lemma 3 ([14]). For random variables X, Y where Y is supported over a set of size T , we have $H_\infty(X|Y) \geq H_\infty(X, Y) - \log T \geq H_\infty(X) - \log T$.

Definition 2 (Extractor [23]). A function $\text{Extract} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a (k, ϵ) strong average min-entropy extractor if, for all jointly distributed random variables (X, Y) where X takes values in $\{0, 1\}^n$ and $H_\infty(X|Y) \geq k$, we have that $(U_d, \text{Extract}(X; U_d), Y)$ is ϵ -close to (s, U_m, Y) , where U_d and U_m are uniformly random strings of length d and m respectively.

Remark 2. Any strong randomness extractor is also a strong average min-entropy extractor, with a constant loss in ϵ .

Definitions of incompressible encryption and functional encryption can be found in the full version [18] of the paper.

3 Multi-instance Randomness Extraction

3.1 Defining Multi-instance Extraction

Definition 3 (Multi-instance Randomness Extraction). A function $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is $(t, \alpha, \beta, \varepsilon)$ -multi-instance extracting if the following holds. Let $X = (X_1, \dots, X_t)$ be any random variable consisting of blocks $X_i \in \{0, 1\}^n$ such that $H_\infty(X) \geq \alpha \cdot tn$. Then, there exists some random variable I_X jointly distributed with X , such that I_X is supported over sets $\mathcal{I} \subseteq [t]$ of size $|\mathcal{I}| \geq \beta \cdot t$ and:

$$(S_1, \dots, S_t, \text{Ext}(X_1; S_1), \dots, \text{Ext}(X_t; S_t)) \approx_\varepsilon (S_1, \dots, S_t, Z_1, \dots, Z_t)$$

where $S_i \in \{0, 1\}^d$ are uniformly random and independent seeds, and $Z_i \in \{0, 1\}^m$ is sampled independently and uniformly random for $i \in I_X$ while $Z_i = \text{Ext}(X_i; S_i)$ for $i \notin I_X$.

In other words, the above definition says that if we use a “multi-instance extracting” extractor with independent seeds to individually extract from t correlated blocks that have a joint entropy-rate of α , then seeing all the extracted outputs is indistinguishable from replacing some carefully chosen β -fraction by uniform.

3.2 Hinting Extractors

Definition 4 (Hinting Extractor). A function $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a (δ, L, h, Q) -hinting extractor if it satisfies the following:

- *List Decodable:* If we think of $\text{ECC}(x) = (\text{Ext}(x; s))_{s \in \{0, 1\}^d}$ as a $(2^d, n)_{\Sigma = \{0, 1\}^m}$ error-correcting code over the alphabet $\Sigma = \{0, 1\}^m$, then the code is $(p = 1 - (1 + \delta)2^{-m}, L)$ -list decodable, meaning that for any $y \in \Sigma^{2^d}$, the number of codewords that are within Hamming distance $p \cdot 2^d$ of y is at most L .
- *Pairwise-Independent Hint:* There exist some functions $\text{hint} : \{0, 1\}^n \times \{0, 1\}^\tau \rightarrow \{0, 1\}^h$, along with rec_0 and rec_1 such that:
 - For all $x \in \{0, 1\}^n, r \in \{0, 1\}^\tau$, if we define $\sigma = \text{hint}(x; r)$, $\{s_1, \dots, s_Q\} = \text{rec}_0(r)$, and $\{y_1, \dots, y_Q\} = \text{rec}_1(\sigma, r)$, then $\text{Ext}(x; s_i) = y_i$ for all $i \in [Q]$.
 - Over a uniformly random $r \leftarrow \{0, 1\}^\tau$, the Q seeds $\{s_1, \dots, s_Q\} = \text{rec}_0(r)$, are individually uniform over $\{0, 1\}^d$ and pairwise independent.

Intuitively, the pairwise-independent hint property says that there is a small (size h) hint about x that allows us to compute $\text{Ext}(x; s_i)$ for a large (size Q) set of pairwise independent seeds s_i . We generally want Q to be exponential in h .

The list-decoding property, on the other hand, is closely related to the standard definition of strong randomness extractors. Namely, if Ext is a (k, ε) -extractor then it is also $(p = 1 - (1 + \delta)2^{-m}, 2^k)$ -list decodable for $\delta = \varepsilon \cdot 2^m$, and conversely, if it is $(p = 1 - (1 + \delta)2^{-m}, 2^k)$ -list decodable then it is a $(k + m + \log(1/\delta), \delta)$ -extractor (see Proposition 6.25 in [26]).

Construction 1: Hadamard. Define $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ via $\text{Ext}(x; s) = \langle x, s \rangle$, where we interpret x, s as elements of $\mathbb{F}_{2^m}^{\hat{n}}$ for $\hat{n} := n/m$ and all the operations are over \mathbb{F}_{2^m} . The seed length is $d = n$ bits and the output length is m bits.

Lemma 4. The above $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ is a (δ, L, h, Q) -hinting extractor for any $h, \delta > 0$ with $Q \geq 2^{h-m}$ and $L \leq 2^{2m}/\delta^2$.

Proof. The list-decoding bounds on δ, L come from the Johnson bound (Lemma 1) with $q = 2^m, \rho = 0$. For pairwise-independent hints, let $\hat{h} = h/m$ and define $\text{hint}(x; R)$ to parse $R \in \mathbb{F}_{2^m}^{\hat{h} \times \hat{n}}$ and output $\sigma = R \cdot x^\top$, which has bit-size h . Let $\mathcal{V} \subseteq \mathbb{F}_{2^m}^{\hat{h}}$ be a set of vectors such that any two distinct vectors $v_1 \neq v_2 \in \mathcal{V}$ are linearly independent. Such a set \mathcal{V} exists of size $Q = (2^m)^{\hat{h}-1} + (2^m)^{\hat{h}-2} + \dots + 2^m + 1 \geq 2^{h-m}$, e.g., by letting \mathcal{V} be the set of all non-zero vectors whose left-most non-zero entry is a 1. Define $\text{rec}_0(R)$ so that it outputs $\{s_v = v \cdot R\}_{v \in \mathcal{V}}$. Correspondingly, $\text{rec}_1(\sigma, R)$ outputs $\{y_v = \langle v, \sigma \rangle\}_{v \in \mathcal{V}}$. It's easy to see that the seeds s_v are individually uniform and pairwise independent, since for any linearly-independent $v_1 \neq v_2 \in \mathcal{V}$ and the value $s_{v_1} = v_1 R$ and $s_{v_2} = v_2 R$ are random and independent over a random choice of the matrix R . Moreover for all seeds s_v we have

$$\text{Ext}(x, s_v) = \langle s_v, x \rangle = v \cdot R \cdot x^\top = \langle v, \sigma \rangle = y_v.$$

□

Construction 2: Hadamard \circ Reed-Muller. Define $\text{Ext}(f; s = (s_1, s_2)) = \langle f(s_1), s_2 \rangle$, where $f \in \mathbb{F}_{2^w}^{\binom{\ell+g}{g}}$ is interpreted as a ℓ -variate polynomial of total degree g over some field of size $2^w > g$, and $s_1 \in \mathbb{F}_{2^w}^\ell$ is interpreted as an input to the polynomial (this is Reed-Muller).⁹ Then $y = f(s_1)$ and s_2 are interpreted as a values in $\mathbb{F}_{2^m}^{w/m}$ and the inner-product $\langle y, s_2 \rangle$ is computed over \mathbb{F}_{2^m} (this is Hadamard). So overall, in bits, the input length is $n = w \cdot \binom{\ell+g}{g}$, the seed length is $d = w(\ell + 1)$ and the output length is m . This code has relative distance $1 - (\frac{1}{2^m} + \frac{g}{2^w}) = 1 - \frac{1}{2^m}(1 + \frac{g}{2^{w-m}})$.

Lemma 5. *For any w, ℓ, g, m, δ such that $2^w > g$ and m divides w , if we set $n = w \cdot \binom{\ell+g}{g}$, $d = w(\ell + 1)$ then the above $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is an (δ, L, h, Q) -hinting extractor with $\delta = \sqrt{g2^{2m}/2^w}$, $L = \frac{2^{2m}}{\delta^2 - g2^{2m}/2^w}$, $h = w \cdot (g + 1)$, $Q = 2^w$.*

In particular, for any n, m, w such that m divides w , we can set $\ell = g = \log n$ to get an (δ, L, h, Q) -hinting extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with $d = O(w \log n)$, $\delta = 2^{m+\log \log n - w/2}$, $h = O(w \log n)$ and $Q = 2^w$.

Proof. The list-decoding bounds on δ, L come from the Johnson bound (Lemma 1) with $q = 2^m$, $\rho = \frac{g}{2^{w-m}}$. On the other hand, for pairwise-independent hints, we can define $\text{hint}(f; r)$ as follows. Parse $r = (r^0, r^1, s_1^1, \dots, s_1^Q)$ with $r^0, r^1 \in \mathbb{F}_{2^w}^\ell$ and $s_1^i \in \mathbb{F}_{2^m}^{w/m}$. Let $\hat{f}(i) = f(r^0 + i \cdot r^1)$ be a univariate polynomial of degree g and define the hint $\sigma = \hat{f}$ to be the description of this polynomial. Define $\{s_i = (s_0^i, s_1^i)\} = \text{rec}_0(r)$ for $i \in \mathbb{F}_{2^w}$ by setting $s_0^i = r^0 + i \cdot r^1$. Define $\{y_i\} = \text{rec}_1(\sigma, r)$ via $y_i = \langle \hat{f}(i), s_1^i \rangle$. It is easy to check correctness and pairwise independence follows from the fact that the values $s_0^i = r^0 + i \cdot r^1$ are pairwise independent over the randomness r^0, r^1 . \square

3.3 Hinting-Extractors Are Multi-instance-Extracting

Lemma 6 (Multi-instance-Extraction Lemma). *Let $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ be a (δ, L, h, Q) -hinting extractor. Then, for any $t, \alpha > 0$ such that $Q \geq 2t \frac{2^{2m}}{\delta^2}$, it is also $(t, \alpha, \beta, \varepsilon)$ -multi-instance extracting with $\varepsilon = 6t\delta$ and $\beta = \alpha - \frac{\log L + h + \log t + \log(1/\varepsilon) + 3}{n}$.*

Proof. Our proof follows a sequence of steps.

Step 0: Relax the Size Requirement. We modify the statement of the lemma as follows. Instead of requiring that $|I_X| \geq \beta \cdot t$ holds with probability 1, we relax this to requiring that $\Pr[|I_X| < \beta \cdot t] \leq \varepsilon/4$. On the other hand, we strengthen the requirement on statistical indistinguishability from ε to $\varepsilon/2$:

$$(S_1, \dots, S_t, \text{Ext}(X_1; S_1), \dots, \text{Ext}(X_t; S_t)) \approx_{\varepsilon/2} (S_1, \dots, S_t, Z_1, \dots, Z_t).$$

This modified variant of the lemma implies the original.

⁹ Since the the input to the extractor is interpreted as a polynomial, we will denote it by f rather than the usual x to simplify notation.

To see this, notice that we can replace the set I_X that satisfies the modified variant with I'_X which is defined as $I'_X := I_X$ when $|I_X| \geq \beta t$ and $I'_X := \{1, \dots, \beta t\}$ else. The set I'_X then satisfies the original variant. In particular, we can prove the indistinguishability guarantee of the original lemma via a hybrid argument: replace I'_X by I_X ($\varepsilon/4$ statistical distance), switch from the left distribution to right distribution ($\varepsilon/2$ statistical distance), replace I_X back by I'_X ($\varepsilon/4$ statistical distance) for a total distance of ε .

Step 1: Change quantifiers. We need to prove that: *for all X with $H_\infty(X) \geq \alpha \cdot tn$, there **exists** some random variable $I_X \subseteq [t]$ with $\Pr[|I_X| < \beta t] \leq \varepsilon/4$ such that **for all** (inefficient) distinguishers D :*

$$|\Pr[D(S_1, \dots, S_t, Y_1, \dots, Y_t) = 1] - \Pr[D(S_1, \dots, S_t, Z_1, \dots, Z_t) = 1]| \leq \varepsilon/2 \quad (1)$$

where we define $Y_i = \text{Ext}(X_i; S_i)$, and the random variables Z_i are defined as in the Lemma. By the min-max theorem, we can switch the order of the last two quantifiers. In particular, it suffices to prove that: *for all X with $H_\infty(X) \geq \alpha \cdot tn$ and **for all** (inefficient, randomized) distinguishers D there **exists** some random variable $I_X \subseteq [t]$ with $\Pr[|I_X| < \beta t] \leq \varepsilon/4$ such that Eq. (1) holds.*

We can apply min-max because a distribution over inefficient distinguishers D is the same as a single randomized inefficient distinguisher D and a distribution over random variables I_X is the same as a single random variable I_X .

Step 2: Define I_X . Fix a (inefficient/randomized) distinguisher D .

For any fixed value $x \in \{0, 1\}^{n \cdot t}$, we define a set $I_x \subseteq [t]$ iteratively as follows. Start with $I_x := \emptyset$. For $i = 1, \dots, t$ add i to I_x if

$$\left(\begin{array}{l} \Pr[D(S_1, \dots, S_t, Z_1^x, \dots, Z_{i-1}^x, Y_i^x, Y_{i+1}^x, \dots, Y_t^x) = 1] \\ - \Pr[D(S_1, \dots, S_t, Z_1^x, \dots, Z_{i-1}^x, U_m, Y_{i+1}^x, \dots, Y_t^x) = 1] \end{array} \right) \leq 3\delta \quad (2)$$

where S_i is uniform over $\{0, 1\}^d$, $Y_j^x = \text{Ext}(x_j; S_j)$ and for $j < i$ we define Z_j^x to be uniformly random over $\{0, 1\}^m$ for $j \in I_x$, while $Z_j^x = Y_j^x$ for $j \notin I_x$. Note that $Y_i^x = (Y_i | X = x)$ and $Z_i^x = (Z_i | X = x)$.

Define I_X to be the random variable over the above sets I_x where x is chosen according to X . With the above definition, Eq. 1 holds since:

$$\begin{aligned} & \Pr[D(S_1, \dots, S_t, Y_1, \dots, Y_t) = 1] - \Pr[D(S_1, \dots, S_t, Z_1, \dots, Z_t) = 1] \\ &= \mathbb{E}_{x \leftarrow X} \Pr[D(S_1, \dots, S_t, Y_1, \dots, Y_t) = 1 | X = x] \\ &\quad - \Pr[[D(S_1, \dots, S_t, Z_1, \dots, Z_t) = 1 | X = x] \\ &= \mathbb{E}_{x \leftarrow X} \Pr[D(S_1, \dots, S_t, Y_1^x, \dots, Y_t^x) = 1] - \Pr[D(S_1, \dots, S_t, Z_1^x, \dots, Z_t^x) = 1] \\ &= \mathbb{E}_{x \leftarrow X} \sum_{i \in [t]} \underbrace{\left(\begin{array}{l} \Pr[D(S_1, \dots, S_t, Z_1^x, \dots, Z_{i-1}^x, Y_i^x, Y_{i+1}^x, \dots, Y_t^x) = 1] \\ - \Pr[D(S_1, \dots, S_t, Z_1^x, \dots, Z_{i-1}^x, Z_i^x, Y_{i+1}^x, \dots, Y_t^x) = 1] \end{array} \right)}_{(*)} \\ &\leq 3t\delta = \varepsilon/2 \end{aligned}$$

The last line follows since, for any x and any $i \in [t]$, if $i \notin I_x$ then $Y_i^x = Z_i^x$ and therefore $(*) = 0$, and if $i \in I_x$ then $(*) \leq 3\delta$ by the way we defined I_x in Eq. (2).

Step 3: Argue I_X is large. We are left to show that

$$\Pr[|I_X| < \beta \cdot t] \leq \varepsilon/4. \quad (3)$$

We do this via a proof by contradiction. Assume otherwise that (3) does not hold. Then we show that we can guess X with high probability, which contradicts the fact that X has high min-entropy. In particular, we define a randomized function $\text{guess}()$ such that, for any x for which $|I_x| < \beta \cdot t$, we have:

$$\Pr_{\hat{x} \leftarrow \text{guess}()}[\hat{x} = x] \geq \frac{1}{4} (t^{\beta t+1} 2^{ht} L^t 2^{\beta tn})^{-1}. \quad (4)$$

Then, assuming (3) does not hold, we have

$$\begin{aligned} \Pr_{\hat{x} \leftarrow \text{guess}(), x \leftarrow X}[\hat{x} = x] &\geq \Pr_{x \leftarrow X}[|I_x| < \beta t] \Pr_{\hat{x} \leftarrow \text{guess}(), x \leftarrow X}[\hat{x} = x \mid |I_x| < \beta t] \\ &\geq \frac{\varepsilon}{16} (t^{\beta t+1} 2^{ht} L^t 2^{\beta tn})^{-1}. \end{aligned}$$

which contradicts $H_\infty(X) \geq \alpha tn$.

Before defining the function $\text{guess}()$, we note that by the definition of I_x in Eq. (2) and the “distinguishing implies predicting” lemma (Lemma 2), there exist some predictors P_i (depending only on D), such that, for all $x \in \{0, 1\}^n$ and $i \notin I_x$, we have:

$$\Pr[P_i(S_1, \dots, S_t, Z_1^x, \dots, Z_{i-1}^x, Y_{i+1}^x, \dots, Y_t^x) = Y_i^x] \geq \frac{1}{2^m} (1 + 3\delta) \quad (5)$$

The guessing strategy. We define $\text{guess}()$ using these predictors P_i as follows:

1. Sample values r_1, \dots, r_t with $r_i \leftarrow \{0, 1\}^\tau$.
2. Sample a set $\hat{I}_x \subseteq [t]$ of size $|\hat{I}_x| \leq \beta t$ uniformly at random.
3. Sample values $\hat{\sigma}_i \leftarrow \{0, 1\}^h$ for $i \notin \hat{I}_x$ uniformly at random.
4. Sample values $\hat{x}_i \leftarrow \{0, 1\}^n$ for $i \in \hat{I}_x$ uniformly at random.
5. Let $\{s_i^1, \dots, s_i^Q\} = \text{rec}_0(r_i)$, and $\{y_i^1, \dots, y_i^Q\} = \text{rec}_1(\hat{\sigma}_i, r_i)$.
6. Use all of the above values to define, for each $i \notin \hat{I}_x$, a randomized function $\hat{P}_i(s)$ which chooses a random $j^* \leftarrow [Q]$ and outputs:

$$\hat{P}_i(s) = P_i(s_1^{j^*}, \dots, s_{i-1}^{j^*}, s, s_{i+1}^{j^*}, \dots, s_t^{j^*}, z_1^{j^*}, \dots, z_{i-1}^{j^*}, y_{i+1}^{j^*}, \dots, y_t^{j^*})$$

where $z_i^{j^*} := y_i^{j^*}$ if $i \notin \hat{I}_x$ and $z_i^{j^*} \leftarrow \{0, 1\}^m$ if $i \in \hat{I}_x$.

7. For each $i \notin \hat{I}_x$, define $\text{cw}_i \in \Sigma^{2^d}$ by setting $\text{cw}_i[s] \leftarrow \hat{P}_i(s)$, where $\Sigma = \{0, 1\}^m$. Let \mathcal{X}_i be the list of at most L values \tilde{x}_i such that the Hamming distance between $\text{ECC}(\tilde{x}_i)$ and cw_i is at most $(1 + \delta)2^d$, as in Definition 4.

8. For each $i \notin \hat{I}_x$, sample $\hat{x}_i \leftarrow \mathcal{X}_i$.
9. Output $\hat{x} = (\hat{x}_1, \dots, \hat{x}_t)$.

Fix any x such that $|I_x| < \beta t$ and let us analyze $\Pr_{\hat{x} \leftarrow \text{guess}(\cdot)}[\hat{x} = x]$.

Event E_0 . Let E_0 be the event that $\hat{I}_x = I_x$ and, for all $i \in I_x$: $\hat{x}_i = x_i$ and $\hat{\sigma}_i = \text{hint}(x_i, r_i)$. Then $\Pr[E_0] \geq (t^{\beta t+1} 2^{ht} 2^{\beta t n})^{-1}$. Let us condition on E_0 occurring for the rest of the analysis. In this case, we can replace all the “hatted” values $\hat{I}_x, \hat{\sigma}_i, \hat{x}_i$ with their “unhatted” counterparts $I_x, \sigma_i = \text{hint}(x_i, r_i), x_i$ and we have $y_i^j = \text{Ext}(x_i; s_i^j)$. Furthermore, since the “hatted” values were chosen uniformly at random, E_0 is independent of the choice of r_1, \dots, r_t and of all the “unhatted” values above; therefore conditioning on E_0 does not change their distribution.

Event E_1 . Now, for any fixed choice of r_1, \dots, r_t , define the corresponding procedure \hat{P}_i to be “good” if

$$\Pr_{s \leftarrow \{0,1\}^d}[\hat{P}_i(s) = \text{Ext}(x_i; s)] \geq (1 + 2\delta) \frac{1}{2^m},$$

where the probability is over the choice of $s \leftarrow \{0,1\}^d$ and the internal randomness of \hat{P}_i (i.e., the choice of the index $j^* \leftarrow [Q]$ and the values $z_i^{j^*} \leftarrow \{0,1\}^m$ for $i \in I_x$). Let E_1 be the event that for all $i \notin I_x$ we have \hat{P}_i is good, where the event is over the choice of r_1, \dots, r_t . Define random variables V_i^j over the choice of r_1, \dots, r_t where

$$\begin{aligned} V_i^j &= \Pr_{s \leftarrow \{0,1\}^d}[\hat{P}_i(s) = \text{Ext}(x_i; s) \mid j^* = j] \\ &= \Pr_{s \leftarrow \{0,1\}^d}[P_i(s_1^j, \dots, s_{i-1}^j, s, s_{i+1}^j, \dots, s_t^j, z_1^j, \dots, z_{i-1}^j, y_{i+1}^j, \dots, y_t^j) = \text{Ext}(x_i; s)]. \end{aligned}$$

and $V_i := \sum_{j \in [Q]} V_i^j$. Then \hat{P}_i is good iff $V_i \geq Q(1 + 2\delta) \frac{1}{2^m}$. By Eq. (5), we have $E[V_i] = \sum_j E[V_i^j] \geq Q(1 + 3\delta) \frac{1}{2^m}$. Furthermore, for any fixed i , the variables V_i^j are pairwise independent by Definition 4 and the fact that V_i^j only depends on s_i^j . Therefore $\text{Var}[V_i] = \sum_j \text{Var}[V_i^j] \leq Q$. We can apply the Chebyshev inequality to get:

$$\begin{aligned} \Pr[E_1|E_0] &\geq 1 - \Pr \left[\exists i \notin I_x : V_i < Q(1 + 2\delta) \frac{1}{2^m} \right] \\ &\geq 1 - \sum_{i \notin I_x} \Pr \left[V_i < Q(1 + 2\delta) \frac{1}{2^m} \right] \\ &\geq 1 - \sum_{i \notin I_x} \Pr \left[|V_i - E[V_i]| > Q\delta \frac{1}{2^m} \right] \geq 1 - t \frac{2^{2m}}{\delta^2 Q} \geq \frac{1}{2} \end{aligned}$$

Event E_2 . Now fix any choice of the values in steps (1)–(6) such that E_0, E_1 hold. Let cw_i be the values sampled in step 7. Define the event E_2 to hold if for

all $i \notin I_x$ the value cw_i agrees with $\text{ECC}(x_i)$ in at least $(1 + \delta)2^{d-m}$ coordinates, where the probability is only over the internal randomness used to sample the components $\text{cw}_i(s) \leftarrow \hat{P}_i(s)$. We can define random variables W_i^s which are 1 if $\text{cw}_i(s) = \text{Ext}(x_i; s)$ and 0 otherwise. These variables are mutually independent (since each invocation of \hat{P}_i uses fresh internal randomness) and $E[\sum_s W_i^s] = 2^d \Pr_s[\hat{P}_i(s) = \text{Ext}(x_i; s)] \geq (1 + 2\delta)2^{d-m}$. Therefore, by the Chernoff bound:

$$\begin{aligned} \Pr[E_2|E_1 \wedge E_0] &= 1 - \Pr[\exists i \notin I_x : \sum_s W_i^s \leq (1 + \delta)2^{d-m}] \\ &\geq 1 - \sum_{i \notin I_x} \Pr[\sum_s W_i^s \leq (1 + \delta)2^{d-m}] \\ &\geq 1 - t \cdot e^{-\delta^2 2^{d-m}/8} \geq \frac{1}{2} \end{aligned}$$

Event E_3 . Finally, fix any choice of the values in steps (1)–(7) such that E_0, E_1, E_2 hold. Let E_3 be the event that for each $i \notin \hat{I}_x$ if $\hat{x}_i \leftarrow \mathcal{X}_i$ is the value sampled in step (8) then $\hat{x}_i = x_i$. Then $\Pr[E_3|E_2 \wedge E_1 \wedge E_0] \geq (\frac{1}{L})^t$. Therefore, our guess is correct if E_0, E_1, E_2, E_3 all occur, which gives us the bound in Eq. (4). \square

Corollary 1. *For any $n, m, t, \varepsilon > 0, \alpha > 0$, there exist extractors $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ that are $(t, \alpha, \beta, \varepsilon)$ -multi-instance extracting with either:*

1. seed length $d = n$ and $\beta = \alpha - \frac{O(m + \log t + \log(1/\varepsilon))}{n}$, or
2. seed length $d = O((\log n)(m + \log \log n + \log t + \log(1/\varepsilon)))$ and $\beta = \alpha - \frac{O(d)}{n}$.

In particular, letting λ denote the security parameter, for any input length $n = \omega(\lambda \log \lambda)$ with $n < 2^\lambda$, for number of blocks $t < 2^\lambda$, any entropy rate $\alpha > 0$, there exists an extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with output length $m = \lambda$ and seed length $d = O(\lambda \log n)$, which is a $(t, \alpha, \beta, \varepsilon = 2^{-\lambda})$ -multi-instance randomness extractor with $\beta = \alpha - o(1)$. In other words, the fraction of extracted values that can be replaced by uniform is nearly α .

4 Multi-user Security for Incompressible Encryption

Utilizing multi-instance randomness extractors, we can now explore the multi-user setting for incompressible encryptions. But first, we need to formally define what it means for an incompressible PKE or SKE scheme to be multi-user secure.

We propose a simulation-based security definition. Roughly, the simulator first needs to simulate all the ciphertexts for all the instances *without* seeing any of the message queries. So far, this is akin to the standard semantic security notion for encryption. But we need to now model the fact that the adversary can store ciphertexts for later decryption, at which point it has all the private keys. We therefore add a second phase where the simulator can query for a *subset* of the messages, and then must simulate *all* the private keys. We require that no space-bounded distinguisher can distinguish between receiving real encryptions/real

private keys vs receiving simulated encryptions/keys. The number of messages the simulator can query is related to the storage bound of the distinguisher.

Put formally, let $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ be a public key encryption scheme, to define simulation-based incompressible ciphertext security for the multiple-instance setting, consider the following two experiments:

- In the real mode experiment, the adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ interacts with the challenger \mathcal{C} , who has knowledge of all the adversary's challenge messages.

Real Mode $\text{ExpReal}_{\mathcal{C}, \mathcal{A}=(\mathcal{A}_1, \mathcal{A}_2)}^{\Pi}(\lambda, \eta, \ell, S)$:

1. For $i \in [\eta]$, the challenger \mathcal{C} runs $\text{Gen}(1^\lambda, 1^S)$ to sample $(\text{pk}_i, \text{sk}_i)$.
 2. The challenger \mathcal{C} sends all the pk_i 's to \mathcal{A}_1 .
 3. For each $i \in [\eta]$, \mathcal{A}_1 can produce up to ℓ message queries $\{m_{i,j}\}_{j \in [\ell]}$. The adversary submits all of the message queries *in one single batch* $\{m_{i,j}\}_{i,j}$ and receives $\{\text{ct}_{i,j}\}_{i,j}$ where $\text{ct}_{i,j} \leftarrow \text{Enc}(\text{pk}_i, m_{i,j})$.
 4. \mathcal{A}_1 produces a state st of size at most S .
 5. On input of $\text{st}, \{m_{i,j}\}_{i,j}, \{(\text{pk}_i, \text{sk}_i)\}_i$, \mathcal{A}_2 outputs a bit $1/0$.
- In the ideal mode experiment, the adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ interacts with a simulator \mathcal{S} , which needs to simulate the view of the adversary with no/partial knowledge of the challenge messages.

Ideal Mode $\text{ExpIdeal}_{\mathcal{S}, \mathcal{A}=(\mathcal{A}_1, \mathcal{A}_2)}^{\Pi}(\lambda, \eta, \ell, q, S)$:

1. For $i \in [\eta]$, the simulator \mathcal{S} samples pk_i .
2. The simulator \mathcal{S} sends all the pk_i 's to \mathcal{A}_1 .
3. For each $i \in [\eta]$, and $j \in [\ell]$, \mathcal{A}_1 produces $m_{i,j}$. All of the queries $\{m_{i,j}\}_{i,j}$ are submitted in one batch and the simulator \mathcal{S} produces $\{\text{ct}_{i,j}\}_{i,j}$ *without seeing* $\{m_{i,j}\}_{i,j}$.
4. \mathcal{A}_1 produces a state st of size at most S .
5. The simulator now submits up to q number of (i, j) index pairs, and receives the corresponding messages $m_{i,j}$'s. Then \mathcal{S} simulates all the secret keys sk_i 's.
6. On input of $\text{st}, \{m_{i,j}\}_{i,j}, \{(\text{pk}_i, \text{sk}_i)\}_i$, \mathcal{A}_2 outputs a bit $1/0$.

Notice that the simulator needs to simulate the ciphertexts first without knowing the corresponding messages, and then sample the secret keys so that the ciphertexts appear appropriate under the given messages.

Definition 5 (Multi-instance Simulation-Based CPA Security). *For security parameters $\lambda, \eta(\lambda), \ell(\lambda), q(\lambda)$ and $S(\lambda)$, a public key encryption scheme $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ is (η, ℓ, q, S) -MULT-SIM-CPA secure if for all PPT adversaries $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, there exists a simulator \mathcal{S} such that:*

$$\left| \Pr \left[\text{ExpReal}_{\mathcal{C}, \mathcal{A}}^{\Pi}(\lambda, \eta, \ell, S) = 1 \right] - \Pr \left[\text{ExpIdeal}_{\mathcal{S}, \mathcal{A}}^{\Pi}(\lambda, \eta, \ell, q, S) = 1 \right] \right| \leq \text{negl}(\lambda).$$

Remark 3. If $\ell = 1$, we say that the scheme has only single-ciphertext-per-user security. For $\ell > 1$, we say that the scheme has multi-ciphertext-per-user security.

Remark 4. Notice that by replacing the underlying PKE scheme with a Symmetric Key Encryption (SKE) scheme and modifying corresponding syntaxes (sample only sk 's instead of (pk, sk) pairs, and remove step 2 of the experiments where the adversary receives the pk 's), we can also get a MULT-SIM-CPA security definition for SKE schemes.

5 Symmetric Key Incompressible Encryption

In this section, we explore the multi-user security of incompressible SKEs, both in the low-rate setting and the rate-1 setting. We also present a generic lifting technique to obtain an SKE with multi-ciphertext-per-user security from an SKE with single-ciphertext-per-user security.

5.1 Low Rate Incompressible SKE

For low rate incompressible SKE, it follows almost immediately from multi-instance randomness extractors that the forward-secure storage by Dziembowski [15] is MULT-SIM-CPA secure (by using multi-instance randomness extractors as the “BSM function” and One Time Pad (OTP) as the underlying SKE primitive).

First, let us recall the construction by Dziembowski [15], with the multi-instance randomness extractors and OTP plugged in.

Construction 1 (Forward-Secure Storage [15]). Let λ and S be security parameters. Given $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^w$ a $(t, \alpha, \beta, \epsilon)$ -multi-instance randomness extractor as defined in Definition 3 where the seed length $d = \text{poly}(\lambda)$, output length $w = \text{poly}(\lambda)$ and $n = \frac{S}{(1-\alpha)t} + \text{poly}(\lambda)$, the construction $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ for message space $\{0, 1\}^w$ works as follows:

- $\text{Gen}(1^\lambda, 1^S)$: Sample a seed $s \leftarrow \{0, 1\}^d$ for the randomness extractor, and a key $k' \leftarrow \{0, 1\}^w$. Output $k = (s, k')$.
- $\text{Enc}(k, m)$: To encrypt a message m , first parse $k = (s, k')$ and sample a long randomness $R \leftarrow \{0, 1\}^n$. Compute the ciphertext as $\text{ct} = (R, \text{ct}' = \text{Ext}(R; s) \oplus k' \oplus m)$.
- $\text{Dec}(k, \text{ct})$: First, parse $\text{ct} = (R, \text{ct}')$ and $k = (s, k')$. Then compute $m = \text{Ext}(R; s) \oplus k' \oplus \text{ct}'$.

Correctness is straightforward. Construction 1 is also MULT-SIM-CPA secure. Essentially, the simulator simply sends ct_i 's as uniformly random strings. Then when the simulator sends the keys, it would use the simulator for the multi-instance randomness extractor to get the index subset $I \subset [\eta]$, and for $i \in I$, send k_i as a uniformly random string. For $i \notin I$, it samples the extractor seed s_i and then compute $k'_i = m_i \oplus \text{Ext}(R_i; s_i) \oplus \text{ct}'_i$. Notice that for $i \notin I$, $\text{ct}'_i = m_i \oplus \text{Ext}(R_i; s_i) \oplus k'_i$, and for $i \in I$, $\text{ct}'_i = m_i \oplus u_i \oplus k'_i$ where u_i is a w -bit uniform string. This is now just the definition of multi-instance randomness extractors.

Theorem 1. *Let λ, S be security parameters. If $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^w$ is a $(t, \alpha, \beta, \epsilon)$ -multi-instance randomness extractor with $d, w = \text{poly}(\lambda)$ and $n = \frac{S}{(1-\alpha)t} + \text{poly}(\lambda)$, then Construction 1 is $(t, 1, (1-\beta)t, S)$ -MULT-SIM-CPA secure.*

For a formal hybrid proof of Theorem 1, see the full version [18].

Remark 5. While MULT-SIM-CPA security only requires that no *PPT* adversaries can distinguish between the real mode and the ideal mode experiments, what we have proved for construction 1 here is that it is actually MULT-SIM-CPA secure against *all* (potentially computationally unbounded) adversaries, and hence is information theoretically MULT-SIM-CPA secure.

5.2 Rate-1 Incompressible SKE

Branco, Döttling and Dujmovic [6] construct rate-1 incompressible SKE from HILL-Entropic Encodings [22], extractors and PRGs. We show that by replacing the extractors with multi-instance randomness extractors and slightly modifying the scheme, we get MULT-SIM-CPA security.

First, we recall the definitions and security requirements of a HILL-Entropic Encoding scheme [22].

Definition 6 (HILL-Entropic Encoding [22]). *Let λ be the security parameter. An (α, β) -HILL-Entropic Encoding in the common random string setting is a pair of *PPT* algorithms $\text{Code} = (\text{Enc}, \text{Dec})$ that works as follows:*

- $\text{Enc}_{\text{crs}}(1^\lambda, m) \rightarrow c$: *On input the common random string crs , the security parameter, and a message, outputs a codeword c .*
- $\text{Dec}_{\text{crs}}(c) \rightarrow m$: *On input the common random string and a codeword, outputs the decoded message m .*

It satisfies the following properties.

Correctness. For all $\lambda \in \mathbb{N}$ and $m \in \{0, 1\}^*$, $\Pr[\text{Dec}_{\text{crs}}(\text{Enc}_{\text{crs}}(1^\lambda, m)) = m] \geq 1 - \text{negl}(\lambda)$.

α -Expansion. For all $\lambda, k \in \mathbb{N}$ and for all $m \in \{0, 1\}^k$, $|\text{Enc}_{\text{crs}}(1^\lambda, m)| \leq \alpha(\lambda, k)$.

β -HILL-Entropy. There exists a simulator algorithm SimEnc such that for all polynomial $k = k(\lambda)$ and any ensemble of messages $m = \{m_\lambda\}$ of length $k(\lambda)$, consider the following real mode experiment:

- $\text{crs} \leftarrow \{0, 1\}^{t(\lambda, k)}$
- $c \leftarrow \text{Enc}_{\text{crs}}(1^\lambda, m_\lambda)$

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

- $(\text{crs}', c') \leftarrow \text{SimEnc}(1^\lambda, m_\lambda)$

and let CRS', C' be the corresponding random variables in the simulated experiment. We require that $(\text{CRS}, C) \approx_c (\text{CRS}', C')$ and that $H_\infty(C' | \text{CRS}') \geq \beta(\lambda, k)$.

Moran and Wichs [22] show that we can construct HILL-Entropic Encodings in the CRS model from either the Decisional Composite Residuosity (DCR) assumption [10, 24] or the Learning with Errors (LWE) problem [25]. Their construction achieves $\alpha(\lambda, k) = k(1 + o(1)) + \text{poly}(\lambda)$ and $\beta(\lambda, k) = k(1 - o(1)) - \text{poly}(\lambda)$, which we call a “good” HILL-entropic encoding.

Now we reproduce the construction from [6] with the multi-instance randomness extractors and some other minor changes (highlighted below).

Construction 2 ([6]). Let λ and S be security parameters. Given $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^w$ a $(t, \alpha, \beta, \epsilon)$ -multi-instance randomness extractor where the seed length $d = \text{poly}(\lambda)$, $w = \text{poly}(\lambda)$ and $n = \frac{S}{(1-\alpha)t} + \text{poly}(\lambda)$, $\text{Code} = (\text{Enc}, \text{Dec})$ a “good” (α', β') -HILL-Entropic Encoding scheme, and $\text{PRG} : \{0, 1\}^w \rightarrow \{0, 1\}^n$ a pseudorandom generator secure against non-uniform adversaries, the construction $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ for message space $\{0, 1\}^n$ works as follows:

- $\text{Gen}(1^\lambda, 1^S)$: Sample a seed $s \leftarrow \{0, 1\}^d$ for the randomness extractor, a common random string $\text{crs} \in \{0, 1\}^{\text{poly}(\lambda, n)}$ for the HILL-Entropic Encoding, and a random pad $r \leftarrow \{0, 1\}^n$. Output $k = (s, r, \text{crs})$.
- $\text{Enc}(k, m)$: To encrypt a message m , first parse $k = (s, r, \text{crs})$ and sample a random PRG seed $s' \leftarrow \{0, 1\}^w$. Compute $c_1 = \text{Code.Enc}_{\text{crs}}(1^\lambda, \text{PRG}(s') \oplus r \oplus m)$ and $c_2 = s' \oplus \text{Ext}(c_1, s)$. The final ciphertext is $\text{ct} = (c_1, c_2)$.
- $\text{Dec}(k, \text{ct})$: First, parse $\text{ct} = (c_1, c_2)$ and $k = (s, r, \text{crs})$. Then compute $s' = \text{Ext}(c_1; s) \oplus c_2$ and obtain $m = \text{Code.Dec}_{\text{crs}}(c_1) \oplus \text{PRG}(s') \oplus r$.

Correctness follows from the original construction and should be easy to verify. Notice that by the α' -expansion of the “good” HILL-entropic encoding, the ciphertexts have length $(1 + o(1))n + w + \text{poly}(\lambda) = (1 + o(1))n + \text{poly}(\lambda)$ (the $\text{poly}(\lambda)$ part is independent of n), while the messages have length n . Hence the scheme achieves an optimal rate of 1 ($(1 - o(1))$ to be exact). The keys are bit longer though, having size $d + n + \text{poly}(\lambda, n) = n + \text{poly}(\lambda, n)$. Furthermore, Moran and Wichs [22] show that the CRS needs to be at least as long as the message being encoded. Thus the key has length at least $2n + \text{poly}(\lambda)$.

Theorem 2. *If $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^w$ is a $(t, \alpha, \beta, \epsilon)$ -multi-instance randomness extractor with $n = \frac{S}{(1-\alpha)t} + \text{poly}(\lambda)$, $\text{Code} = (\text{Enc}, \text{Dec})$ is a “good” HILL-entropic encoding with β' -HILL-entropy, and PRG is a pseudo-random generator secure against non-uniform adversaries, then Construction 2 is $(t, 1, (1 - \beta)t, S)$ -MULT-SIM-CPA secure.*

The hybrid proof essentially follows the same structure from [6], except for a different extractor step, the inclusion of the random pad r and the requirement of PRG to be secure against non-uniform attackers. For the detailed hybrid proof of Theorem 2, see the full version [18].

5.3 Dealing with Multiple Messages per User

Above we have showed MULT-SIM-CPA security for SKE schemes where the number of messages per user ℓ is equal to 1. Here, we show how we can generically lift a SKE scheme with single-message-per-user MULT-SIM-CPA security to multiple-messages-per-user MULT-SIM-CPA security.

Construction 3. Let λ, S be security parameters. Given $\text{SKE} = (\text{Gen}, \text{Enc}, \text{Dec})$ a $(\eta, 1, q, S)$ -MULT-SIM-CPA secure SKE with key space $\{0, 1\}^n$ ¹⁰ and \mathcal{F} a

¹⁰ Here we assume SKE’s keys are uniformly random n -bit strings. This is without loss of generality since we can always take the key to be the random coins for Gen.

class of ℓ -wise independent functions with range $\{0, 1\}^n$, we construct $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ as follows.

- $\text{Gen}(1^\lambda, 1^S)$: Sample a random function $f \leftarrow \mathcal{F}$. Output $k = f$.
- $\text{Enc}(k = f, m)$: Sample a short random string r with $|r| = \text{polylog}(\ell)$, compute $k' = f(r)$, and get $c \leftarrow \text{SKE.Enc}(k', m)$. Output $\text{ct} = (r, c)$.
- $\text{Dec}(k = f, \text{ct} = (r, c))$: Compute $k' = f(r)$, and output $m \leftarrow \text{SKE.Dec}(k', c)$.

Correctness should be easy to verify given the correctness of the underlying SKE scheme and the deterministic property of the ℓ -wise independent functions.

Lemma 7. *If SKE is a $(\eta, 1, q, S)$ -MULT-SIM-CPA secure SKE with key space $\{0, 1\}^n$ and \mathcal{F} is a class of ℓ -wise independent functions with range $\{0, 1\}^n$, then Construction 3 is $(\eta/\ell, \ell, q, S - \eta \cdot \text{polylog}(\ell))$ -MULT-SIM-CPA secure.*

Proof. We prove this through a reduction. We show that if there is an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ that breaks the $(\eta/\ell, \ell, q, S - \eta \cdot \text{polylog}(\ell))$ -MULT-SIM-CPA security of Π , then we can construct an adversary $\mathcal{A}' = (\mathcal{A}'_1, \mathcal{A}'_2)$ that breaks the $(\eta, 1, q, S)$ -MULT-SIM-CPA security of SKE. $\mathcal{A}' = (\mathcal{A}'_1, \mathcal{A}'_2)$ works as follows:

- \mathcal{A}'_1 : First, run \mathcal{A}_1 to get a list of message queries $\{m_{i,j}\}_{i \in [\eta/\ell], j \in [\ell]}$. Let $m'_i = m_{(i/\ell)+1, ((i-1) \bmod \ell)+1}$ for $i \in [\eta]$. Notice that here we are essentially flattening the list of messages. Submit the list $\{m'_i\}_{i \in [\eta]}$ and receive $\{\text{ct}'_i\}_{i \in [\eta]}$. Reconstruct $\text{ct}_{i,j} = (r_{i,j}, \text{ct}'_{(i-1)\cdot\ell+j})$ for $i \in [\eta/\ell]$ and $j \in [\ell]$, where $r_{i,j}$ is a uniformly random string sampled from $\{0, 1\}^{\text{polylog}(\ell)}$. Notice that the $r_{i,j}$'s have no collisions under the same i with overwhelming probability. Send the list of ciphertexts $\{\text{ct}_{i,j}\}_{i,j}$ back to \mathcal{A}_1 and receive a state st . Output the state $\text{st}' = (\text{st}, \{r_{i,j}\}_{i,j})$. The size of the state is $|\text{st}'| + \eta \cdot \text{polylog}(\ell) \leq S - \eta \cdot \text{polylog}(\ell) + \eta \cdot \text{polylog}(\ell) = S$.
- \mathcal{A}'_2 : First receive $\text{st}' = (\text{st}, \{r_{i,j}\}_{i,j}, \{m'_i\}_{i \in [\eta]}, \{k'_i\}_{i \in [\eta]})$ from the challenger / simulator. Reorganize $m_{i,j} = m'_{(i-1)\cdot\ell+j}$ for $i \in [\eta/\ell]$ and $j \in [\ell]$. Construct k_i as an ℓ -wise independent function f_i s.t. for all $i \in [\eta/\ell]$ and $j \in [\ell]$, $f_i(r_{i,j}) = k'_{(i-1)\cdot\ell+j}$. Send $\text{st}, \{m_{i,j}\}_{i \in [\eta/\ell], j \in [\ell]}, \{k_i = f_i\}_{i \in [\eta/\ell]}$ to \mathcal{A}_2 and receive a bit b . Output b .

Notice that \mathcal{A}' perfectly simulates the view for \mathcal{A} . If \mathcal{A} says it is in the real mode, this means the ciphertexts are faithful encryptions of the message queries, and hence \mathcal{A}' should be in the real mode as well, and vice versa. Therefore, construction 3 is $(\eta/\ell, \ell, q, S - \eta \cdot \text{polylog}(\ell))$ -MULT-SIM-CPA secure. \square

6 Public Key Incompressible Encryption

Here we explore multi-user security of incompressible Public Key Encryptions (PKEs), considering constructions from [6, 17]. Unlike the SKE setting, where we can generically lift single-ciphertext-per-user security to multi-ciphertext-per-user security, here we show how to obtain multi-ciphertext security by modifying each construction specifically.

6.1 Low Rate Incompressible PKE

For low rate incompressible PKE, we show that the construction from [17] is MULT-SIM-CPA secure by plugging in the multi-instance randomness extractor. Then, we upgrade the construction to have multi-ciphertext-per-user security by upgrading the functionality of the underlying functional encryption scheme.

Construction 4 ([17] with Multi-Instance Randomness Extractor).

Given $FE = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ a single-key selectively secure functional encryption scheme and a $(t, \alpha, \beta, \epsilon)$ -multi-instance randomness extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^w$, with $d = \text{poly}(\lambda)$, $w = \text{poly}(\lambda)$ and $n = \frac{S}{(1-\alpha)t} + \text{poly}(\lambda)$, the construction $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ with message space $\{0, 1\}^w$ works as follows:

- $\text{Gen}(1^\lambda, 1^S)$: First, obtain $(FE.\text{mpk}, FE.\text{msk}) \leftarrow FE.\text{Setup}(1^\lambda)$. Then, generate the secret key for the following function f_v with a hardcoded $v \in \{0, 1\}^{d+w}$:

$$f_v(s' = (s, \text{pad}), \text{flag}) = \begin{cases} s' & \text{if flag} = 0 \\ s' \oplus v & \text{if flag} = 1 \end{cases}.$$

Output $\text{pk} = FE.\text{mpk}$ and $\text{sk} = FE.\text{sk}_{f_v} \leftarrow FE.\text{KeyGen}(FE.\text{msk}, f_v)$.

- $\text{Enc}(\text{pk}, m)$: Sample a random tuple $s' = (s, \text{pad})$ where $s \in \{0, 1\}^d$ is used as a seed for the extractor and $\text{pad} \in \{0, 1\}^w$ is used as a one-time pad. The ciphertext consists of three parts: $FE.\text{ct} \leftarrow FE.\text{Enc}(FE.\text{mpk}, (s', 0))$, a long randomness $R \in \{0, 1\}^n$, and $z = \text{Ext}(R; s) \oplus \text{pad} \oplus m$.
- $\text{Dec}(\text{sk}, \text{ct} = (FE.\text{ct}, R, z))$: First, obtain $s' \leftarrow FE.\text{Dec}(FE.\text{sk}_{f_v}, FE.\text{ct})$, and then use the seed s to compute $\text{Ext}(R; s) \oplus z \oplus \text{pad}$ to recover m .

The correctness follows from the original construction.

Theorem 3. *If FE is a single-key selectively secure functional encryption scheme and $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^w$ is a $(t, \alpha, \beta, \epsilon)$ -multi-instance randomness extractor with $d, w = \text{poly}(\lambda)$ and $n = \frac{S}{(1-\alpha)t} + \text{poly}(\lambda)$, then Construction 4 is $(t, 1, (1 - \beta)t, S)$ -MULT-SIM-CPA secure.*

For the sequence of hybrids, see the full version [18]. The proofs of the hybrid arguments are identical to those from [17], except for the extractor step, which is analogous to the proof of Lemma 5.2 in the full version [18].

Upgrading to Multiple Ciphertexts per User. Additionally, We show that the constructions from [17] can be upgraded to have multi-ciphertext-per-user security. Essentially, all we need is to upgrade the functionality of the underlying functional encryption scheme to work for a slightly more generalized class of functions. We will need functions $f_{\{v_i\}_i}(s, \text{flag}) = s \oplus v_{\text{flag}}$ for hard coded values v_1, \dots, v_ℓ and a special v_0 being the all 0 string. Notice that the original GWZ construction [17] can be viewed as using functions that are a special case where $\ell = 1$. We show how to construct FE schemes for such $f_{\{v_i\}_i}$ functions from plain PKE in the full version [18]. With this new class of functions,

we can achieve $(t, \ell, (1 - \beta)\ell t, S)$ -MULT-SIM-CPA security. In the hybrid proof where we replace $\text{FE.Enc}(\text{FE.mpk}, (s', 0))$ with $\text{FE.Enc}(\text{FE.mpk}, (s' \oplus v, 1))$, now for the j -th message query for the i -th user where $i \in [t]$ and $j \in [\ell]$, we replace $\text{FE.Enc}(\text{FE.mpk}_i, (s'_{i,j}, 0))$ with $\text{FE.Enc}(\text{FE.mpk}_i, (s'_{i,j} \oplus v_{i,j}, j))$. The rest of the hybrid proof follows analogously.

6.2 Rate-1 Incompressible PKE

For rate-1 incompressible PKE, we first show that we can easily plug in the multi-instance randomness extractor to the construction by Guan, Wichs and Zhandry [17]. We also provide a generalization on the construction by Branco, Döttling and Dujmovic [6] using a Key Encapsulation Mechanism (KEM) with a special *non-committing* property. For both constructions, we show how to adapt them to allow for multi-ciphertext-per-user security.

Construction by [17]. We first reproduce the rate-1 PKE construction from [17], with the multi-instance randomness extractors plugged in.

Construction 5 ([17]). Given $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ a rate-1 functional encryption scheme satisfying single-key semi-adaptive security, $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^w$ a $(t, \alpha, \beta, \epsilon)$ -multi-instance randomness extractor with $d, w = \text{poly}(\lambda)$, $n = \frac{S}{(1-\alpha)t} + \text{poly}(\lambda)$ and $\text{PRG} : \{0, 1\}^w \rightarrow \{0, 1\}^n$ a secure PRG against non-uniform adversaries, the construction $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ for message space $\{0, 1\}^n$ works as follows:

- $\text{Gen}(1^\lambda, 1^S)$: First, obtain $(\text{FE.mpk}, \text{FE.msk}) \leftarrow \text{FE.Setup}(1^\lambda)$. Then, generate the secret key for the following function $f_{v,s}$ with a hardcoded large random pad $v \in \{0, 1\}^n$ and a small extractor seed $s \in \{0, 1\}^d$:

$$f_{v,s}(x, \text{flag}) = \begin{cases} x & \text{if flag} = 0 \\ \text{PRG}(\text{Extract}(x; s)) \oplus v & \text{if flag} = 1 \end{cases}$$

- Output $\text{pk} = \text{FE.mpk}$ and $\text{sk} = \text{FE.sk}_{f_{v,s}} \leftarrow \text{FE.KeyGen}(\text{FE.msk}, f_{v,s})$.
- $\text{Enc}(\text{pk}, m)$: The ciphertext is simply an encryption of $(m, 0)$ using the underlying FE scheme, i.e. $\text{FE.ct} \leftarrow \text{FE.Enc}(\text{FE.mpk}, (m, 0))$.
- $\text{Dec}(\text{sk}, \text{ct})$: Decryption also corresponds to FE decryption. The output is simply $\text{FE.Dec}(\text{FE.sk}_{f_{v,s}}, \text{ct}) = f_{v,s}(m, 0) = m$ as desired.

Correctness easily follows from the original construction. The rate of the construction is the rate of the underlying FE multiplied by $\frac{n}{n+1}$. If the FE has rate $(1 - o(1))$, the construction has rate $(1 - o(1))$ as desired.

Theorem 4. *If $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ is a single-key semi-adaptively secure FE scheme, $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^w$ is a $(t, \alpha, \beta, \epsilon)$ -multi-instance randomness extractor, with $d, w = \text{poly}(\lambda)$ and $n = \frac{S}{(1-\alpha)t} + \text{poly}(\lambda)$, and $\text{PRG} : \{0, 1\}^w \rightarrow \{0, 1\}^n$ is a PRG secure against non-uniform adversaries, then Construction 5 is $(t, 1, (1 - \beta)t, S)$ -MULT-SIM-CPA secure.*

For the sequence of hybrids to prove Theorem 4, see the full version [18]. For the proofs of each hybrid argument, see the original [17] paper, since they are identical except for the extractor step (analogous to Lemma 5.2 in the full version [18]) and the PRG against non-uniform attackers step (analogous to Lemma 5.8 in the full version [18]).

Upgrading to Multiple Ciphertexts per User. Upgrading Construction 5 to multi-ciphertext-per-user security is rather straightforward. Since the construction already requires a full functionality FE scheme, we just modify the class of functions that the underlying FE scheme uses, without introducing any new assumptions. Specifically, we now use $f_{\{v_j\}_j, \{s_j\}_j}$ with hard-coded values $v_j \in \{0, 1\}^n$ and $s_j \in \{0, 1\}^d$ for $j \in [\ell]$ that behaves as follows:

$$f_{\{v_j\}_j, \{s_j\}_j}(x, \text{flag}) = \begin{cases} x & \text{if flag} = 0 \\ \text{PRG}(\text{Extract}(x; s_{\text{flag}})) \oplus v_{\text{flag}} & \text{if flag} \in [\ell] \end{cases}.$$

This gives us $(t, \ell, (1 - \alpha)\ell t, S)$ -MULT-SIM-CPA security. Notice that this modification does slightly harm the rate of the scheme, since the flag is now $\log(\ell)$ bits instead of one bit, but asymptotically the rate is still $(1 - o(1))$.

The hybrid proof works analogously to that of Theorem 4, except that in the hybrid proof where we swap the FE encryption of $(m, 0)$ to $(R, 1)$, we now swap from $(m_{i,j}, 0)$ to $(R_{i,j}, j)$ for the j -th ciphertext from the i -th user.

Generalization of Construction by [6]. [6] show how to lift a rate-1 incompressible SKE scheme to a rate-1 incompressible PKE scheme using a Key Encapsulation Mechanism [8] from programmable Hash Proof Systems (HPS) [7, 21]. Their construction satisfies CCA2 security. We show that if we are to relax the security notion to only CPA security, all we need for the lifting is a Key Encapsulation Mechanism with a *non-committing* property, defined below.

Definition 7 (Key Encapsulation Mechanism [8]). Let λ be the security parameters, a Key Encapsulation Mechanism (KEM) is a tuple of algorithms $\Pi = (\text{KeyGen}, \text{Encap}, \text{Decap})$ that works as follows:

- $\text{KeyGen}(1^\lambda, 1^{\mathcal{L}_k}) \rightarrow (\text{pk}, \text{sk})$: The key generation algorithm takes as input the security parameter and the desired symmetric key length \mathcal{L}_k , outputs a pair of public key and private key (pk, sk) .
- $\text{Encap}(\text{pk}) \rightarrow (k, c)$: The encapsulation algorithm takes the public key pk , produces a symmetric key $k \in \{0, 1\}^{\mathcal{L}_k}$, and a header c that encapsulates k .
- $\text{Decap}(\text{sk}, c) \rightarrow k$: The decapsulation algorithm takes as input the private key sk and a header c , and decapsulates the header to get the symmetric key k .

Definition 8 (Correctness of KEM). A key encapsulation mechanism $\text{KEM} = (\text{KeyGen}, \text{Encap}, \text{Decap})$ is said to be correct if:

$$\Pr \left[k' = k : \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(1^\lambda, 1^{\mathcal{L}_k}) \\ (k, c) \leftarrow \text{Encap}(\text{pk}) \\ k' \leftarrow \text{Decap}(\text{sk}, c) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

Definition 9 (Non-Committing). A key encapsulation mechanism $\text{KEM} = (\text{KeyGen}, \text{Encap}, \text{Decap})$ is said to be non-committing if there exists a pair of simulator algorithm $(\text{Sim}_1, \text{Sim}_2)$ such that $\text{Sim}_1(1^\lambda, 1^{\mathcal{L}_k})$ outputs a simulated public key pk' , a header c' and a state st with $|\text{st}| = \text{poly}(\lambda, \mathcal{L}_k)$, and for any given target key $k' \in \{0, 1\}^{\mathcal{L}_k}$, $\text{Sim}_2(\text{st}, k')$ outputs the random coins r^{KeyGen} and r^{Encap} . We require that if we run the key generation and encapsulation algorithm using these random coins, we will get the desired pk' , c' , and k' , ..

$$\Pr \left[\begin{array}{l} \text{pk}' = \text{pk} \\ k' = k \\ c' = c \end{array} : \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(1^\lambda, 1^{\mathcal{L}_k}; r^{\text{KeyGen}}) \\ (k, c) \leftarrow \text{Encap}(\text{pk}; r^{\text{Encap}}) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

Kindly notice that by the correctness property, $\text{Decap}(\text{sk}, c') \rightarrow k'$.

This *non-committing* property allows us to commit to a public key and header first, but then later able to reveal it as an encapsulation of an arbitrary symmetric key in the key space. And it will be impossible to distinguish the simulated public key and header from the ones we get from faithfully running KeyGen and Encap .

Using this non-committing KEM, we are able to construct rate-1 incompressible PKE from rate-1 incompressible SKE, with multi-user security in mind. This is a generalization of the construction by [6].

Construction 6 (Generalization of [6]). For security parameters λ, S , given $\text{KEM} = (\text{KeyGen}, \text{Encap}, \text{Decap})$ a non-committing KEM and $\text{SKE} = (\text{Gen}, \text{Enc}, \text{Dec})$ a rate-1 incompressible SKE for message space $\{0, 1\}^n$, we construct rate-1 incompressible PKE $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ for message space $\{0, 1\}^n$ as follows:

- $\text{Gen}(1^\lambda, 1^S)$: First, run $\text{SKE.Gen}(1^\lambda, 1^S)$ to determine the required symmetric key length \mathcal{L}_k under security parameters λ, S . Then run $(\text{pk}, \text{sk}) \leftarrow \text{KEM.KeyGen}(1^\lambda, 1^{\mathcal{L}_k})$ and output (pk, sk) .
- $\text{Enc}(\text{pk}, m)$: First, run $(k, c_0) \leftarrow \text{KEM.Encap}(\text{pk})$ to sample a symmetric key k , and encapsulate it into a header c_0 . Then compute $c_1 \leftarrow \text{SKE.Enc}(k, m)$. The ciphertext is the tuple (c_0, c_1) .
- $\text{Dec}(\text{sk}, \text{ct} = (c_0, c_1))$: Decapsulate c_0 with sk to obtain $k \leftarrow \text{KEM.Decap}(\text{sk}, c_0)$, and then use k to decrypt c_1 and get $m \leftarrow \text{SKE.Dec}(k, c_1)$.

Correctness follows from the correctness of the underlying incompressible SKE and the KEM scheme. In terms of the rate, to achieve a rate-1 incompressible PKE, we would require the KEM to produce “short” headers, i.e. $|c_0| = \text{poly}(\lambda)$ independent of \mathcal{L}_k (notice that $\mathcal{L}_k = \text{poly}(\lambda, n)$ and needs to be at least as large as n). We can build such KEMs using various efficient encapsulation techniques [2, 3, 5]. With the short header and an incompressible SKE with rate $(1 - o(1))$, the ciphertext length is $n/(1 - o(1)) + \text{poly}(\lambda)$, yielding an ideal rate of $(1 - o(1))$ for the construction. However, these KEMs require long public keys, as opposed to the short public keys in Construction 5.

For security, we prove that if the underlying SKE has MULT-SIM-CPA security, then Construction 6 has MULT-SIM-CPA security as well.

Theorem 5. *If KEM is a non-committing KEM, and SKE is a $(\eta, 1, q, S)$ -MULT-SIM-CPA secure SKE with message space $\{0, 1\}^n$, then Construction 6 is $(\eta, 1, q, S - \eta \cdot \text{poly}(\lambda, n))$ -MULT-SIM-CPA secure.*

Proof. We prove this through a reduction. We show that if there is an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ that breaks the $(\eta, 1, q, S - \eta \cdot \text{poly}(\lambda, n))$ -MULT-SIM-CPA security of Π , then we can construct an adversary $\mathcal{A}' = (\mathcal{A}'_1, \mathcal{A}'_2)$ that breaks the $(\eta, 1, q, S)$ -MULT-SIM-CPA security of SKE. $\mathcal{A}' = (\mathcal{A}'_1, \mathcal{A}'_2)$ works as follows:

- \mathcal{A}'_1 : Use the security parameters λ, S to determine the key length \mathcal{L}_k for the underlying SKE¹¹. For each $i \in [\eta]$, obtain $(\text{pk}_i, c_{0,i}, \text{KEM.st}_i) \leftarrow \text{KEM.Sim}_1(1^\lambda, 1^{\mathcal{L}_k})$. Send $\{\text{pk}_i\}_i$ to \mathcal{A}_1 to get a list of message queries $\{m_i\}_i$. Then, forward the list $\{m_i\}_i$ to the challenger/simulator and receive a list of ciphertexts $\{\text{ct}'_i\}_i$. Construct $\text{ct}_i = (c_{0,i}, \text{ct}'_i)$, and send all $\{\text{ct}_i\}_i$ to \mathcal{A}_1 to receive a state st . Output the state $\text{st}' = (\text{st}, \{\text{KEM.st}_i\}_i)$. The size of the state is $|\text{st}| + \eta \cdot \text{poly}(\lambda, \mathcal{L}_k) \leq S - \eta \cdot \text{poly}(\lambda, n) + \eta \cdot \text{poly}(\lambda, n) = S$.
- \mathcal{A}'_2 : First receive $\text{st}' = (\text{st}, \{\text{KEM.st}_i\}_i), \{m_i\}_i, \{k_i\}_i$ from the challenger/simulator. For each $i \in [\eta]$, run $(r_i^{\text{KeyGen}}, r_i^{\text{Encap}}) \leftarrow \text{KEM.Sim}_2(\text{KEM.st}_i, k_i)$, and $(\text{pk}_i, \text{sk}_i) \leftarrow \text{KEM.KeyGen}(1^\lambda, 1^{\mathcal{L}_k}; r_i^{\text{KeyGen}})$. Notice that pk_i matches the pk_i produced previously by \mathcal{A}'_1 due to the non-committing property of the KEM. Send $\text{st}, \{m_i\}_i, \{(\text{pk}_i, \text{sk}_i)\}_i$ to \mathcal{A}_2 and receive a bit b . Output b .

Notice that \mathcal{A}' perfectly simulates the view for \mathcal{A} . If \mathcal{A} says it is in the real mode interacting with the challenger, this means the ciphertexts ct_i 's are faithful encryptions of the message queries m_i 's for all $i \in [\eta]$. Then we have $\text{SKE.Dec}(k_i, \text{ct}'_i) = m_i$, and hence \mathcal{A}' is also in the real mode. The converse also holds true. Therefore, construction 6 is $(\eta, 1, q, S - \eta \cdot \text{poly}(\lambda, n))$ -MULT-SIM-CPA secure. \square

Upgrading to Multiple Ciphertexts per User. Next we show how to upgrade Construction 6 to have multi-ciphertext-per-user security. All we need is to upgrade the KEM to be ℓ -strongly non-committing, defined as below.

Definition 10 (ℓ -Strongly Non-Committing). *A key encapsulation mechanism $\text{KEM} = (\text{KeyGen}, \text{Encap}, \text{Decap})$ is said to be ℓ -strongly non-committing if there exists a pair of simulator algorithm $(\text{Sim}_1, \text{Sim}_2)$ such that $\text{Sim}_1(1^\lambda, 1^{\mathcal{L}_k})$ outputs a simulated public key pk' , a set of simulated headers $\mathcal{C}' = \{c'_1, c'_2, \dots, c'_\ell\}$ and a state st with $|\text{st}| = \text{poly}(\lambda, \mathcal{L}_k, \ell)$, and for any given set of target keys $\mathcal{K}' = \{k'_1, k'_2, \dots, k'_\ell\}$ where $k'_i \in \{0, 1\}^{\mathcal{L}_k}$ for all $i \in [\ell]$, $\text{Sim}_2(\text{st}, \mathcal{K}')$ outputs a set of random coin pairs $\{(r_i^{\text{KeyGen}}, r_i^{\text{Encap}})\}_{i \in [\ell]}$. We require that if we run the key generation and encapsulation algorithm using the i -th pair of these random coins, we will get the desired pk', c'_i , and k'_i , i.e. for all $i \in [\ell]$:*

$$\Pr \left[\begin{array}{l} \text{pk}' = \text{pk} \\ k'_i = k \\ c'_i = c \end{array} : \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(1^\lambda, 1^{\mathcal{L}_k}; r_i^{\text{KeyGen}}) \\ (k, c) \leftarrow \text{Encap}(\text{pk}; r_i^{\text{Encap}}) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

Kindly notice that by the correctness property, $\text{Decap}(\text{sk}, c'_i) \rightarrow k'_i$.

¹¹ For the ease of syntax, we imagine the security parameters to be part of the public parameters always accessible to the adversary.

We show how to construct ℓ -strongly non-committing KEMs by composing plain non-committing KEMs in the full version [18].

To get multi-ciphertext security, we simply plug in the ℓ -strongly non-committing KEM in place of the plain non-committing KEM in construction 6. The resulting construction has $(\eta/\ell, \ell, q, S-\eta\text{-poly}(\lambda, n, \ell))$ -MULT-SIM-CPA security. The security proof follows analogously from that of Theorem 5.

References

1. Aggarwal, D., Obremski, M., Ribeiro, J., Siniscalchi, L., Visconti, I.: How to extract useful randomness from unreliable sources. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020. LNCS, vol. 12105, pp. 343–372. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45721-1_13
2. Albrecht, M., Cid, C., Paterson, K.G., Tjhai, C.J., Tomlinson, M.: Nts-kem. NIST Submissions **2**, 4–13 (2019)
3. Bardet, M., et al.: Big quake binary goppa quasi-cyclic key encapsulation. NIST Submissions (2017)
4. Barnett Jr., T.: The zettabyte era officially begins (how much is that?). <https://blogs.cisco.com/sp/the-zettabyte-era-officially-begins-how-much-is-that>
5. Bernstein, D.J., et al.: Classic mceliece: conservative code-based cryptography. NIST Submissions (2017)
6. Branco, P., Döttling, N., Dujmovic, J.: Rate-1 incompressible encryption from standard assumptions. In: Kiltz, E., Vaikuntanathan, V. (eds.) TCC 2022, Part II. LNCS, vol. 13748, pp. 33–69. Springer, Heidelberg (2022). https://doi.org/10.1007/978-3-031-22365-5_2
7. Cramer, R., Shoup, V.: Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 45–64. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-46035-7_4
8. 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)
9. Damgård, I.B., Fehr, S., Renner, R., Salvail, L., Schaffner, C.: A tight high-order entropic quantum uncertainty relation with applications. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 360–378. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-74143-5_20
10. Damgård, I., Jurik, M.: A generalisation, a simplification and some applications of Paillier’s probabilistic public-key system. In: Kim, K. (ed.) PKC 2001. LNCS, vol. 1992, pp. 119–136. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-44586-2_9
11. Department, S.R.: Data center storage capacity worldwide from 2016 to 2021, by segment. <https://www.statista.com/statistics/638593/worldwide-data-center-storage-capacity-cloud-vs-traditional/>
12. Dinur, I., Stemmer, U., Woodruff, D.P., Zhou, S.: On differential privacy and adaptive data analysis with bounded space. Cryptology ePrint Archive, Report 2023/171 (2023). <https://eprint.iacr.org/2023/171>
13. Dodis, Y., Quach, W., Wichs, D.: Authentication in the bounded storage model. In: Dunkelman, O., Dziembowski, S. (eds.) EUROCRYPT 2022, Part III. LNCS, vol. 13277, pp. 737–766. Springer, Heidelberg (2022). https://doi.org/10.1007/978-3-031-07082-2_26

14. Dodis, Y., Reyzin, L., Smith, A.: Fuzzy extractors: how to generate strong keys from biometrics and other noisy data. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 523–540. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24676-3_31
15. Dziembowski, S.: On forward-secure storage. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 251–270. Springer, Heidelberg (2006). https://doi.org/10.1007/11818175_15
16. Dziembowski, S., Kazana, T., Zdanowicz, M.: Quasi chain rule for min-entropy. *Inf. Process. Lett.* **134**, 62–66 (2018). <https://doi.org/10.1016/j.ipl.2018.02.007>. <https://www.sciencedirect.com/science/article/pii/S002001901830036X>
17. Guan, J., Wicks, D., Zhandry, M.: Incompressible cryptography. In: Dunkelman, O., Dziembowski, S. (eds.) EUROCRYPT 2022, Part I. LNCS, vol. 13275, pp. 700–730. Springer, Heidelberg (2022). https://doi.org/10.1007/978-3-031-06944-4_24
18. Guan, J., Wicks, D., Zhandry, M.: Multi-instance randomness extraction and security against bounded-storage mass surveillance. *Cryptology ePrint Archive* (2023)
19. Günther, C.G.: An identity-based key-exchange protocol. In: Quisquater, J.-J., Vandewalle, J. (eds.) EUROCRYPT 1989. LNCS, vol. 434, pp. 29–37. Springer, Heidelberg (1990). https://doi.org/10.1007/3-540-46885-4_5
20. Guruswami, V.: List Decoding of Error-Correcting Codes. LNCS, vol. 3282. Springer, Heidelberg (2005). <https://doi.org/10.1007/b104335>
21. Kalai, Y.T.: Smooth projective hashing and two-message oblivious transfer. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 78–95. Springer, Heidelberg (2005). https://doi.org/10.1007/11426639_5
22. Moran, T., Wicks, D.: Incompressible encodings. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020. LNCS, vol. 12170, pp. 494–523. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-56784-2_17
23. Nisan, N.: Pseudorandom generators for space-bounded computation. In: 22nd ACM STOC, pp. 204–212. ACM Press (1990). <https://doi.org/10.1145/100216.100242>
24. Paillier, P.: Public-key cryptosystems based on composite degree residuosity classes. In: Stern, J. (ed.) EUROCRYPT 1999. LNCS, vol. 1592, pp. 223–238. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-48910-X_16
25. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: Gabow, H.N., Fagin, R. (eds.) 37th ACM STOC, pp. 84–93. ACM Press (2005). <https://doi.org/10.1145/1060590.1060603>
26. Vadhan, S.P., et al.: Pseudorandomness. *Found. Trends® Theor. Comput. Sci.* **7**(1–3), 1–336 (2012)