

Selective-Opening Security in the Presence of Randomness Failures

Viet Tung Hoang¹, Jonathan Katz², Adam O’Neill³, and Mohammad Zaheri³

¹ Dept. of Computer Science, Florida State University, USA
tvhoang@cs.fsu.edu

² Dept. of Computer Science, University of Maryland, USA
jkatz@cs.umd.edu

³ Dept. of Computer Science, Georgetown University, USA
adam@cs.georgetown.edu, mz394@georgetown.edu

Abstract. We initiate the study of public-key encryption (PKE) secure against selective-opening attacks (SOA) in the presence of *randomness failures*, i.e., when the sender may (inadvertently) use low-quality randomness. In the SOA setting, an adversary can adaptively corrupt senders; this notion is natural to consider in tandem with randomness failures since an adversary may target senders by multiple means.

Concretely, we first treat SOA security of *nonce-based PKE*. After formulating an appropriate definition of SOA-secure nonce-based PKE, we provide efficient constructions in the non-programmable random-oracle model, based on lossy trapdoor functions.

We then lift our notion of security to the setting of “hedged” PKE, which ensures security as long as the sender’s seed, message, and nonce *jointly* have high entropy. This unifies the notions and strengthens the protection that nonce-based PKE provides against randomness failures even in the non-SOA setting. We lift our definitions and constructions of SOA-secure nonce-based PKE to the hedged setting as well.

1 Introduction

Imagine that an adversary wants to gain access to encrypted communication that various senders are transmitting to a receiver. There are various ways to go about doing this. One is to try to subvert the random-number generator used by the senders. Another is to break-in to the senders’ machines, possibly in an adaptive fashion. Encryption schemes resisting the first sort of attack have been studied in the context of *security under randomness failures* [18, 3, 23, 7, 10] while resistance to the second sort of attack corresponds to the notion of *security against selective-opening attacks* (SOA) [9, 5, 15, 14, 11, 16].¹ However, as far as we are aware, these notions have so far only been considered separately.

¹ There are two forms of SOA security, called coin-revealing (corresponding to sender corruption) and key-revealing (corresponding to receiver corruption). This paper concerns the first one.

We initiate the study of SOA-secure encryption in the presence of randomness failures, providing new definitions and constructions achieving these definitions in the public-key setting.

There are currently three main approaches in the literature to dealing with randomness failures for PKE: (1) *deterministic* PKE [2], which does not use randomness at all but guarantees security only if plaintexts have high entropy, (2) *hedged* PKE, which is randomized and guarantees security as long as plaintexts and the randomness *jointly* have high entropy, and (3) the recently introduced notion of *nonce-based* PKE by Bellare and Tackmann (BT) [10], where each sender uses a uniform seed² in addition to a nonce, and security is guaranteed if either the seed is secret and the nonces are unique, or the seed is revealed and the nonces have high entropy. Hedged PKE and nonce-based PKE are incomparable and are useful in different scenarios, and part of our contribution is to unify them into a single primitive. We start by adding consideration of SOA security to nonce-based PKE. We then lift the resulting notions to the setting of hedged PKE (which subsumes deterministic PKE) as well, thereby adding consideration of SOA to a unified primitive with the guarantees of both nonce-based and hedged PKE.

1.1 Our Results

SELECTIVE-OPENING SECURITY FOR NONCE-BASED PKE. As explained above, the first notion we consider for protecting against randomness failures is *nonce-based* PKE, recently introduced by Bellare and Tackmann [10]. For consistency with the definitions of SOA security we introduce for later notions (where new technical challenges arise), we formulate an indistinguishability-based (rather than simulation-based) definition, which we call N-SO-CPA, along the lines of the indistinguishability-based definition of SOA security for standard PKE [9]. Under our definition, the adversary can (i) learn the seeds of some senders, (ii) choose the nonces for all the other senders, as long as nonces of each individual sender do not repeat. Then, *after* seeing the ciphertexts, the adversary can adaptively corrupt some senders to learn their messages *together with* seeds and nonces. The definition asks that the adversary cannot distinguish between the plaintexts of the uncorrupted senders and a *resampling* of these plaintexts conditioned on the revealed plaintexts.

The next question is whether N-SO-CPA security is achievable. Throughout our work, we focus on constructions in the so-called non-programmable random-oracle model (NPROM) [20]. Intuitively, this means that in a security proof, the constructed adversary must honestly answer (i.e., cannot program) the random oracle queries of the assumed adversary. The NPROM is arguably closer to the standard (random oracle devoid) model than the programmable random oracle model (PROM), since real-world hash functions are not programmable. In this

² The idea is that because a seed is chosen infrequently, it can be generated using high-quality randomness.

model, we give an efficient construction of N-SO-CPA-secure³ nonce-based PKE based on any lossy trapdoor function [21]. The idea is to modify the nonce-based PKE scheme of Bellare and Tackmann, which encrypts a message m using public-key pk , seed xk , and nonce N by encrypting m using any standard (randomized) PKE scheme with public key pk and “synthetic” coins derived from a hash of (xk, N, m) . Here, we use a *specific* randomized encryption scheme based on any lossy trapdoor function. The security proof of the resulting scheme, which we call NE1, relies on switching to the lossy key-generation algorithm and then using the random oracle to argue that the adversary’s choice of which senders to corrupt must be independent of the plaintexts.

SOA+HEDGED SECURITY FOR NONCE-BASED PKE. Unlike nonce-based PKE, *hedged* PKE [3] guarantees security as long as the message and randomness used by the sender *jointly* have high entropy. Indeed, viewing the sender’s seed and nonce together as the sender’s randomness, nonce-based PKE as defined in [10] *lacks* such a guarantee. To get the best of both worlds, we would like to add such a guarantee to nonce-based PKE. This strengthens the protection provided against randomness failures even in the absence of SOA; however, sticking with the main theme of this work, we aim to achieve it in the SOA setting as well. This leads to a definition that we call HN-SO-CPA, which incorporates both hedged and SOA security into the existing notion of nonce-based PKE.

Modeling SOA in the hedged setting is technically challenging. Indeed, Bellare et al. [4] recently showed that a simulation-based notion of SOA security for deterministic PKE (which is a special case of hedged PKE) is impossible to achieve. They also noted that a natural indistinguishability-based definition is (for different reasons) trivially impossible to achieve, and left open the problem of defining a meaningful (yet achievable) definition. To that end, we introduce a novel “comparison-based” definition of SOA for nonce-based PKE, inspired by the comparison-based definition of SOA for deterministic PKE [2, 6] combined with the indistinguishability-based definition of SOA for standard PKE [9]. Roughly, the definition requires that the adversary cannot predict any function of all the plaintexts (i.e., including those of the uncorrupted senders) with much better probability than by computing the same function on a resampling of all the plaintexts conditioned on the revealed plaintexts. For technical reasons, HN-SO-CPA does not protect partial information about the messages depending on the public key, so we still require N-SO-CPA to hold in addition.

We provide two approaches for achieving HN-SO-CPA + N-SO-CPA-secure nonce-based PKE. The first is a generic transform inspired by the “randomized-then-deterministic” transform of [3] in the setting of hedged security. Namely, we propose a “Nonce-then-Deterministic” (NtD) transform in which one obtains a new nonce-based PKE scheme by composing an underlying nonce-based PKE scheme with a deterministic PKE scheme. We require that the underlying deterministic PKE scheme meet a corresponding special case of the HN-SO-CPA definition that we call D-SO-CPA, and achieve it via a scheme DE1 in the NPROM.

³ In the main body of the paper we treat both CPA and CCA security. For simplicity, we do not discuss CCA here.

Interestingly, the scheme DE1 is exactly the recent construction of Bellare and Hoang [7], except that they assume the hash function is UCE-secure [8] and achieve standard security (not SOA). Again, the analysis is quite involved and deals with subtleties neither present in SOA for randomized PKE nor in prior work on deterministic PKE. Alternatively, we show that the scheme NE1 directly achieves both HN-SO-CPA and N-SO-CPA in the NPROM.

SEPARATION RESULTS. Finally, to justify our developing new schemes in the setting of selective-opening security in the presence of randomness failures rather than using existing ones, we show that the N-SO-CPA and D-SO-CPA are not implied by the standard notions (non-SOA) of nonce-based PKE [10] and D-PKE [2], respectively. Our counter-examples rely on the recent result of Hofheinz, Rao, and Wichs (HRW) [15] that separates IND-CCA security from SOA security for randomized PKE. We also show that N-SO-CPA does not imply HN-SO-CPA for nonce-based PKE, meaning the hedged security does strengthen the notion considered for nonce-based PKE in [10].

OPEN QUESTION. We leave obtaining standard-model (versus NPROM) schemes achieving our notions as an open question. Note that our NtD transform is in the standard model, so if we had standard-model instantiations of the underlying primitives we would get a standard-model HN-SO-CPA + N-SO-CPA-secure nonce-based PKE as well.

1.2 Organization

In contrast to the order in which we explained the results above, in the main body of the paper we first present our results on SOA security for deterministic PKE, then move to our results on SOA security for nonce-based PKE, and then finally present our results on hedged security for SOA-secure nonce-based PKE. This is because the results for deterministic PKE constitute the technical core of our work, and form a basis for the results that follow.

2 Preliminaries

NOTATION AND CONVENTIONS. An adversary is an algorithm or tuple of algorithms. All algorithms may be randomized and are required to be efficient unless otherwise indicated; we let PPT stand for “probabilistic, polynomial time.” For an algorithm A we denote by $x \leftarrow^s A(\dots)$ the experiment that runs A on the elided inputs with uniformly random coins and assigns the output to x , and $x \leftarrow^s A(\dots; r)$ to denote the same experiment, but under the coins r instead of randomly chosen ones. If A is deterministic we denote this instead by $x \leftarrow A(\dots)$. We let $[A(\dots)]$ denote the set of all possible outputs of A when run on the elided arguments. If S is a finite set then $s \leftarrow^s S$ denotes choosing a uniformly random element from S and assigning it to s . We denote by $\Pr[P(x) : \dots]$ the probability that some predicate P is true of x after executing the elided experiment.

Let \mathbb{N} denote the set of all non-negative integers. For any $n \in \mathbb{N}$ we denote by $[n]$ the set $\{1, \dots, n\}$. For a vector \mathbf{x} , we denote by $|\mathbf{x}|$ its length (number of

components) and by $\mathbf{x}[i]$ its i -th component. For a vector \mathbf{x} of length n and any $I \subseteq [n]$, we denote by $\mathbf{x}[I]$ the vector of length $|I|$ such that $\mathbf{x}[I] = (\mathbf{x}[i])_{i \in I}$. For a string X , we let $|X|$ denote its length. For any integer $1 \leq i \leq j \leq |X|$, we write $X[i]$ to denote the i th bit of X , and $X[i, j]$ the substring from the i -th to the j -th bit (inclusive) of X .

PUBLIC-KEY ENCRYPTION. A *public-key encryption scheme* PKE with message space Msg is a tuple of algorithms $(\text{Kg}, \text{Enc}, \text{Dec})$. The key-generation algorithm Kg on input 1^k outputs a public key pk and secret key sk . The encryption algorithm Enc on inputs a public key pk and message $m \in \text{Msg}(k)$ outputs a ciphertext c . The deterministic decryption algorithm Dec on inputs a secret key sk and ciphertext c outputs a message m or \perp . We require that for all $(pk, sk) \in [\text{Kg}(1^k)]$ and all $m \in \text{Msg}(1^k)$, the probability that $\text{Dec}(sk, (\text{Enc}(pk, m))) = m$ is 1. We say PKE is *deterministic* if Enc is deterministic.

LOSSY TRAPDOOR FUNCTION. A *lossy trapdoor function* [21] with domain LDom and range LRng is a tuple of algorithms $\text{LT} = (\text{LT.IKg}, \text{LT.LKg}, \text{LT.Eval}, \text{LT.Inv})$ that work as follows. Algorithm LT.IKg on input a unary encoding of the security parameter 1^k outputs an “injective” evaluation key ek and matching trapdoor td . Algorithm LT.LKg on input 1^k outputs a “lossy” evaluation key lk . Algorithm LT.Eval on inputs an (either injective or lossy) evaluation key ek and $x \in \text{LDom}(k)$ outputs $y \in \text{LRng}(1^k)$. Algorithm LT.Inv on inputs a trapdoor td and a $y' \in \text{LRng}(k)$ outputs $x' \in \text{LDom}(k)$. We require the following properties.

Correctness: For all $k \in \mathbb{N}$ and any $(ek, td) \in [\text{LT.IKg}(1^k)]$, it holds that $\text{Inv}(td, \text{LT.Eval}(ek, x)) = x$ for every $x \in \text{LDom}(k)$.

Key indistinguishability: For every distinguisher D , the advantage $\text{Adv}_{\text{LT}, D}^{\text{tdf}}(k) = \Pr[D(ek) \Rightarrow 1 : (ek, td) \leftarrow_s \text{LT.IKg}(1^k)] - \Pr[D(lk) \Rightarrow 1 : lk \leftarrow_s \text{LT.LKg}(1^k)]$ is negligible.

Lossiness: The size of the co-domain of $\text{LT.Eval}(lk, \cdot)$ is at most $|\text{LRng}(k)|/2^{\tau(k)}$ for all $k \in \mathbb{N}$ and all $lk \in [\text{LT.LKg}(1^k)]$. We call τ the *lossiness* of LT .

If the function $\text{LT.Eval}(ek, \cdot)$ is a permutation for any $k \in \mathbb{N}$ and any $(ek, td) \in [\text{LT.IKg}(1^k)]$ then we call LT a *lossy trapdoor permutation*. Both RSA and Rabin are lossy trapdoor permutations under appropriate assumptions [19, 22].

MESSAGE SAMPLERS. A *message sampler* \mathcal{M} is a PPT algorithm that takes as input 1^k and a string $param \in \{0, 1\}^*$, and outputs a vector \mathbf{m} of messages and a vector \mathbf{a} of the same length. Each $\mathbf{a}[i]$ is the auxiliary information that an adversary gains in addition to $\mathbf{m}[i]$, if it breaks into the machine of the sender of $\mathbf{m}[i]$. For example, if each $\mathbf{m}[i]$ is a signature of some string $\mathbf{x}[i]$, then the adversary may be able to obtain even $\mathbf{x}[i]$ in its break-in. We require that \mathcal{M} be associated with functions $v(\cdot)$ and $n(\cdot)$ such that for any $param \in \{0, 1\}^*$, for any $k \in \mathbb{N}$, and any $\mathbf{m} \in [\mathcal{M}(1^k, param)]$, we have $|\mathbf{m}| = v(k)$ and $|\mathbf{m}[i]| = n(k)$, for every $i \leq |\mathbf{m}|$.

A message sampler \mathcal{M} is (μ, d) -*entropic* if

- For any $k \in \mathbb{N}$, any $I \subseteq \{1, \dots, v(k)\}$ such that $|I| \leq d$, any $param \in \{0, 1\}^*$, and $(\mathbf{m}, \mathbf{a}) \leftarrow_s \mathcal{M}(1^k, param)$, conditioning on messages $\mathbf{m}[I]$ and

their auxiliary information $\mathbf{a}[I]$ and $param$, each other message $\mathbf{m}[j]$ (with $j \in \{1, \dots, v(k)\} \setminus I$) must have conditional min-entropy at least μ . Note that here (\mathbf{m}, \mathbf{a}) is sampled independent of the set I .

- Messages $\mathbf{m}[1], \dots, \mathbf{m}[\mathbf{m}]$ must be distinct, for any $param \in \{0, 1\}^*$ and any $\mathbf{m} \in [\mathcal{M}(1^k, param)]$.

In this definition d can be ∞ , which corresponds to a message sampler in which the conditional distribution of each message, given $param$ and all other messages and their corresponding auxiliary information, has at least μ bits of min-entropy.

RESAMPLING. Following [9], let $\text{Coins}[k]$ be the set of coins for $\mathcal{M}(1^k, \cdot)$, and $\text{Coins}[k, \mathbf{m}^*, \mathbf{a}^*, I, param] = \{\omega \in \text{Coins}[k] \mid \mathbf{m}'[I] = \mathbf{m}^* \text{ and } \mathbf{a}'[I] = \mathbf{a}^*, \text{ where } (\mathbf{m}', \mathbf{a}') \leftarrow \mathcal{M}(1^k, param; \omega)\}$. Let $\text{Resamp}_{\mathcal{M}}(1^k, I, \mathbf{m}^*, \mathbf{a}^*, param)$ be the algorithm that first samples $r \leftarrow_s \text{Coins}[k, \mathbf{m}^*, \mathbf{a}^*, I, param]$, then runs $(\mathbf{m}', \mathbf{a}') \leftarrow \mathcal{M}(1^k, param; r)$, and then returns \mathbf{m}' . (Note that $\text{Resamp}_{\mathcal{M}}$ may run in exponential time.) A *resampling algorithm* of \mathcal{M} is an algorithm Rsmp such that $\text{Rsmp}(1^k, I, \mathbf{m}^*, \mathbf{a}^*, param)$ and $\text{Resamp}_{\mathcal{M}}(1^k, I, \mathbf{m}^*, \mathbf{a}^*, param)$ are identically distributed.⁴ A message sampler \mathcal{M} is *fully resamplable* if it admits a PPT resampling algorithm.

PARTIAL RESAMPLING. We also introduce a new notion of “partial resampling.” Let δ be a function and let $\text{Resamp}_{\mathcal{M}, \delta}(1^k, I, \mathbf{m}^*, \mathbf{a}^*, param)$ be the algorithm that samples $r \leftarrow_s \text{Coins}[k, \mathbf{m}^*, \mathbf{a}^*, I, param]$, runs $(\mathbf{m}', \mathbf{a}') \leftarrow \mathcal{M}(1^k, param; r)$, and then returns $\delta(\mathbf{m}', param)$. We say that \mathcal{M} is δ -*partially resamplable* if there is a PT algorithm Rsmp such that $\text{Rsmp}(1^k, I, \mathbf{m}^*, \mathbf{a}^*, param)$ is identically distributed as $\text{Resamp}_{\mathcal{M}, \delta}(1^k, I, \mathbf{m}^*, \mathbf{a}^*, param)$. Such an algorithm Rsmp is called a δ -*partial resampling algorithm* of \mathcal{M} . If a message sampler is already fully resamplable then it’s δ -partially resamplable for any PT function δ .

3 Selective-Opening Security for D-PKE

3.1 Security Notions

Bellare, Dowsley, and Keelveedhi [4] were the first to consider selective-opening security of deterministic PKE (D-PKE). They propose a “simulation-based” semantic security notion, but then show that this definition is unachievable in both the standard model and the non-programmable random-oracle model (NPROM), even if the messages are uniform and independent. To address this, we introduce an alternative, “comparison-based” semantic-security notion that generalizes the original PRIV definition for D-PKE of Bellare, Boldyreva, and O’Neill [2]. In particular, our notion follows the IND-SO-CPA notion of Bellare, Hofheinz, and

⁴ Here for simplicity, we only consider \mathcal{M} and Rsmp such that the distributions of $\text{Rsmp}(1^k, I, \mathbf{m}^*, \mathbf{a}^*, param)$ and $\text{Resamp}_{\mathcal{M}}(1^k, I, \mathbf{m}^*, \mathbf{a}^*, param)$ are identical. Following [9], one might also consider \mathcal{M} and Rsmp such that the two distributions above are statistically close.

Yilek (BHY) [9] in the sense that we compare what partial information the adversary learns from the unopened messages, versus messages resampled from the same conditional distribution.

D-SO-CPA1 SECURITY. Let $\text{PKE} = (\text{Kg}, \text{Enc}, \text{Dec})$ be a D-PKE scheme. To a message sampler \mathcal{M} and an adversary $A = (A.\text{pg}, A.\text{cor}, A.\text{g}, A.\text{f})$, we associate the experiment in Fig. 1 for every $k \in \mathbb{N}$. We say that DE is D-SO-CPA1 secure for a class \mathcal{M} of resamplable message samplers and a class \mathcal{A} of adversaries if for every $\mathcal{M} \in \mathcal{M}$ and any $A \in \mathcal{A}$,

$$\begin{aligned} & \text{Adv}_{\text{DE}, A, \mathcal{M}}^{\text{d-so-cpa1}}(\cdot) \\ &= \Pr \left[\text{D-CPA1-REAL}_{\text{DE}}^{A, \mathcal{M}}(\cdot) \Rightarrow 1 \right] - \Pr \left[\text{D-CPA1-IDEAL}_{\text{DE}}^{A, \mathcal{M}}(\cdot) \Rightarrow 1 \right] \end{aligned}$$

is negligible. In these games, the adversary $A.\text{pg}$ first creates some parameter *param* to feed the message sampler \mathcal{M} . Note that $A.\text{pg}$ is not given the public key, and thus messages \mathbf{m}_1 created by \mathcal{M} are independent of the public key, a necessary restriction of D-PKE pointed out by Bellare et al. [2]. Next, adversary $A.\text{cor}$ will be given both the public key and the ciphertexts \mathbf{c} , and decides which set I of indices that it'd like to open $\mathbf{c}[I]$. It then passes its state to adversary $A.\text{g}$. The latter is also given $(\mathbf{m}_1[I], \mathbf{a}[I])$ and has to output some partial information ω of the message vector \mathbf{m}_1 .

Game $\text{D-CPA1-REAL}_{\text{DE}}^{A, \mathcal{M}}$ returns 1 if the string ω above matches the output of $A.\text{f}(\mathbf{m}_1, \text{param})$ which is the partial information of interest to the adversary. On the other hand, game $\text{D-CPA1-IDEAL}_{\text{DE}}^{A, \mathcal{M}}$ returns 1 if ω matches the output of $A.\text{f}(\mathbf{m}_0, \text{param})$, where \mathbf{m}_0 is the resampled message vector by $\text{Resamp}_{\mathcal{M}}(1^k, \mathbf{m}_1[I], \mathbf{a}[I], I, \text{param})$. Note that in both games, $A.\text{f}$ is not given the public key pk , otherwise it can encrypt the messages it receives and output the resulting ciphertexts, while $A.\text{g}$ outputs \mathbf{c} . Again, this issue is pointed out in [2]: since encryption is deterministic, the ciphertexts themselves are some partial information about the messages. D-PKE can only hope to protect partial information of \mathbf{m} that is independent of pk , and $A.\text{f}$ is therefore stripped of access to pk .

DISCUSSION. For selective-opening attacks against a D-PKE scheme in which an adversary can open d messages, it is clear that the message sampler must be (μ, d) -entropic, where $2^{-\mu(\cdot)}$ is a negligible function, for any meaningful privacy to be achievable. For convenience of discussion, let's say that a scheme is D-SO-CPA1[d] secure if it's D-SO-CPA1 secure for all (μ, d) -entropic, fully resamplable message samplers and all PT adversaries that open at most d ciphertexts, for any μ such that $2^{-\mu(\cdot)}$ is a negligible function. (The resamplability restriction is dropped for $d = 0$.) The D-SO-CPA1[0] security corresponds to the PRIV notion of Bellare et al. [2].⁵

⁵ A technical difference is that, to be consistent with [4], we require the “partial information” to be an efficiently computable function of the messages. This formulation can be shown equivalent to a definition in the style of [2] up to a difference of one in the size of the message vectors output by \mathcal{M} , following [6, Appendix A].

Game D-CPA1-REAL $_{\text{DE}}^{A, \mathcal{M}}(k)$ $param \leftarrow_s A.pg(1^k)$ $(pk, sk) \leftarrow_s \text{Kg}(1^k)$ $(\mathbf{m}_1, \mathbf{a}) \leftarrow_s \mathcal{M}(1^k, param)$ For $i = 1$ to $ \mathbf{m} $ do $\mathbf{c}[i] \leftarrow \text{Enc}(pk, \mathbf{m}_1[i])$ $(state, I) \leftarrow_s A.cor(pk, \mathbf{c}, param)$ $\omega \leftarrow_s A.g(state, \mathbf{m}_1[I], \mathbf{a}[I])$ Return $(\omega = A.f(\mathbf{m}_1, param))$	Game D-CPA1-IDEAL $_{\text{DE}}^{A, \mathcal{M}}(k)$ $param \leftarrow_s A.pg(1^k); (pk, sk) \leftarrow_s \text{Kg}(1^k)$ $(\mathbf{m}_1, \mathbf{a}) \leftarrow_s \mathcal{M}(1^k, param)$ For $i = 1$ to $ \mathbf{m} $ do $\mathbf{c}[i] \leftarrow \text{Enc}(pk, \mathbf{m}_1[i])$ $(state, I) \leftarrow_s A.cor(pk, \mathbf{c}, param)$ $\mathbf{m}_0 \leftarrow_s \text{Resamp}_{\mathcal{M}}(1^k, \mathbf{m}_1[I], \mathbf{a}[I], I, param)$ $\omega \leftarrow_s A.g(state, \mathbf{m}_1[I], \mathbf{a}[I])$ Return $(\omega = A.f(\mathbf{m}_0, param))$
---	--

Fig. 1: Games to define D-SO-CPA1 security.

We note that it is unclear if D-SO-CPA1 $[\infty]$ security implies the classic PRIV security: the latter doesn’t allow opening, but it can handle a broader class of message samplers. Our goal is to find D-PKE schemes that offer D-SO-CPA1 $[d]$ security for any value of d , including the important special cases $d = 0$ (PRIV security) and $d = \infty$ (unbounded opening).

SEPARATION. In the full version, we show that the standard PRIV notion of D-PKE doesn’t imply D-SO-CPA1. Our construction relies on the recent result of Hofheinz, Rao, and Wichs [15] that separates the standard IND-CPA notion and IND-SO-CPA of randomized PKE. Specifically, we build a contrived D-PKE scheme that is PRIV-secure in the standard model, but subject to the following D-SO-CPA1 attack. The message sampler picks a string $s \leftarrow_s \{0, 1\}^{\ell(k)}$ and then secret-share it to $v(k)$ shares $\mathbf{x}[1], \dots, \mathbf{x}[v(k)]$ such that any $t(k)$ shares reveal no information about the secret s . Let $\mathbf{m}[i] \leftarrow \mathbf{x}[i] \parallel \mathbf{u}[i]$ for every $i \in \{1, \dots, v(k)\}$, where $\mathbf{u}[i] \leftarrow_s \{0, 1\}^{2\ell(k)}$. Since s is uniform, any $t + 1$ shares $\mathbf{x}[i]$ are uniform and independent. Thus, this message sampler is $(3\ell, t)$ -entropic. We show that it is also efficiently resamplable. Surprisingly, there is an efficient SOA adversary $(A.cor, A.g)$ that opens just t ciphertexts and can recover all strings $\mathbf{x}[i]$. Next, $A.g$ outputs $\mathbf{x}[1] \oplus \dots \oplus \mathbf{x}[v(k)]$, and $A.f$ outputs the checksum of the first ℓ bits of the given messages. The adversary A thus wins with advantage $1 - 2^{-\ell(k)}$.

D-SO-CPA2 SECURITY. The D-SO-CPA1 security notion only guarantees to protect messages that are fully resamplable. The D-SO-CPA2 notion strengthens that protection, requiring privacy of $\delta(\mathbf{m}, param)$ for any entropic message sampler \mathcal{M} and any δ such that \mathcal{M} is δ -partially resamplable. In Section 5, we’ll see a concrete use of this extra protection, where (i) we have a sampler \mathcal{M} that is not fully resamplable, but (ii) each message itself is a ciphertext, and there’s a function δ such that the plaintexts underneath \mathbf{m} are $\delta(\mathbf{m}, param)$ and \mathcal{M} is δ -partially resamplable. Formally, let

$$\text{Adv}_{\text{DE}, A, \mathcal{M}, \delta}^{\text{d-so-cpa2}}(\cdot)$$

$$= \Pr \left[\text{D-CPA2-REAL}_{\text{DE}}^{A, \mathcal{M}, \delta}(\cdot) \Rightarrow 1 \right] - \Pr \left[\text{D-CPA2-IDEAL}_{\text{DE}}^{A, \mathcal{M}, \delta}(\cdot) \Rightarrow 1 \right],$$

<p>Game D-CPA2-REAL$_{\text{DE}}^{A, \mathcal{M}, \delta}(k)$</p> <p>$param \leftarrow_s A.pg(1^k)$ $(pk, sk) \leftarrow_s \text{Kg}(1^k)$ $(\mathbf{m}, \mathbf{a}) \leftarrow_s \mathcal{M}(1^k, param)$ For $i = 1$ to \mathbf{m} do $\mathbf{c}[i] \leftarrow \text{Enc}(pk, \mathbf{m}[i])$ $(state, I) \leftarrow_s A.cor(pk, \mathbf{c}, param)$ $\omega \leftarrow_s A.g(state, \mathbf{m}[I], \mathbf{a}[I])$ Return $(\omega = A.f(\delta(\mathbf{m}), param))$</p>	<p>Game D-CPA2-IDEAL$_{\text{DE}}^{A, \mathcal{M}, \delta}(k)$</p> <p>$param \leftarrow_s A.pg(1^k); (pk, sk) \leftarrow_s \text{Kg}(1^k)$ $(\mathbf{m}, \mathbf{a}) \leftarrow_s \mathcal{M}(1^k, param)$ For $i = 1$ to \mathbf{m} do $\mathbf{c}[i] \leftarrow \text{Enc}(pk, \mathbf{m}[i])$ $(state, I) \leftarrow_s A.cor(pk, \mathbf{c}, param)$ $z \leftarrow_s \text{Resamp}_{\mathcal{M}, \delta}(1^k, \mathbf{m}[I], \mathbf{a}[I], I, param)$ $\omega \leftarrow_s A.g(state, \mathbf{m}[I], \mathbf{a}[I])$ Return $(\omega = A.f(z, param))$</p>
---	--

Fig. 2: Games to define D-SO-CPA2 security.

where games $\text{D-CPA2-REAL}_{\text{DE}}^{A, \mathcal{M}, \delta}$ and $\text{D-CPA2-IDEAL}_{\text{DE}}^{A, \mathcal{M}, \delta}$ are defined in Fig. 2. In these games, adversary $A.f$ is given either $\delta(\mathbf{m}_1)$ in the real game, or the output of $\text{Resamp}_{\mathcal{M}, \delta}(1^k, \mathbf{m}_1[I], \mathbf{a}[I], I, param)$ in the ideal game. We say that DE is D-SO-CPA2 secure if $\text{Adv}_{\text{DE}, A, \mathcal{M}, \delta}^{\text{d-so-cpa2}}(\cdot)$ is negligible for any (μ, d) -entropic message sampler \mathcal{M} such that $2^{-\mu}$ is a negligible function, any PT adversary A that opens at most d ciphertexts, and any PT functions δ such that \mathcal{M} is δ -partially resamplable.

WEAK EQUIVALENCE. Clearly, the D-SO-CPA2 notion implies D-SO-CPA1: the latter is the special case of the former for fully resamplable samplers, and for a specific function $\delta(\mathbf{m}, param)$ that simply returns \mathbf{m} . Below, we'll show that if we just restrict to fully resamplable samplers, the D-SO-CPA1 notion actually implies D-SO-CPA2. This is expected, because on an entropic, fully resamplable \mathcal{M} , both notions promise to protect any partial information of \mathbf{m} that is independent of the public key.

Proposition 1. Let \mathcal{M} be a fully resamplable sampler, and let δ be a PT function. Then for any adversary A , there is an adversary B such that

$$\text{Adv}_{\text{DE}, A, \mathcal{M}, \delta}^{\text{d-so-cpa2}}(\cdot) \leq \text{Adv}_{\text{DE}, B, \mathcal{M}}^{\text{d-so-cpa1}}(\cdot) .$$

The adversary B opens as many ciphertexts as A , and its running time is about that of A plus the time to run δ .

Proof. Let B be the adversary that is identical to A , but $B.f$ behaves as follows. When it's given a vector \mathbf{m} and parameter $param$, it'll run $z \leftarrow \delta(\mathbf{m}, param)$ and then outputs $A.f(z, param)$. Then $\text{Adv}_{\text{DE}, B, \mathcal{M}}^{\text{d-so-cpa1}}(\cdot) = \text{Adv}_{\text{DE}, A, \mathcal{M}, \delta}^{\text{d-so-cpa2}}(\cdot)$. \square

In the remainder of the paper, we'll have 6 other notions. Any notion xxx considers an arbitrary message sampler \mathcal{M} with a function δ such that \mathcal{M} is δ -partially resamplable. One can consider a variant xxx1 of xxx, in which the message sampler is fully resamplable and only the specific function $\delta(\mathbf{m}, param) = \mathbf{m}$ is considered, and then establish a weak equivalence between xxx1 and xxx. However, it will lead to a proliferation of 12 definitions. We therefore choose to present just the stronger notion xxx.

$\text{DE.Kg}(1^k)$	$\text{DE.Enc}(pk, m)$	$\text{DE.Dec}(sk, c)$
$(ek, td) \leftarrow_s \text{LT.IKg}(1^k)$	$(hk, ek) \leftarrow pk$	$(hk, td) \leftarrow sk$
$hk \leftarrow_s \{0, 1\}^k$	$r \leftarrow H(hk \parallel 0 \parallel m, \text{LT.il}(k))$	$(trap, y) \leftarrow c$
Return $((hk, ek), (hk, td))$	$trap \leftarrow \text{LT.Eval}(ek, r)$	$r \leftarrow \text{LT.Inv}(td, trap)$
	$y \leftarrow H(hk \parallel 1 \parallel r, m) \oplus m$	Return $H(hk \parallel 1 \parallel r, y) \oplus y$
	Return $(trap, y)$	

Fig. 3: **D-PKE** scheme $\text{DE1}[H, \text{LT}]$.

CCA EXTENSION. To add a CCA flavor to D-SO-CPA2, a notion which we call D-SO-CCA, one would allow adversaries $A.\text{cor}$ and $A.g$ oracle access to $\text{Dec}(sk, \cdot)$ with the restriction that they are forbidden from querying a ciphertext in the given \mathbf{c} to this oracle. Let D-CCA-REAL and D-CCA-IDEAL be the corresponding experiments, and define

$$\begin{aligned} & \mathbf{Adv}_{\text{DE}, A, \mathcal{M}, \delta}^{\text{d-so-cca}}(\cdot) \\ &= \Pr \left[\text{D-CCA-REAL}_{\text{DE}}^{A, \mathcal{M}, \delta}(\cdot) \Rightarrow 1 \right] - \Pr \left[\text{D-CCA-IDEAL}_{\text{DE}}^{A, \mathcal{M}, \delta}(\cdot) \Rightarrow 1 \right]. \end{aligned}$$

We say that DE is D-SO-CCA secure if $\mathbf{Adv}_{\text{DE}, A, \mathcal{M}, \delta}^{\text{d-so-cca}}(\cdot)$ is negligible for any (μ, d) -entropic message sampler \mathcal{M} such that $2^{-\mu}$ is a negligible function, any PT adversary A that opens at most d ciphertexts, and any PT functions δ such that \mathcal{M} is δ -partially resamplable.

3.2 Achieving D-SO-CPA2 Security

While the simulation-based definition of Bellare et al. [4] is impossible to achieve even in the non-programmable random-oracle model (NPRM), we show that it is possible to build a D-SO-CPA2 secure scheme in the NPRM. A close variant of our scheme is shown to be PRIV-secure in the standard model [7]. Our scheme can handle messages of any length, and is highly efficient: the asymmetric cost is fixed and thus the amortized cost is about as cheap as a symmetric encryption. It’s also highly practical on short messages. The only public-key primitive that it uses is a lossy trapdoor function [21], which has practical instantiations, e.g., both Rabin and RSA are lossy [19, 22].

ACHIEVING D-SO-CPA2 SECURITY. To handle arbitrary-length messages, we use a hash function H of arbitrary input and output length. On input $(x, \ell) \in \{0, 1\}^* \times \mathbb{N}$, the hash returns $y = H(x, \ell) \in \{0, 1\}^\ell$. Our scheme $\text{DE1}[H, \text{LT}]$ is shown in Fig. 3, where LT is a lossy trapdoor function with domain $\{0, 1\}^{\text{LT.il}}$. Theorem 2 below shows that DE1 is D-SO-CPA2 secure in the NPRM. The proof is in the full version. We stress that for (μ, ∞) -entropic message samplers, our scheme allows the adversary to open as many ciphertexts as it wishes.

Theorem 2. Let LT be a lossy trapdoor function with lossiness τ . Let \mathcal{M} be a (μ, d) -entropic message sampler, and let δ be a function such that \mathcal{M} is δ -partially resamplable. Let $\text{DE1}[H, \text{LT}]$ be as above. In the NPRM, for any

adversary A opening at most d ciphertexts, there is an adversary D such that

$$\mathbf{Adv}_{\text{DE1}[H, \text{LT}], A, \mathcal{M}, \delta}^{\text{d-so-cpa2}}(k) \leq \frac{4q(k)}{2^k} + \frac{4q(k)v(k)}{2^{\mu(k)}} + \frac{v(k)(v(k) + 4q(k))}{2^{\tau(k)}} + 2\mathbf{Adv}_{\text{LT}, D}^{\text{ltdf}}(k),$$

where $q(k)$ is the total number of random-oracle queries of A and \mathcal{M} , and $v(k)$ is the number of messages that \mathcal{M} produces. The running time of D is about that of A plus the time to run δ and an efficient δ -partial resampling algorithm of \mathcal{M} plus the time to run $\text{DE1}[H, \text{LT}]$ to encrypt \mathcal{M} 's messages. Adversary D makes at most q random-oracle queries.

PROOF IDEAS. Let $\text{RO}_1, \text{RO}_2, \text{RO}_3$, and RO_4 denote the oracle interface of $(A.\text{pg}, \mathcal{M}), A.\text{cor}, A.\text{g}$, and $A.\text{f}$ respectively. Initially, each interface simply calls RO . In game-based proofs of ROM-based D-PKE constructions, one often considers the event that $A.\text{pg}$ or \mathcal{M} queries $(hk \parallel x, \ell)$ to RO_1 , and then let the interface lie, instead of calling $\text{RO}(hk \parallel x, \ell)$. This allows the coins $\mathbf{r}[i] \leftarrow \text{RO}(hk \parallel 0 \parallel \mathbf{m}[i], \text{LT}.\text{il}(k))$ to be independent of the messages \mathbf{m} . The discrepancy due to the lying is tiny, since the chance that $A.\text{pg}$ or \mathcal{M} can make such a query is at most $q(k)/2^k$. However, in the SOA setting, this strategy creates the following subtlety. For the resampling algorithm to behave correctly, one has to give it access to RO_1 . Yet the adversary $A.\text{cor}$ can embed some information of hk in I , and therefore it's well possible that the resampling algorithm queries $\text{RO}_1(hk \parallel \cdot, \cdot)$. This issue is unique to SOA security of D-PKE: prior papers of SOA security for randomized PKE never have to deal with this. While getting around the subtlety above is not too difficult, it shows that a rigorous proof for Theorem 2 is not as simple as one might expect.

Suppose that $A.\text{pg}$ and \mathcal{M} never query $\text{RO}_1(hk \parallel \cdot, \cdot)$. The first step in the proof is to move from an injective key ek of LT to a lossy key lk . Next, recall that the adversary $A.\text{cor}$ is given $\text{LT}.\text{Eval}(lk, \mathbf{r}[i])$. Since each synthetic coin $\mathbf{r}[i]$ is uniformly random and LT has lossiness τ , in the view of $A.\text{cor}$, each $\mathbf{r}[i]$ has min-entropy at least $\tau(k)$. Suppose that $A.\text{cor}$ doesn't make any query in $\{hk \parallel 0 \parallel \mathbf{m}[i], hk \parallel 1 \parallel \mathbf{r}[i] \mid 1 \leq i \leq |\mathbf{m}|\}$; this happens with probability at least $1 - q(k)v(k)/2^{\mu(k)} - q(k)v(k)/2^{\tau(k)}$. Then $A.\text{cor}$ knows nothing about \mathbf{m} , and thus I is conditionally independent of \mathbf{m} , given param . Hence in the view of $A.\text{g}$, each $\mathbf{m}[i]$ (for $i \notin I$) still has min-entropy μ , and thus the chance that $A.\text{g}$ can make a query in $\{hk \parallel 0 \parallel \mathbf{m}[i] \mid i \notin I\}$ is at most $v(k)q(k)/2^{\mu(k)}$.

The core of the proof is to bound the probability that the adversary $A.\text{g}$ can make a query in $\{hk \parallel 1 \parallel \mathbf{r}[i] \mid i \notin I\}$. Let X_i be the random variable for the number of pre-images of $\text{LT}.\text{Eval}(lk, \mathbf{r}[i])$. Although in the view of $A.\text{cor}$, the average conditional min-entropy of each $\mathbf{r}[i]$ is $\tau(k)$, the same claim may *not* hold in the view of $A.\text{g}$. For example, the adversary $A.\text{cor}$ may choose to open all but the ciphertext of $\mathbf{m}[j]$, where j is chosen so that $X_j = \min\{X_1, \dots, X_{v(k)}\}$: while $\mathbf{E}(1/X_i) \leq 2^{-\tau(k)}$ for each fixed $i \in \{1, \dots, v(k)\}$, the same bound doesn't work for $\mathbf{E}(1/\min\{X_1, \dots, X_{v(k)}\})$. To get around this, note that the chance that $A.\text{g}$

can make a query in $\{hk \parallel 1 \parallel \mathbf{r}[i] \mid i \notin I\}$ is at most

$$q(k) \cdot \mathbf{E}\left(\sum_{i \notin I} \frac{1}{X_i}\right) \leq q(k) \cdot \mathbf{E}\left(\sum_{i=1}^{v(k)} \frac{1}{X_i}\right) \leq q(k) \cdot \sum_{i=1}^{v(k)} \mathbf{E}\left(\frac{1}{X_i}\right) \leq \frac{q(k)v(k)}{2^{\tau(k)}}.$$

Finally, if I is conditionally independent of \mathbf{m} given $param$, then the re-sampled string z is identically distributed as $\delta(\mathbf{m}, param)$, even conditioning on hk, I , and $param$.⁶ Hence $A.f$ can query $\text{RO}_4(hk \parallel \cdot, \cdot)$ with probability at most $q(k)/2^k$. If all bad events above don’t happen then (i) in the joint view of $A.g$ and $A.f$, the strings $\delta(\mathbf{m}, param)$ and z are identically distributed, and (ii) the output of $A.f$ will be conditionally independent of the ciphertexts and the public key, given $param$. This means the d-so-cpa2 advantage of A is 0.

3.3 Achieving D-SO-CCA Security

To achieve D-SO-CCA security, we modify DE1 construction as follows: In the decryption, once we recover the message m , we’ll re-encrypt it and return \perp if the resulting ciphertext doesn’t match the given one, or the hash image of the message doesn’t match the string obtained via inverting the trapdoor function. The resulting construction DE2 is shown in Fig. 4. The scheme $\text{DE} = \text{DE2}[H, \text{LT}]$ is *unique-ciphertext*, as formalized by Bellare and Hoang [7]: for every $k \in \mathbb{N}$, every $(pk, sk) \in [\text{DE.Kg}(1^k)]$, and every $m \in \{0, 1\}^*$, there is at most a string c such that $\text{DE.Dec}(sk, c) = m$. Theorem 3 below shows that DE2 is D-SO-CCA secure in the NPROM. The re-encrypting trick for lifting CPA to CCA security in the random-oracle model dates back to a paper of Fujisaki and Okamoto [13], but that work only considers randomized PKE and there’s no opening. Still, the proof ideas are quite similar.

Theorem 3. Let LT be a lossy trapdoor function with lossiness τ . Let \mathcal{M} be a (μ, d) -entropic message sampler and let δ be a function such that \mathcal{M} is δ -partially resamplable. Let $\text{DE2}[H, \text{LT}]$ be as above. In the NPROM, for any adversary A opening at most d ciphertexts, there is an adversary D such that

$$\begin{aligned} \text{Adv}_{\text{DE2}[H, \text{LT}], A, \mathcal{M}, \delta}^{\text{d-so-cca}}(k) &\leq \frac{2p(k)}{2^{\text{LT.il}(k)}} + \frac{10q(k)}{2^k} + \frac{4q(k)v(k)}{2^{\mu(k)}} \\ &\quad + \frac{v(k)(v(k) + 8q(k))}{2^{\tau(k)}} + 2\text{Adv}_{\text{LT}, D}^{\text{ltidf}}(k). \end{aligned}$$

where $p(k)$ is the number of decryption-oracle queries of A , $q(k)$ is the total number of random-oracle queries of A and \mathcal{M} , and $v(k)$ is the number of messages

⁶ Even for the simple case that \mathcal{M} is fully resamplable and outputs empty auxiliary information, and $\delta(\mathbf{m}, param) = \mathbf{m}$, note that if I is correlated to \mathbf{m} then \mathbf{m} and the re-sampled \mathbf{m}' may have completely different distributions. For example, consider \mathcal{M} that outputs (m_1, m_2) , with $m_1 \leftarrow_{\$} \{00, 01\}$ and $m_2 \leftarrow_{\$} \{10, 11\}$. Since m_1 and m_2 are independent, \mathcal{M} is fully resamplable. Let $I = \{1\}$ if $m_1 = 00$, and $I = \{2\}$ otherwise. Then $\Pr[\mathbf{m}' = (00, 11)] = 3/8$, whereas $\Pr[\mathbf{m} = (00, 11)] = 1/4$.

DE.Kg(1^k)	DE.Enc(pk, m)	DE.Dec(sk, c)
$(ek, td) \leftarrow_s \text{LT.IKg}(1^k)$ $hk \leftarrow_s \{0, 1\}^k$ $pk \leftarrow (hk, ek)$ $sk \leftarrow (hk, ek, td)$ Return (pk, sk)	$(hk, ek) \leftarrow pk$ $r \leftarrow H(hk \parallel 0 \parallel m, \text{LT.il}(k))$ $trap \leftarrow \text{LT.Eval}(ek, r)$ $y \leftarrow H(hk \parallel 1 \parallel r, m) \oplus m$ Return $(trap, y)$	$(hk, ek, td) \leftarrow sk$ $(trap, y) \leftarrow c$ $r \leftarrow \text{LT.Inv}(td, trap)$ $trap' \leftarrow \text{LT.Eval}(ek, r)$ $m \leftarrow H(hk \parallel 1 \parallel r, y) \oplus y$ $r' \leftarrow H(hk \parallel 0 \parallel m, \text{LT.il}(k))$ If $r' \neq r$ or $trap' \neq trap$ then Return \perp Return m

Fig. 4: **D-PKE scheme** $\text{DE} = \text{DE2}[H, \text{LT}]$. If LT is a lossy trapdoor permutation then in the decryption algorithm, the computation of $trap'$ and the check $trap' \neq trap$ can be omitted.

that \mathcal{M} produces. The running time of D is about that of A plus the time to run δ and an efficient δ -partial resampling algorithm of \mathcal{M} , plus the time to run $\text{DE2}[H, \text{LT}]$ to encrypt \mathcal{M} 's messages and decrypt A 's decryption queries. Adversary D makes at most $2q$ random-oracle queries.

Proof. Let RsmP be an efficient δ -partial resampling algorithm for \mathcal{M} . Consider games G_1 and G_2 in Fig. 5. Then

$$\text{Adv}_{\text{DE2}[H, \text{LT}], A, \mathcal{M}}^{\text{d-so-cca}}(\cdot) = 2 \Pr[G_1(\cdot) \Rightarrow 1] - 1 .$$

Game G_2 is identical to game G_1 , except for the following. In procedure $\text{DEC}(c)$, instead of using the decryption of DE2 to decrypt c , we maintain the set Dom of the suffixes of random-oracle queries (x, ℓ) that $A.\text{cor}$ and $A.g$ make such that $x[1, k+1] = hk \parallel 0$ and $\ell = \text{LT.il}(k)$. If there's $m \in \text{Dom}$ such that the corresponding ciphertext of m is c then we return m ; otherwise return \perp . Wlog, assume that $A.\text{cor}$ stores all random-oracle queries/answers in its state; that is, both $A.\text{cor}$ and $A.g$ also can track Dom and implement the DEC procedure of game G_2 on their own, without calling the decryption oracle.⁷ Let $\text{Range} = \{\text{DE2.Enc}(pk, m) \mid m \in \text{Dom}\}$. On a query $c \in \text{Range}$, the procedures DEC of both games have the same behavior, due to the correctness of the decryption of DE2 . Wlog, assume that both $A.\text{cor}$ and $A.g$ never query $c \in \text{Range}$ to the decryption oracle. (Adversaries $A.\text{cor}$ and $A.g$ are thus assumed to maintain the corresponding ciphertexts of messages in Dom . But this needs additional queries to the random oracle, so the total random-oracle queries of these two adversaries is now at most $2q$.)

⁷ This assumption crucially relies on our use of a domain separation in hashing the coins r and the messages m : we employ $H(hk \parallel 0 \parallel \cdot, \cdot)$ for m , but $H(hk \parallel 1 \parallel \cdot, \cdot)$ for r . In contrast, BH's variant [7] doesn't use domain separation, and one can't make this assumption anymore: building the corresponding ciphertexts may create additional queries to $H(hk \parallel 0 \parallel \cdot, \text{LT.il}(k))$, leading to a possible exponential blowup on the number of random-oracle queries.

<p>Games $G_1(k), G_2(k)$</p> <p>$param \leftarrow_s A.pg^{RO}(1^k); (\mathbf{m}, \mathbf{a}) \leftarrow_s \mathcal{M}^{RO}(1^k, param); z_1 \leftarrow \delta(\mathbf{m}, param)$</p> <p>$hk \leftarrow_s \{0, 1\}^k; (ek, td) \leftarrow_s LT.IKg(1^k)$</p> <p>For $i = 1$ to \mathbf{m} do</p> <p style="padding-left: 20px;">$\mathbf{r}[i] \leftarrow RO(hk \parallel 0 \parallel \mathbf{m}[i], LT.il(k)); trap \leftarrow LT.Eval(ek, \mathbf{r}[i])$</p> <p style="padding-left: 20px;">$y \leftarrow RO_1(hk \parallel 1 \parallel \mathbf{r}[i], \mathbf{m}[i]) \oplus \mathbf{m}[i]; \mathbf{c}[i] \leftarrow (trap, y)$</p> <p>$Dom \leftarrow \emptyset; (state, I) \leftarrow_s A.cor^{DEC, ROSIM}((hk, ek), \mathbf{c}, param)$</p> <p>$\omega \leftarrow_s A.g^{DEC, ROSIM}(state, \mathbf{m}[I], \mathbf{a}[I])$</p> <p>$z_0 \leftarrow_s Rsmpr^{RO}(1^k, \mathbf{m}[I], \mathbf{a}[I], I, param); b \leftarrow_s \{0, 1\}; t \leftarrow_s A.f^{RO}(z_b, param)$</p> <p>If $(\omega = t)$ then return b else return $1 - b$</p> <p>Procedure $ROSIM(x, \ell)$</p> <p style="padding-left: 20px;">If $x > k + 1$ and $x[1, k + 1] = hk \parallel 0$ then $Dom \leftarrow Dom \cup \{x[k + 2, x]\}$</p> <p style="padding-left: 20px;">Return $RO(x, \ell)$</p>	
<p>Procedure $DEC(c)$ // of game G_1</p> <p style="padding-left: 20px;">$sk \leftarrow (hk, td)$</p> <p style="padding-left: 20px;">$m \leftarrow DE1[H, LT].Dec(sk, c)$</p> <p style="padding-left: 20px;">Return m</p>	<p>Procedure $DEC(c)$ // of game G_2</p> <p style="padding-left: 20px;">For $m \in Dom$ do</p> <p style="padding-left: 40px;">If $c = DE1[H, LT].Enc(pk, m)$ then</p> <p style="padding-left: 40px;">Return m</p> <p style="padding-left: 20px;">Return \perp</p>

Fig. 5: **Games G_1 and G_2 of the proof of Theorem 3.** Their procedures DEC are in the bottom-left and bottom-right panels, respectively.

Assume that $A.pg$ and \mathcal{M} never make a random-oracle query (x, ℓ) such that the k -bit suffix of x is hk . This happens with probability at least $1 - q(k)/2^k$. The adversaries can distinguish the games if and only if they can trigger DEC of game G_1 to produce non- \perp output. Let $c = (trap, y)$ be a decryption-oracle query. Let $r = LT.Inv(td, trap)$ and $m = RO(hk \parallel 1 \parallel r) \oplus y$. Due to the unique-ciphertext property of DE2, if this can trigger the DEC procedure of game G_1 to return a non- \perp answer, we must have $m \notin \{\mathbf{m}[1], \dots, \mathbf{m}[|\mathbf{m}|]\} \cup Dom$. Then there is no prior random-oracle query $(x, LT.il(k))$ such that $x = hk \parallel 0 \parallel m$. Hence procedure DEC of game G_1 will return a non- \perp answer only if $r = RO(hk \parallel 0 \parallel m, LT.il(k))$, which happens with probability $2^{-LT.il(k)}$. Multiplying for $p(k)$ decryption-oracle queries,

$$\Pr[G_1(k) \Rightarrow 1] - \Pr[G_2(k) \Rightarrow 1] \leq q(k)/2^k + p(k)/2^{LT.il(k)} .$$

Now in game G_2 , the decryption oracle always return \perp , and thus wlog, assume that the adversaries never make a decryption query, meaning that they only launch a D-SO-CPA2 attack. Hence

$$2 \Pr[G_2(\cdot) \Rightarrow 1] = \mathbf{Adv}_{DE2[H, LT], A, \mathcal{M}}^{d-so-cpa2}(\cdot) .$$

But DE2 and DE1 only differ in the decryption algorithms, which doesn't affect the D-SO-CPA security. Hence from Theorem 2, we can construct a distin-

guisher D of the claimed running time such that

$$\begin{aligned} \mathbf{Adv}_{\text{DE2}[H, \text{LT}], A, \mathcal{M}, \delta}^{\text{d-so-cpa2}}(k) &\leq \frac{8q(k)}{2^k} + \frac{8q(k)v(k)}{2^{\mu(k)}} + \frac{v(k)(v(k) + 8q(k))}{2^{\tau(k)}} \\ &\quad + 2\mathbf{Adv}_{\text{LT}, D}^{\text{tdf}}(k) . \end{aligned}$$

(Note that the bound above is for adversaries who make at most $2q$ random-oracle queries.) Summing up,

$$\begin{aligned} \mathbf{Adv}_{\text{DE2}[H, \text{LT}], A, \mathcal{M}, \delta}^{\text{d-so-cca}}(k) &\leq \frac{2p(k)}{2^{\text{LT.ii}(k)}} + \frac{10q(k)}{2^k} + \frac{4q(k)v(k)}{2^{\mu(k)}} \\ &\quad + \frac{v(k)(v(k) + 8q(k))}{2^{\tau(k)}} + 2\mathbf{Adv}_{\text{LT}, D}^{\text{tdf}}(k) . \end{aligned}$$

□

4 Selective-Opening Security for Nonce-based PKE

Recall that D-PKE protects only unpredictable messages, but in practice, messages often have very limited entropy [12]. Hedge PKE tries to improve this situation by adding the unpredictability of coins. However, the coins generated by Dual EC are completely determined by Big Brother, and those by the buggy Debian RNG have only about 15 bits of min-entropy. In a recent work, Bellare and Tackmann (BT) [10] propose the notion of nonce-based PKE to address this limitation, supporting arbitrary messages. In this section, we extend the notion of nonce-based PKE for SOA setting, and then show how to achieve this.

4.1 Security Notions

NONCE GENERATORS. A *nonce generator* NG with nonce space \mathcal{N} is an algorithm that takes as input the unary encoding 1^k of the security parameter, a current state St , and a *nonce selector* σ . It then probabilistically produces a nonce $N \in \mathcal{N}$ together with an updated state St . That is, $(N, St) \leftarrow_s \text{NG}(1^k, St, \sigma)$. A good nonce generator needs to satisfy the following properties: (i) nonces should never repeat, and (ii) each nonce is unpredictable, even if all nonce selectors are adversarially chosen. Formally, let $\mathbf{Adv}_{\text{NG}, A}^{\text{rp}}(k) = \Pr[\text{RP}_{\text{NG}}^A(k)]$, where game RP is defined in Fig. 6. We say that NG is RP -secure if for any PT adversary A , $\mathbf{Adv}_{\text{NG}, A}^{\text{rp}}(\cdot)$ is a negligible function.

NONCE-BASED PKE. A nonce-based PKE with nonce space \mathcal{N} is a tuple $\text{NE} = (\text{NE.Kg}, \text{NE.Sg}, \text{NE.Enc}, \text{NE.Dec})$. The key generator $\text{NE.Kg}(1^k)$ generates a public key pk and an associated secret key sk . The seed generator $\text{NE.Sg}(1^k)$ produces a sender seed xk . The encryption algorithm NE.Enc takes as input a public key pk , a sender seed xk , a nonce $N \in \mathcal{N}$, and a message m , and then *deterministically* returns a ciphertext c . The decryption algorithm $\text{NE.Dec}(sk, \cdot)$ plays the same role as in traditional randomized PKE; it's not given the nonce or the sender seed.

Game $\text{RP}_{\text{NG}}^A(k)$	Procedure $\text{GEN}(\sigma)$
$St \leftarrow \varepsilon$; $coll \leftarrow \text{false}$	$(N, St) \leftarrow_s \text{NG}(1^k, St, \sigma)$
$\text{Dom} \leftarrow \emptyset$; $N \leftarrow_s A^{\text{GEN}}(1^k)$	If $N \in \text{Dom}$ then $coll \leftarrow \text{true}$
Return $(N \in \text{Dom}) \vee coll$	$\text{Dom} \leftarrow \text{Dom} \cup \{N\}$; Return N

Fig. 6: Game to define security of a nonce generator NG.

Nonce-based PKE can be viewed as a way to harden the randomness at the sender side; the receiver is oblivious to this change. Security of nonce-based PKE should hold when either (i) the seed xk is secret and the nonces are unique, or (ii) the seed is leaked to the adversary, but the nonces are unpredictable to the adversary.⁸

DISCUSSION. To formalize security of nonce-based PKE, BT define two notions, NBP1 and NBP2. Both notions are in the single-sender setting and use nonces generated from a nonce generator NG. The former notion considers the situation when the seed xk is secret, and there’s no security requirement from NG, except the uniqueness of nonces. The latter notion considers the case when the seed xk is given to the adversary; now nonces generated from NG have to satisfy RP security.

When we bring SOA extension to nonce-based PKE below, there will be many changes. First, since there are multiple senders and only some of them can keep their seeds secret, one has to merge the SOA variants of NBP1 and NBP2 into a single definition. Next, because the adversary learns the seeds of some senders, the nonce generator NG must be RP-secure. If we let senders whose seeds are secret use unpredictable nonces from NG then our notion will fail to model the possibility that the adversary can corrupt the nonce generator. Therefore, in our notion, for senders whose seeds are secret, we’ll let the adversary specify their nonces. We require the adversary to be *nonce-respecting*, meaning that the nonces of every single sender must be distinct.

N-SO-CPA. Let NE be a nonce-based PKE scheme and NG be a nonce generator of the same nonce space \mathcal{N} . Let \mathcal{M} be a message sampler, but the generated messages don’t have to be distinct or unpredictable. Let δ be a function such that \mathcal{M} is δ -partially resamplable. The game N-SO-CPA defining the N-SO-CPA security is specified in Fig. 7.

Initially, the game picks seed $\mathbf{xk}[j] \leftarrow_s \text{NE.Sg}(1^k)$ and sets state $\mathbf{st}[j] \leftarrow \varepsilon$ for sender j , with $j = 1, 2, \dots$. The adversary is then given the public key pk and has to specify the list J of senders that it wishes to get the seeds. It’s then

⁸ The definition of BT [10] requires that if the seed xk is secret then security should hold as long as the message/nonce pairs are unique. If one directly extends this to the SOA setting, there will be some pesky issue, as the adversary can detect equality within the message vectors by repeating the nonces. Here for simplicity, we only demand that nonces should be unique, which is analogous to nonce-based symmetric encryption. Nevertheless, our constructions are specific instantiations of BT construction, and thus meet their requirement.

Game $\text{N-SO-CPA}_{\text{NE,NG}}^{A,\mathcal{M},\delta}(k)$

For $j = 1, 2, \dots$ do $\mathbf{xk}[j] \leftarrow_{\$} \text{NE.Sg}(1^k)$; $\mathbf{st}[j] \leftarrow \varepsilon$
 $(pk, sk) \leftarrow_{\$} \text{NE.Kg}(1^k)$; $(J, state) \leftarrow_{\$} A(1^k, pk)$
 $(param, \mathbf{N}, U, \boldsymbol{\sigma}, state) \leftarrow_{\$} A(state, \mathbf{xk}[J])$
 $(\mathbf{m}, \mathbf{a}) \leftarrow_{\$} \mathcal{M}(1^k, param)$; $z_1 \leftarrow \delta(\mathbf{m}, param)$

For $i = 1$ to $|\mathbf{m}|$ do
 $j \leftarrow U[i]$
 If $j \in J$ then $(N, \mathbf{st}[j]) \leftarrow_{\$} \text{NG}(1^k, \mathbf{st}[j], \boldsymbol{\sigma}[i])$; $\mathbf{N}[i] \leftarrow N$
 $\mathbf{c}[i] \leftarrow \text{NE.Enc}(pk, \mathbf{xk}[j], \mathbf{N}[i], \mathbf{m}[i])$
 $(I, state) \leftarrow_{\$} A(state, \mathbf{c})$; $b \leftarrow_{\$} \{0, 1\}$

For $i \in I, j = 1$ to $|\mathbf{m}|$ do
 If $U[i] = U[j]$ then $I \leftarrow I \cup \{j\}$

$z_0 \leftarrow_{\$} \text{Resamp}_{\mathcal{M},\delta}(1^k, \mathbf{m}[I], \mathbf{a}[I], param)$
 $b' \leftarrow_{\$} A(state, \mathbf{m}[I], \mathbf{a}[I], \mathbf{N}[I], \mathbf{xk}[U[I]], z_b)$; Return $(b = b')$

Fig. 7: Game defining N-SO-CPA security.

granted $\mathbf{xk}[J]$ and then has to provide some parameter $param$ for generating $(\mathbf{m}, \mathbf{a}) \leftarrow_{\$} \mathcal{M}(1^k, param)$, together with a vector \mathbf{N} of nonces, a map U that specifies message $\mathbf{m}[i]$ belongs to sender $U[i]$, and a vector $\boldsymbol{\sigma}$ of nonce selectors for NG. Note that the messages \mathbf{m} here can depend on the public key. We require that the adversary be *nonce-respecting*, meaning that $(\mathbf{N}[1], U[1]), (\mathbf{N}[2], U[2]), \dots$ are distinct.

The game then iterates over $i = 1, \dots, |\mathbf{m}|$ to encrypt each message $\mathbf{m}[i]$. If $i \in J$ then $\mathbf{N}[i]$ is overwritten by a nonce N generated by NG as follows. Let $j \leftarrow U[i]$. The nonce generator NG will read the current state $\mathbf{st}[j]$ of sender j and the nonce selector $\boldsymbol{\sigma}[i]$ for the message $\mathbf{m}[i]$, to generate a nonce N and update $\mathbf{st}[j]$. The adversary then is given the ciphertexts and has to output a set I to indicate which ciphertexts it wants to open. Note that opening $\mathbf{c}[i]$ returns not only $(\mathbf{m}[i], \mathbf{a}[i])$ but also the associated nonce and sender seed. Moreover, if the adversary opens a message belonging to sender j , then any other messages of this sender are considered open. Finally, the game resamples z_0 , and let $z_1 \leftarrow \delta(\mathbf{m}, param)$. It picks $b \leftarrow_{\$} \{0, 1\}$, and gives the adversary z_b and $(\mathbf{m}[I], \mathbf{a}[I], \mathbf{xk}[U[I]], \mathbf{N}[I])$. The adversary has to guess the challenge bit b . Define

$$\mathbf{Adv}_{\text{NE,NG},A,\mathcal{M},\delta}^{\text{n-so-cpa}}(k) = 2 \Pr[\text{N-SO-CPA}_{\text{NE,NG}}^{A,\mathcal{M},\delta}(k)] - 1 .$$

We say that NE is N-SO-CPA secure, with respect to NG, if for any message sampler \mathcal{M} and any PT adversary A , and any PT function δ such that \mathcal{M} is δ -partially resamplable, $\mathbf{Adv}_{\text{NE,NG},A,\mathcal{M},\delta}^{\text{n-so-cpa}}(\cdot)$ is a negligible function.

N-SO-CCA. To add a CCA flavor to N-SO-CPA, one would give the adversary oracle access to $\text{Dec}(sk, \cdot)$. Once it's given the ciphertexts \mathbf{c} , it's not allowed to query any $\mathbf{c}[i]$ to the decryption oracle. Let N-SO-CCA be the corresponding game, and define $\mathbf{Adv}_{\text{NE,NG},A,\mathcal{M},\delta}^{\text{n-so-cca}}(k) = 2 \Pr[\text{N-SO-CCA}_{\text{NE,NG}}^{A,\mathcal{M},\delta}(k)] - 1$.

SIMULATION-BASED SECURITY. One could also define an appropriate simulation-based notion of SOA security for nonce-based PKE, which unlike N-SO-CPA would not require the unrevealed messages to be efficiently resampleable, analogously to the SIM-SOA definition for randomized PKE in [9]. However, we conjecture that such a definition is impossible to achieve. We leave this as an open question. In any case, a simulation-based definition of SOA security for nonce-based PKE will indeed be impossible to achieve later when we lift the primitive to the hedged setting, where an existing impossibility result for a simulation-based notion of SOA security for deterministic PKE [4] applies (because hedged PKE generalizes deterministic PKE).

4.2 Separation

We now show that the standard notions for nonce-based PKE of BT [10] do not imply N-SO-CPA. Our separation is based on the recent result of Hofheinz, Rao, and Wichs (HRW) [15] to show that IND-CCA doesn’t imply the notion IND-SO-CPA for randomized PKE.

HRW CONSTRUCTION. Our counterexample is based on the recent (contrived) construction $\text{RE}_{\text{bad}} = (\text{RE}_{\text{bad}}.\text{Kg}, \text{RE}_{\text{bad}}.\text{Enc}, \text{RE}_{\text{bad}}.\text{Dec})$ of HRW. The scheme RE_{bad} is IND-CCA secure, but is vulnerable to the following SOA attack. The message sampler $\mathcal{M}(1^k, \text{param})$ ignores param , picks a secret $s \leftarrow_{\$} \{0, 1\}^\ell$ and then secret-shares it to $v(k)$ messages $\mathbf{m}[1], \dots, \mathbf{m}[v(k)]$ so that any $t(k)$ shares reveal no information of the secret s . In other words, it picks a_0, a_1, \dots, a_t uniformly from $\text{GF}(2^\ell)$, the finite field of size 2^ℓ , and computes $\mathbf{m}[i] \leftarrow f(i)$ for every $i \in \{1, \dots, v(k)\}$, where $f(x) = a_0 + a_1x + \dots + a_tx^t$ is the corresponding polynomial in $\text{GF}(2^\ell)[X]$. Recall that any $t + 1$ shares will uniquely determine the polynomial f (via polynomial interpolation), and thus any $t + 1$ shares are uniformly and independently random. The auxiliary information is empty. Surprisingly, there’s an efficient adversary that opens only t ciphertexts and can recover all messages. We note that HRW’s counter-example is based on public-coin differing-inputs obfuscation [17], which is a very strong assumption.

RESULTS. Let H be a hash function. One can model it as a random oracle, or, for a standard-model result, a primitive that BT call *hedged extractor*. BT show that one can build a nonce-based PKE achieving their notions from an arbitrary IND-CCA secure PKE RE as follows. Given seed xk , nonce N , and message m , one uses $H((xk, N, m))$ to extract synthetic coins r , and then encrypt m via RE under coins r . Now, use the scheme RE_{bad} above to instantiate RE , and let $\text{NE}_{\text{bad}}[H, \text{RE}_{\text{bad}}]$ be the resulting nonce-based PKE. This $\text{NE}_{\text{bad}}[H, \text{RE}_{\text{bad}}]$ achieves BT’s notions.

We now break the N-SO-CPA security of NE_{bad} . The message sampler \mathcal{M} is as described in HRW attack, and let A be the adversary attacking RE_{bad} as above. Note that \mathcal{M} is fully resampleable, and let $\delta(\mathbf{m}, \text{param}) = \mathbf{m}$. Consider the following adversary B attacking NE_{bad} . It specifies $J = \emptyset$, meaning that it doesn’t want to get any sender seed before the opening. It then lets $N[i] = U[i] = i$, for every i . That is, each sender has only a single message. Then, when B gets the

$\underline{\text{NE1.Kg}(1^k)}$ $(ek, td) \leftarrow_s \text{LT.IKg}(1^k); hk \leftarrow_s \{0, 1\}^k$ $pk \leftarrow (hk, ek); sk \leftarrow (hk, td); \text{Return } (pk, sk)$	$\underline{\text{NE1.Sg}(1^k)}$ $xk \leftarrow_s \{0, 1\}^k; \text{Return } xk$
$\underline{\text{NE1.Enc}(pk, xk, N, m)}$ $(hk, ek) \leftarrow pk$ $r \leftarrow H(hk \parallel 0 \parallel (xk, N, m), \text{LT.il}(k))$ $trap \leftarrow \text{LT.Eval}(pk, r); y \leftarrow H(hk \parallel 1 \parallel r, m) \oplus m$ $\text{Return } (trap, y)$	$\underline{\text{NE1.Dec}(sk, c)}$ $(hk, td) \leftarrow sk; (trap, y) \leftarrow c$ $r \leftarrow \text{LT.Inv}(sk, trap)$ $m \leftarrow H(hk \parallel 1 \parallel r, y) \oplus y$ $\text{Return } m$

 Fig. 8: **Nonce-based PKE scheme NE1[H, LT].**

ciphertexts \mathbf{c} , it runs A on those \mathbf{c} . Note that \mathbf{c} are ciphertexts of RE_{bad} , although the coins are only pseudorandom. Still, adversary A can recover all messages by opening just t ciphertexts. When B is given the messages (real or resampled), it compares that with what A recovers. Then $\text{Adv}_{\text{NE,NG},B,\mathcal{M},\delta}^{\text{n-so-cpa}}(k) \geq 1 - 2^{-\ell(k)}$, where ℓ is the length of each message.

4.3 Achieving N-SO-CPA Security

BT's construction of nonce-based PKE is simple. To encrypt a message m under a seed xk , a nonce N , and public key pk , we hash (xk, N, m) to derive a string r , and then uses a traditional randomized PKE to encrypt m under the synthetic coins r and public key pk . Here we'll use BT's construction, but the underlying randomized PKE is a randomized counterpart of the D-PKE scheme DE1 in Section 3.2.

Formally, let H be a hash of arbitrary input and output length, meaning that $H(x, \ell)$ returns an ℓ -bit string. Let LT be a lossy trapdoor function. Our nonce-based PKE $\text{NE1}[H, \text{LT}]$ is described in Fig. 8; it has nonce space $\{0, 1\}^*$ and message space $\{0, 1\}^*$. Theorem 4 below shows that $\text{NE1}[H, \text{LT}]$ is N-SO-CPA secure in the NPROM; the proof is in the full version.

Theorem 4. Let LT be a lossy trapdoor function with lossiness τ . Let \mathcal{M} be a message sampler and let δ be a function such that \mathcal{M} is δ -partially resamplable. Let $\text{NE1}[H, \text{LT}]$ be as above, and let NG be a nonce generator. In the NPROM, for any adversary A , there are adversaries B and D such that

$$\begin{aligned} \text{Adv}_{\text{NE1}[H,\text{LT}],\text{NG},A,\mathcal{M},\delta}^{\text{n-so-cpa}}(k) &\leq 2\text{Adv}_{\text{LT},B}^{\text{ldf}}(k) + 8q(k)v(k) \cdot \text{Adv}_{\text{NG},D}^{\text{rp}}(k) \\ &\quad + \frac{7v(k)(q(k) + v(k))}{2^k} + \frac{12v(k)(q(k) + v(k))}{2^{\tau(k)}}, \end{aligned}$$

where v is the number of messages that \mathcal{M} generates, and q bounds the total number of random-oracle queries that A and \mathcal{M} make. The running time of B or D is about the time to run game $\text{N-SO-CPA}_{\text{NE,NG}}^{A,\mathcal{M},\delta}$, but using an efficient δ -partial resampling algorithm of \mathcal{M} instead of $\text{Resamp}_{\mathcal{M},\delta}$. Each of B and D makes at most q random-oracle queries.

<u>NE2.Kg(1^k)</u> $(ek, td) \leftarrow_s \text{LT.IKg}(1^k)$; $hk \leftarrow_s \{0, 1\}^k$ $pk \leftarrow (hk, ek)$; $sk \leftarrow (hk, ek, td)$ Return (pk, sk)	<u>NE2.Sg(1^k)</u> $xk \leftarrow_s \{0, 1\}^k$; Return xk
<u>NE2.Enc(pk, xk, N, m)</u> $(hk, ek) \leftarrow pk$ $r \leftarrow H(hk \parallel 00 \parallel (xk, N, m), \text{LT.il}(k))$ $y \leftarrow H(hk \parallel 01 \parallel r, m) \oplus m$ $z \leftarrow H(hk \parallel 10 \parallel r \parallel m, k)$ $trap \leftarrow \text{LT.Eval}(pk, r)$ Return $(trap, y)$	<u>NE2.Dec(sk, c)</u> $(hk, ek, td) \leftarrow sk$; $(trap, y, z) \leftarrow c$ $r \leftarrow \text{LT.Inv}(sk, trap)$ $trap' \leftarrow \text{LT.Eval}(pk, r)$ $m \leftarrow H(hk \parallel 01 \parallel r, y) \oplus y$ $z' \leftarrow H(hk \parallel 10 \parallel r \parallel m, k)$ If $(z' \neq z) \vee (trap' \neq trap)$ then return \perp Return m

Fig. 9: Nonce-based PKE scheme $\text{NE2}[H, \text{LT}]$.

4.4 Achieving N-SO-CCA Security

To strengthen NE1 with CCA capability, in the encryption, we append to the ciphertext a hash image of $r \parallel m$. When we decrypt a ciphertext, we'll recover both r and m , and check if the hash image of $r \parallel m$ matches with what's given in the ciphertext. The resulting scheme $\text{NE2}[H, \text{LT}]$ is shown in Fig. 9. The underlying randomized PKE of NE2 is a textbook IND-CCA construction in the ROM (but LT just needs to be an ordinary trapdoor function). Theorem 5 below shows that $\text{NE2}[H, \text{LT}]$ is N-SO-CCA secure in the NPROM; the proof is in the full version.

Theorem 5. Let LT be a lossy trapdoor function with lossiness τ . Let \mathcal{M} be a message sampler and let δ be a function such that \mathcal{M} is δ -partially resamplable. Let $\text{NE2}[H, \text{LT}]$ be as above, and let NG be a nonce generator. In the NPROM, for any adversary A , there are adversaries B and D such that

$$\begin{aligned} \text{Adv}_{\text{NE2}[H, \text{LT}], \text{NG}, A, \mathcal{M}, \delta}^{\text{n-so-cca}}(k) &\leq 2\text{Adv}_{\text{LT}, B}^{\text{ltdf}}(k) + 8v(k)Q(k) \cdot \text{Adv}_{\text{NG}, D}^{\text{TP}}(k) \\ &\quad + \frac{2p(k)}{2^k} + \frac{7v(k)Q(k)}{2^k} + \frac{12v(k)Q(k)}{2^{\tau(k)}}, \end{aligned}$$

where v is the number of messages that \mathcal{M} generates, p is the number of A 's queries to the decryption oracle, q bounds the total number of random-oracle queries that A and \mathcal{M} make, and $Q = q + 2p + v$. The running time of B or D is about the time to run game $\text{N-SO-CCA}_{\text{NE}, \text{NG}}^{A, \mathcal{M}, \delta}$, but using an efficient δ -partial resampling algorithm of \mathcal{M} instead of $\text{Resamp}_{\mathcal{M}, \delta}$. Each of B and D makes at most $q + 2p$ random-oracle queries.

5 Hedged Security for Nonce-based PKE

Recall that the security of nonce-based PKE relies on the assumption that the adversary cannot obtain the secret seeds and corrupt the nonce generator si-

multaneously. Still, this assumption may fail in practice, and it's desirable to retain some security guarantee when seeds and nonces are bad. We capture this via the notion HN-SO-CPA that is a variant of the notion D-SO-CPA2, adapted for the nonce-based setting. A good nonce-based PKE thus has to satisfy both N-SO-CPA and HN-SO-CPA simultaneously. We then extend this treatment to the CCA setting.

5.1 Security Notions

UNPREDICTABLE SAMPLERS. Let \mathcal{M} be a message sampler. We say that \mathcal{M} is (μ, d) -unpredictable if for any $param \in \{0, 1\}^*$,

- (i) For any $(\mathbf{m}, \mathbf{a}) \in [\mathcal{M}(1^k, param)]$, each $\mathbf{a}[i]$ is a tuple (a_i, xk_i, N_i) , where xk_i is a seed and N_i is a nonce. Moreover, $(xk_1, N_1, \mathbf{m}[1]), (xk_2, N_2, \mathbf{m}[2]), \dots$ must be distinct.
- (ii) For any $I \subseteq \{1, \dots, v(k)\}$ such that $|I| \leq d$, and any $i \in \{1, \dots, v(k)\} \setminus I$, for $(\mathbf{m}, \mathbf{a}) \leftarrow_s \mathcal{M}(1^k, param)$ the conditional min-entropy of $(\mathbf{m}[i], xk_i, N_i)$ given $(\mathbf{m}[I], \mathbf{a}[I], param)$ is at least μ , where $v(k)$ is the number of messages that \mathcal{M} produces and xk_i and N_i are the seed and nonce specified by $\mathbf{a}[i]$.

Defining unpredictable samplers allows us to model the situation when the seeds, nonces, and messages are related, and quantify security based on the combined min-entropy of each message with its nonce and seed.

HN-SO-CPA SECURITY. Let \mathbf{NE} be a nonce-based PKE scheme, and let \mathcal{M} be an unpredictable message sampler. Let δ be a function such that \mathcal{M} is δ -partially resamplable. Let $A = (A.pg, A.cor, A.g, A.f)$ be an adversary. Define

$$\begin{aligned} & \mathbf{Adv}_{\mathbf{NE}, A, \mathcal{M}, \delta}^{\text{hn-so-cpa}}(\cdot) \\ &= \Pr \left[\text{HN-CPA-REAL}_{\mathbf{NE}}^{A, \mathcal{M}, \delta}(\cdot) \Rightarrow 1 \right] - \Pr \left[\text{HN-CPA-IDEAL}_{\mathbf{NE}}^{A, \mathcal{M}, \delta}(\cdot) \Rightarrow 1 \right], \end{aligned}$$

where the games are defined in Fig. 10.

HN-SO-CCA SECURITY. To add a CCA flavor to HN-SO-CPA, one would give $A.cor$ and $A.g$ oracle access to $\text{Dec}(sk, \cdot)$. They are not allowed to query any $\mathbf{c}[i]$ to the decryption oracle. Let HN-CCA-REAL and HN-CCA-IDEAL be the corresponding games, and define

$$\begin{aligned} & \mathbf{Adv}_{\mathbf{NE}, A, \mathcal{M}, \delta}^{\text{hn-so-cca}}(\cdot) \\ &= \Pr \left[\text{HN-CCA-REAL}_{\mathbf{NE}}^{A, \mathcal{M}, \delta}(\cdot) \Rightarrow 1 \right] - \Pr \left[\text{HN-CCA-IDEAL}_{\mathbf{NE}}^{A, \mathcal{M}, \delta}(\cdot) \Rightarrow 1 \right]. \end{aligned}$$

SEPARATION. We now show that N-SO-CCA doesn't imply HN-SO-CPA, even if \mathcal{M} picks $\mathbf{m}[i] \leftarrow_s \{0, 1\}^k$ and $\mathbf{a}[i] = (i, i, i)$, and there's no opening. Note that \mathcal{M} is fully resamplable, and consider the function δ such that $\delta(\mathbf{m}, param) = param$. Let H be a hash and LT be a lossy trapdoor function. Let $\mathbf{NE}_{\text{bad}}[H, \text{LT}]$ be the following variant of $\mathbf{NE2}[H, \text{LT}]$. To encrypt message m under public key pk ,

<p>Game $\text{HN-CPA-REAL}_{\text{NE}}^{A, \mathcal{M}, \delta}(k)$ $param \leftarrow_s A.\text{pg}(1^k)$; $(pk, sk) \leftarrow_s \text{Kg}(1^k)$; $(\mathbf{m}, \mathbf{a}) \leftarrow_s \mathcal{M}(1^k, param)$ For $i = 1$ to \mathbf{m} do $(a, xk, N) \leftarrow \mathbf{a}[i]$; $\mathbf{c}[i] \leftarrow \text{Enc}(pk, xk, N, \mathbf{m}[i])$ $(state, I) \leftarrow_s A.\text{cor}(pk, \mathbf{c}, param)$; $\omega \leftarrow_s A.\text{g}(state, \mathbf{m}[I], \mathbf{a}[I])$ $z_1 \leftarrow \delta(\mathbf{m}, param)$ Return $(\omega = A.f(z_1, param))$</p> <p>Game $\text{HN-CPA-IDEAL}_{\text{NE}}^{A, \mathcal{M}, \delta}(k)$ $param \leftarrow_s A.\text{pg}(1^k)$; $(pk, sk) \leftarrow_s \text{Kg}(1^k)$; $(\mathbf{m}, \mathbf{a}) \leftarrow_s \mathcal{M}(1^k, param)$ For $i = 1$ to \mathbf{m} do $(a, xk, N) \leftarrow \mathbf{a}[i]$; $\mathbf{c}[i] \leftarrow \text{Enc}(pk, xk, N, \mathbf{m}[i])$ $(state, I) \leftarrow_s A.\text{cor}(pk, \mathbf{c}, param)$; $\omega \leftarrow_s A.\text{g}(state, \mathbf{m}[I], \mathbf{a}[I])$ $z_0 \leftarrow_s \text{Resamp}_{\mathcal{M}, \delta}(1^k, \mathbf{m}[I], \mathbf{a}[I], I, param)$ Return $(\omega = A.f(z_0, param))$</p>

Fig. 10: Games to define HN-SO-CPA security.

seed xk and nonce N , instead of hashing (xk, N, m) to derive synthetic coins r , we just hash (xk, N) . The proof of Theorem 5 can be recast to justify the N-SO-CCA security NE_{bad} . However, without even opening, one can trivially break HN-SO-CPA security of NE_{bad} as follows. First, adversary $A.\text{pg}$ outputs an arbitrary $param$. Next, adversary $A.\text{cor}$ stores the ciphertexts and the public key in its state, and outputs $I = \emptyset$. Adversary $A.\text{g}$ computes $r \leftarrow H(hk \parallel 00 \parallel (1, 1))$, parses $(trap, y, z) \leftarrow \mathbf{c}[1]$, and outputs $\mathbf{m}[1] = y \oplus H(hk \parallel 01 \parallel r, |y|)$. Finally, adversary $A.f(\mathbf{m}^*, param)$ simply outputs $\mathbf{m}^*[1]$. The adversaries win with advantage $1 - 2^{-k}$.

5.2 Achieving HN-SO-CPA Security

NtD TRANSFORM. We first give a transform Nonce-then-Deterministic (**NtD**). Let DE be a D-SO-CPA2 secure D-PKE and NE be an N-SO-CPA secure nonce-based PKE. Then $\text{NtD}[\text{NE}, \text{DE}]$ achieves both HN-SO-CPA and N-SO-CPA security simultaneously. The resulting nonce-based PKE $\overline{\text{NE}}$ is a double encryption: it first encrypts via NE , and then uses DE to encrypt the resulting ciphertext.⁹ The transform **NtD** is shown in Fig. 11, and Theorem 6 below confirms that it works as claimed.

DISCUSSION. To explain why **NtD** works, note that using an outer D-PKE on the ciphertext of NE doesn't affect its N-SO-CPA security, and thus $\overline{\text{NE}} = \text{NtD}[\text{NE}, \text{DE}]$ inherits the N-SO-CPA security of NE . For HN-SO-CPA security, there are some subtle points as follows.

⁹ For simplicity, we assume that the ciphertext length of $\overline{\text{NE}}$ is the plaintext length of DE . One may also consider a more generalized setting in which the ciphertext length of NE is smaller than the plaintext length of DE . In this case one needs to pad 10^* to the ciphertexts of NE before feeding them to DE .

$\overline{\text{NE}}.\text{Kg}(1^k)$	$\overline{\text{NE}}.\text{Enc}(pk, xk, N, m)$	$\overline{\text{NE}}.\text{Dec}(sk, c)$
$(pk_n, sk_n) \leftarrow_s \text{NE.Kg}(1^k)$	$(pk_n, pk_d) \leftarrow pk$	$(sk_n, sk_d) \leftarrow sk$
$(pk_d, sk_d) \leftarrow_s \text{DE.Kg}(1^k)$	$m' \leftarrow (m \parallel xk \parallel N)$	$y \leftarrow \text{DE.Dec}(sk_d, c)$
$pk \leftarrow (pk_n, pk_d); sk \leftarrow (sk_n, sk_d)$	$y \leftarrow \text{NE.Enc}(pk_n, xk, N, m')$	$m' \leftarrow \text{NE.Dec}(sk_n, y)$
Return (pk, sk)	$c \leftarrow \text{DE.Enc}(pk_d, y)$	$(m \parallel xk \parallel N) \leftarrow m'$
	Return c	Return $(m \parallel xk \parallel N)$

Fig. 11: **Nonce-based PKE scheme** $\overline{\text{NE}} = \text{NtD}[\text{NE}, \text{DE}]$. It uses the same seed-generating algorithm as NE.

First, the “messages” for DE are the ciphertexts produced by NE. Now, the D-SO-CPA2 security demands that those “messages” must have good min-entropy, but we only know that the combined min-entropy of each message with its nonce and seed is μ . We need a bound, call it $\text{NE.Guess}(\mu)$, to quantify the min-entropy of the ciphertexts of NE. Therefore, let $\text{NE.Guess}(\mu(k))$ be biggest number that, for any seed xk , any nonce N , any message m , and any random variable X such that the conditional min-entropy of (m, xk, N) given X is at least $\mu(k)$, and $(pk, sk) \leftarrow_s \text{NE.Kg}(1^k)$ independent of (m, xk, N, X) , the conditional min-entropy of $\text{NE.Enc}(pk, xk, N, m)$ given X is at least $\text{NE.Guess}(\mu(k))$. We say that NE is *entropy-preserving* if for any μ such that $2^{-\mu}$ is negligible, so is $2^{-\text{NE.Guess}(\mu)}$. For example, one can show that $\text{NE1}[H, \text{LT}].\text{Guess}(\mu(k)) \geq \min\{k, \mu(k)/2\} - 1$, by modeling $h_{hk}(\cdot) = H(hk \parallel 0 \parallel \cdot, \text{LT.il}(k))$ as a universal hash function, and using the Generalized Leftover Hash Lemma [1, Lemma 3.4]. Hence NE1 is entropy-preserving.

Next, we need to build an adversary B attacking DE from an adversary A that attacks $\overline{\text{NE}}$. Then $B.\text{pg}$ will run $param \leftarrow_s A.\text{pg}(1^k)$, pick $(pk, sk) \leftarrow_s \text{NE.Kg}(1^k)$, and outputs $pars = (pk, sk, param)$, asking its sampler $\overline{\mathcal{M}}$ to run \mathcal{M} and encrypt the resulting messages, nonces, and seeds under pk . At some point, $A.\text{cor}$ will ask to open some ciphertexts $\mathbf{c}[I]$ to get the corresponding $\mathbf{m}[I], \mathbf{xk}[I], \mathbf{N}[I]$, but the opened “messages” that $B.g$ receives are $\text{NE.Enc}(pk, \mathbf{xk}[i], \mathbf{N}[i], \mathbf{m}[i])$. Although $B.g$ knows the secret key sk of NE, if we use NE1 to instantiate NE then one can’t recover $(\mathbf{N}[i], \mathbf{xk}[i])$ from just $c_i = \text{NE.Enc}(pk, \mathbf{xk}[i], \mathbf{N}[i], \mathbf{m}[i])$ and sk . Thanks to our explicit modeling of the auxiliary information, adversary B does get $(\mathbf{xk}[i], \mathbf{N}[i])$ when it opens $\mathbf{c}[i]$.

Finally, one has to reason about the resamplability of the constructed sampler $\overline{\mathcal{M}}$. Had we restricted our notions to fully resamplable samplers and the function $\delta(\mathbf{m}, param) = \mathbf{m}$, we would have run into problem here. Why so? The resampling algorithm Rsmp of $\overline{\mathcal{M}}$ has to generate $\text{NE.Enc}(pk, \mathbf{xk}'[i], \mathbf{N}'[i], \mathbf{m}'[i])$, but it only knows pk and another algorithm Rsmp to generate \mathbf{m}' . That is, it’s unclear how to resample the seeds \mathbf{xk}' and nonces \mathbf{N}' . Using partial resamplability solves this issue. To justify this, suppose that we need to justify the HN-SO-CPA security of $\overline{\text{NE}}$ with respect to function δ . Then, we’ll find *another* function $\overline{\delta}$ such that $\text{Adv}_{\overline{\text{NE}}, A, \overline{\mathcal{M}}, \delta}^{\text{hn-so-cpa}}(\cdot) \leq \text{Adv}_{\text{DE}, B, \overline{\mathcal{M}}, \overline{\delta}}^{\text{d-so-cpa2}}(\cdot)$, and at the same time, $\overline{\mathcal{M}}$ is $\overline{\delta}$ -partially

Algorithm $B(1^k, pk)$ Return $A(1^k, pk)$	Algorithm $B(state, \mathbf{x}k^*)$ Return $A(state, \mathbf{x}k^*)$
Algorithm $B(state, \mathbf{c})$ $(pk_d, sk_d) \leftarrow \text{DE.Kg}(1^k)$ For $i = 1$ to $ \mathbf{c} $ do $\bar{c} \leftarrow \text{DE.Enc}(pk_d, \mathbf{c}[i])$ Return $A(state, \bar{\mathbf{c}})$	Algorithm $B(state, \mathbf{m}^*, \mathbf{a}^*, \mathbf{N}^*, \mathbf{x}k^*)$ Return $A(state, \mathbf{m}^*, \mathbf{a}^*, \mathbf{N}^*, \mathbf{x}k^*)$

Fig. 12: N-SO-CPA adversary B in the proof of Theorem 6.

resamplable. The function $\bar{\delta}(\mathbf{x}, pars)$ works as follows. It first parses $pars$ as $(pk, sk, param)$, runs $\mathbf{m}[i] \leftarrow \text{NE.Dec}(sk, \mathbf{x}[i])$, and then outputs $\delta(\mathbf{m}, param)$.

We stress that the NtD transform works in both the standard model and the NPRM. (Of course, this assumes that there are standard-model D-SO-CPA2 secure D-PKE and N-SO-CPA secure entropy-preserving nonce-based PKE.)

Theorem 6. Let NE be a nonce-based PKE, and let DE be a D-PKE scheme such that the ciphertext length of the former is a plaintext length of the latter. Let $\bar{\text{NE}} = \text{NtD}[\text{NE}, \text{DE}]$.

N-SO-CPA security: For any adversary A , any message sampler \mathcal{M} , any function $\bar{\delta}$ such that \mathcal{M} is $\bar{\delta}$ -partially resamplable, and any nonce generator NG, there is an adversary B such that

$$\text{Adv}_{\bar{\text{NE}}, \text{NG}, A, \mathcal{M}, \bar{\delta}}^{\text{n-so-cpa}}(\cdot) \leq \text{Adv}_{\text{NE}, \text{NG}, B, \mathcal{M}, \delta}^{\text{n-so-cpa}}(\cdot) .$$

The running time of B is about that of A plus the running time of DE.Kg plus the time to run DE.Enc on the messages that \mathcal{M} produces.

HN-SO-CPA security: For any (μ, d) -unpredictable message sampler \mathcal{M} , any function $\bar{\delta}$ such that \mathcal{M} is $\bar{\delta}$ -partially resamplable, and any adversary A , there are an adversary B that opens the same number of ciphertexts, another function $\bar{\delta}$, and another $(\text{NE.Guess}(\mu), d)$ -entropic, $\bar{\delta}$ -partially resamplable message sampler $\bar{\mathcal{M}}$ such that

$$\text{Adv}_{\bar{\text{NE}}, A, \mathcal{M}, \bar{\delta}}^{\text{hn-so-cpa}}(\cdot) \leq \text{Adv}_{\text{DE}, B, \bar{\mathcal{M}}, \bar{\delta}}^{\text{d-so-cpa2}}(\cdot) .$$

The running time of B is about that of A plus the running time of $\bar{\text{NE}}.\text{Kg}$ plus the time to run $\bar{\text{NE}}.\text{Dec}$ on v ciphertexts, where v is the number of messages that \mathcal{M} produces. The running time of $\bar{\mathcal{M}}$ is about that of \mathcal{M} plus the time to run NE.Enc on v messages.

Proof. For the first part, consider an arbitrary adversary A . Consider the adversary B in Fig. 12 attacking NE. Then game $\text{N-SO-CPA}_{\bar{\text{NE}}, \text{NG}}^{B, \mathcal{M}, \bar{\delta}}$ coincides with game $\text{N-SO-CPA}_{\bar{\text{NE}}, \text{NG}}^{A, \mathcal{M}, \delta}$, and thus $\text{Adv}_{\bar{\text{NE}}, \text{NG}, A, \mathcal{M}, \delta}^{\text{n-so-cpa}}(\cdot) = \text{Adv}_{\text{NE}, \text{NG}, B, \mathcal{M}, \delta}^{\text{n-so-cpa}}(\cdot)$.

For the second part, consider an arbitrary adversary A and a message sampler \mathcal{M} . Consider the following message sampler $\bar{\mathcal{M}}(1^k, pars)$. It parses $param$ as a

<p>Algorithm $B.\text{pg}(1^k)$ $param \leftarrow_s A.\text{pg}(1^k)$ $(pk_n, sk_n) \leftarrow_s \text{NE.Kg}(1^k)$ $pars \leftarrow (pk_n, sk_n, param)$ Return $pars$</p> <p>Algorithm $B.\text{cor}(pk_d, \mathbf{c}, pars)$ $(pk_n, sk_n, param) \leftarrow pars$ $pk \leftarrow (pk_d, pk_n)$ $(state, I) \leftarrow_s A.\text{cor}(pk, \mathbf{c}, param)$ $t \leftarrow (state, pars)$; Return (t, I)</p>	<p>Algorithm $B.g(t, \mathbf{y}^*, \mathbf{a}^*)$ $(state, pars) \leftarrow t$ $(pk_n, sk_n, param) \leftarrow pars$ For $i \leftarrow 1$ to \mathbf{y}^* do $\mathbf{m}[i] \leftarrow \text{NE.Dec}(sk_n, \mathbf{y}[i])$ $\omega \leftarrow_s A.g(state, \mathbf{m}^*, \mathbf{a}^*)$ Return ω</p> <p>Algorithm $B.f(z, pars)$ $(pk_n, sk_n, param) \leftarrow pars$ $t \leftarrow_s A.f(z, param)$; Return t</p>
<p>Algorithm $\overline{\mathcal{M}}(1^k, pars)$ $(pk_n, sk_n, param) \leftarrow pars$ $(\mathbf{m}, \mathbf{a}) \leftarrow_s \mathcal{M}(1^k, param)$ For $i = 1$ to \mathbf{m} do $(a, xk, N) \leftarrow \mathbf{a}[i]$ $\mathbf{y}[i] \leftarrow \text{NE.Enc}(pk_n, xk, N, \mathbf{m}[i])$ Return (\mathbf{y}, \mathbf{a})</p>	<p>Algorithm $\overline{\text{Rsm}}(1^k, \mathbf{y}^*, \mathbf{a}^*, I, pars)$ $(pk_n, sk_n, param) \leftarrow pars$ For $i = 1$ to \mathbf{y}^* do $\mathbf{m}^*[i] \leftarrow \text{NE.Dec}(sk_n, \mathbf{y}^*[i])$ $z \leftarrow_s \text{Rsm}(1^k, \mathbf{m}^*, \mathbf{a}^*, I, param)$ Return z</p>

Fig. 13: D-SO-CPA2 adversary B , constructed sampler $\overline{\mathcal{M}}$, and its partial resampling algorithm $\overline{\text{Rsm}}$ in the proof of Theorem 6.

triple $(pk_n, sk_n, param)$, where pk_n and sk_n are public and secret keys for NE. It then runs $\mathcal{M}(1^k, param)$ to generate (\mathbf{m}, \mathbf{a}) . Since \mathcal{M} is unpredictable, each $\mathbf{a}[i]$ can be parsed as (a_i, xk_i, N_i) . Now the “messages” of \mathcal{M} is the vector \mathbf{y} , where each $\mathbf{y}[i] = \text{NE.Enc}(pk_n, xk_i, N_i, \mathbf{m}[i])$, and the corresponding auxiliary information is still $\mathbf{a}[i]$. The code of $\overline{\mathcal{M}}$ is given in Fig. 13. Since \mathcal{M} is (μ, d) -unpredictable, $\overline{\mathcal{M}}$ is $(\text{NE.Guess}(\mu), d)$ -entropic. Let δ be a function such that \mathcal{M} is δ -partially resamplable. Let $\overline{\delta}(\mathbf{y}, pars)$ be the following function. It parses $pars$ as $(pk_n, sk_n, param)$, decrypts $\mathbf{m}[i] \leftarrow \text{NE.Dec}(sk_n, \mathbf{y}[i])$, and then returns $\delta(\mathbf{m}, param)$. Then $\overline{\mathcal{M}}$ is $\overline{\delta}$ -partially resamplable: given any δ -partial resampling algorithm Rsm for \mathcal{M} , we can construct a $\overline{\delta}$ -partial resampling algorithm $\overline{\text{Rsm}}$ for $\overline{\mathcal{M}}$ as in Fig. 13.

Now, consider the adversary B attacking DE as given in Fig. 13. It targets message sampler $\overline{\mathcal{M}}$, with respect to function $\overline{\delta}$. Initially, $B.\text{pg}(1^k)$ runs $param \leftarrow A(1^k)$, and then generates public and secret keys pk_n and sk_n for NE. It then outputs $pars \leftarrow (pk_n, sk_n, param)$. When $B.g$ receives its “messages” \mathbf{y}^* , it extracts the secret key sk_n from its state and decrypts $\mathbf{m}^*[i] \leftarrow \text{NE.Dec}(sk_n, \mathbf{y}^*[i])$, and then gives \mathbf{m}^* to $A.g$ together with the auxiliary information \mathbf{a}^* . Then game $\text{HN-CPA-REAL}_{\text{NE}}^{A, \mathcal{M}, \delta}$ coincides with game $\text{D-CPA2-REAL}_{\text{DE}}^{B, \overline{\mathcal{M}}, \overline{\delta}}$. Moreover, game $\text{HN-CPA-IDEAL}_{\text{NE}}^{A, \mathcal{M}, \delta}$ coincides with $\text{D-CPA2-IDEAL}_{\text{DE}}^{B, \overline{\mathcal{M}}, \overline{\delta}}$. Hence $\text{Adv}_{\text{NE}, A, \mathcal{M}, \delta}^{\text{hn-so-cpa}}(\cdot) \leq \text{Adv}_{\text{DE}, B, \overline{\mathcal{M}}, \overline{\delta}}^{\text{d-so-cpa2}}(\cdot)$. \square

NE1 ALONE IS ENOUGH. Constructions via NtD transform will be at least twice slower than NE1, because we need to run public primitives twice. But in the NPROM, NE1[H, LT] alone achieves both N-SO-CPA and HN-SO-CPA security simultaneously. In Theorem 7 below, we’ll show that NE1 is HN-SO-CPA secure. See the full version for the proof. We stress that for (μ, ∞) -unpredictable message samplers, NE1 allows the adversary to open as many ciphertexts as it wishes.

Theorem 7. Let LT be a lossy trapdoor function with lossiness τ . Let \mathcal{M} be a (μ, d) -unpredictable resamplable message sampler, and let δ be a function such that \mathcal{M} is δ -partially resamplable. Let NE1[H, LT] be as above. In the NPROM, for any adversary A opening at most d ciphertexts, there is an adversary D such that

$$\mathbf{Adv}_{\text{NE1}[H, LT], A, \mathcal{M}, \delta}^{\text{hn-so-cpa}}(k) \leq \frac{4q(k)}{2^k} + \frac{4q(k)v(k)}{2^{\mu(k)}} + \frac{v(k)(v(k) + 4q(k))}{2^{\tau(k)}} + 2\mathbf{Adv}_{LT, D}^{\text{ldf}}(k),$$

where $q(k)$ is the total number of random-oracle queries of A and \mathcal{M} , and $v(k)$ is the number of messages that \mathcal{M} produces. The running time of D is about that of A plus the time to run δ and an efficient δ -partial resampling algorithm of \mathcal{M} plus the time to run NE1[H, LT] to encrypt \mathcal{M} ’s messages. Adversary D makes at most q random-oracle queries.

5.3 Achieving HN-SO-CCA Security

In proving that NtD[NE, DE] achieves HN-SO-CPA security, we don’t need any property of the D-PKE scheme DE. This no longer holds for HN-SO-CCA. Indeed, consider a scheme DE_{bad} such that $\text{DE}_{\text{bad}}.\text{Enc}$ appends 0 to the ciphertexts, and $\text{DE}_{\text{bad}}.\text{Dec}$ ignores the last bit of the ciphertexts. An adversary thus can obtain the plaintexts by modifying the last bits of the ciphertexts, and querying those to the decryption oracle. Hence to obtain HN-SO-CCA, one has to exploit some property of DE. We’ll need DE to be *unique-ciphertext*, a property formalized by Bellare and Hoang [7].

Formally, a D-PKE scheme DE is *unique-ciphertext* if for every $k \in \mathbb{N}$, every $(pk, sk) \in [\text{DE.Kg}(1^k)]$, and every $m \in \{0, 1\}^*$, there is at most a string c such that $\text{DE.Dec}(sk, c) = m$. The D-PKE scheme DE_{bad} above is not unique-ciphertext. The unique-ciphertext property of DE ensures that if one modifies a ciphertext of NtD[NE, DE], the underneath ciphertext of NE will be changed.

Bellare and Hoang also show how to efficiently transform a D-PKE scheme DE to a unique-ciphertext one UE: in the decryption, we first recover the message, and then re-encrypt it and return \perp if the newly constructed ciphertext doesn’t match the given one. The transform UniqueCtx is given in Fig. 14. Note that this transform doesn’t affect the D-SO-CCA security of DE. Indeed, for any message sampler \mathcal{M} , any PT adversary A attacking $\text{UE} = \text{UniqueCtx}[\text{DE}]$, it’s trivial to construct another PT adversary B attacking DE such that $\mathbf{Adv}_{\text{UE}, B, \mathcal{M}}^{\text{d-so-cca}}(\cdot) \leq \mathbf{Adv}_{\text{DE}, A, \mathcal{M}}^{\text{d-so-cca}}(\cdot)$.

$\text{UE.Kg}(1^k)$	$\text{UE.Enc}(pk, m)$	$\text{UE.Dec}((pk, sk), c)$
$(pk, sk) \leftarrow_s \text{DE.Kg}(1^k)$ Return $(pk, (pk, sk))$	$c \leftarrow \text{DE.Enc}(pk, m)$ Return c	$m \leftarrow \text{DE.Dec}(sk, c)$ If $m \neq \perp$ then $c' \leftarrow \text{DE.Enc}(pk, m)$ If $c' \neq c$ then return \perp Return m

Fig. 14: **Unique-ciphertext D-PKE scheme** $\text{UE} = \text{UniqueCtx}[\text{DE}]$ **constructed from D-PKE scheme** DE .

Let DE be unique-ciphertext and D-SO-CCA secure D-PKE and NE be an N-SO-CCA secure, entropy-preserving nonce-based PKE. Then, Theorem 8 confirms $\text{NtD}[\text{NE}, \text{DE}]$ achieves both HN-SO-CCA and N-SO-CCA security simultaneously; the proof is in the full version. To instantiate DE , one can either apply the UniqueCtx transform on a D-SO-CCA secure D-PKE scheme, or directly use our construction DE2 in Section 3.3.

Theorem 8. Let NE be a nonce-based PKE as above, and let DE be a unique-ciphertext D-PKE scheme such that the ciphertext length of the former is a plaintext length of the latter. Let $\overline{\text{NE}} = \text{NtD}[\text{NE}, \text{DE}]$.

N-SO-CCA security: For any adversary A , any message sampler \mathcal{M} , any function $\bar{\delta}$ such that \mathcal{M} is $\bar{\delta}$ -partially resamplable, and any nonce generator NG , there is an adversary B such that

$$\text{Adv}_{\overline{\text{NE}}, \text{NG}, A, \mathcal{M}, \bar{\delta}}^{\text{n-so-cca}}(\cdot) \leq \text{Adv}_{\text{NE}, \text{NG}, B, \mathcal{M}, \delta}^{\text{n-so-cca}}(\cdot) .$$

The running time of B is about that of A plus the running time of DE.Kg plus the time to run DE.Enc on the messages that A produces, and the time to run DE.Dec on the decryption queries of A . Adversary B makes as many decryption-oracle queries as A .

HN-SO-CCA security: For any adversary A , any (μ, d) -unpredictable message sampler \mathcal{M} , and any function $\bar{\delta}$ such that \mathcal{M} is $\bar{\delta}$ -partially resamplable, there are an adversary B that opens the same number of ciphertexts, a function $\bar{\delta}$, and an $(\text{NE.Guess}(\mu), d)$ -entropic, $\bar{\delta}$ -partially resamplable message sampler $\overline{\mathcal{M}}$ such that

$$\text{Adv}_{\overline{\text{NE}}, A, \mathcal{M}, \bar{\delta}}^{\text{hn-so-cca}}(\cdot) \leq \text{Adv}_{\text{DE}, B, \overline{\mathcal{M}}, \bar{\delta}}^{\text{d-so-cca}}(\cdot) .$$

The running time of B is about that of A plus the running time of $\overline{\text{NE.Kg}}$ plus the time to run $\overline{\text{NE.Dec}}$ on $v + p$ ciphertexts, where v is the number of messages that \mathcal{M} produces and p is the number of A 's decryption-oracle queries. Adversary B makes as many decryption-oracle queries as A . The running time of $\overline{\mathcal{M}}$ is about that of \mathcal{M} plus the time to run NE.Enc on v messages.

Alternatively, we can use NE2 directly. In Theorem 9 below, we'll show that NE2 is HN-SO-CCA secure. See the full version for the proof.

Theorem 9. Let LT be a lossy trapdoor function with lossiness τ . Let \mathcal{M} be a (μ, d) -unpredictable message sampler, and let δ be a function such that \mathcal{M} is δ -partially resamplable. Let $\text{NE2}[H, \text{LT}]$ be as above. In the NPROM, for any adversary A opening at most d ciphertexts, there is an adversary D such that

$$\begin{aligned} & \mathbf{Adv}_{\text{NE2}[H, \text{LT}], A, \mathcal{M}, \delta}^{\text{hn-so-cca}}(k) \\ & \leq \frac{6Q(k)}{2^k} + \frac{4Q(k)v(k)}{2^{\mu(k)}} + \frac{v(k)(v(k) + 4Q(k))}{2^{\tau(k)}} + 2\mathbf{Adv}_{\text{LT}, D}^{\text{ltdf}}(k), \end{aligned}$$

where $q(k)$ is the total number of random-oracle queries of A and \mathcal{M} , $v(k)$ is the number of messages that \mathcal{M} produces, and $p(k)$ is the number of decryption queries of A , and $Q(k) = q(k) + 2p(k)$. The running time of D is about that of A plus the time to run δ and a δ -partial resampling algorithm of \mathcal{M} plus the time to run $\text{NE2}[H, \text{LT}]$ to encrypt \mathcal{M} ’s messages. Adversary D makes at most Q random-oracle queries.

Acknowledgments

We thank the Asiacrypt reviewers for their insightful comments. This work was supported in part by NSF awards CNS-1223623, CNS-1423566, and CNS-1553758 (CAREER), as well as the Glen and Susanne Culler Chair. The work of Viet Tung Hoang was done while at the University of Maryland, Georgetown University, and UC Santa Barbara.

References

1. B. Barak, Y. Dodis, H. Krawczyk, O. Pereira, K. Pietrzak, F.-X. Standaert, and Y. Yu. Leftover hash lemma, revisited. In P. Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 1–20, Santa Barbara, CA, USA, Aug. 14–18, 2011. Springer, Berlin, Germany.
2. M. Bellare, A. Boldyreva, and A. O’Neill. Deterministic and efficiently searchable encryption. In A. Menezes, editor, *CRYPTO 2007*, volume 4622 of *LNCS*, pages 535–552, Santa Barbara, CA, USA, Aug. 19–23, 2007. Springer, Berlin, Germany.
3. M. Bellare, Z. Brakerski, M. Naor, T. Ristenpart, G. Segev, H. Shacham, and S. Yilek. Hedged public-key encryption: How to protect against bad randomness. In M. Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 232–249, Tokyo, Japan, Dec. 6–10, 2009. Springer, Berlin, Germany.
4. M. Bellare, R. Dowsley, and S. Keelveedhi. How secure is deterministic encryption? In *Public-Key Cryptography — PKC 2015*, volume 9020 of *LNCS*, pages 52–73. Springer, Berlin, Germany, 2015.
5. M. Bellare, R. Dowsley, B. Waters, and S. Yilek. Standard security does not imply security against selective-opening. In D. Pointcheval and T. Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 645–662, Cambridge, UK, Apr. 15–19, 2012. Springer, Berlin, Germany.
6. M. Bellare, M. Fischlin, A. O’Neill, and T. Ristenpart. Deterministic encryption: Definitional equivalences and constructions without random oracles. In D. Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 360–378, Santa Barbara, CA, USA, Aug. 17–21, 2008. Springer, Berlin, Germany.

7. M. Bellare and V. T. Hoang. Resisting randomness subversion: Fast deterministic and hedged public-key encryption in the standard model. In *Advances in Cryptology—EUROCRYPT 2015*, volume 9057 of *LNCS*, pages 627–656. Springer, Berlin, Germany, 2015.
8. M. Bellare, V. T. Hoang, and S. Keelveedhi. Instantiating random oracles via UCEs. In R. Canetti and J. A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 398–415, Santa Barbara, CA, USA, Aug. 18–22, 2013. Springer, Berlin, Germany.
9. M. Bellare, D. Hofheinz, and S. Yilek. Possibility and impossibility results for encryption and commitment secure under selective opening. In A. Joux, editor, *EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 1–35, Cologne, Germany, Apr. 26–30, 2009. Springer, Berlin, Germany.
10. M. Bellare and B. Tackmann. Nonce-based cryptography: Retaining security when randomness fails. In *Advances in Cryptology—EUROCRYPT 2016*, 2016.
11. F. Böhl, D. Hofheinz, and D. Kraschewski. On definitions of selective opening security. In M. Fischlin, J. Buchmann, and M. Manulis, editors, *PKC 2012*, volume 7293 of *LNCS*, pages 522–539, Darmstadt, Germany, May 21–23, 2012. Springer, Berlin, Germany.
12. D. Cash, P. Grubbs, J. Perry, and T. Ristenpart. Leakage-abuse attacks against searchable encryption. In *CCS*, pages 668–679, 2015.
13. E. Fujisaki and T. Okamoto. How to enhance the security of public-key encryption at minimum cost. In H. Imai and Y. Zheng, editors, *PKC’99*, volume 1560 of *LNCS*, pages 53–68, Kamakura, Japan, Mar. 1–3, 1999. Springer, Berlin, Germany.
14. F. Heuer, E. Kiltz, and K. Pietrzak. Standard security does imply security against selective opening for markov distributions. Cryptology ePrint Archive, Report 2015/853, 2015. <http://eprint.iacr.org/2015/853>.
15. D. Hofheinz, V. Rao, and D. Wichs. Standard security does not imply indistinguishability under selective opening. Cryptology ePrint Archive, Report 2015/792, 2015. <http://eprint.iacr.org/2015/792>.
16. D. Hofheinz and A. Rupp. Standard versus selective opening security: Separation and equivalence results. In Y. Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 591–615, San Diego, CA, USA, Feb. 24–26, 2014. Springer, Berlin, Germany.
17. Y. Ishai, O. Pandey, and A. Sahai. Public-coin differing-inputs obfuscation and its applications. In *Theory of Cryptography Conference*, pages 668–697. Springer, 2015.
18. S. Kamara and J. Katz. How to encrypt with a malicious random number generator. In K. Nyberg, editor, *FSE 2008*, volume 5086 of *LNCS*, pages 303–315, Lausanne, Switzerland, Feb. 10–13, 2008. Springer, Berlin, Germany.
19. E. Kiltz, A. O’Neill, and A. Smith. Instantiability of RSA-OAEP under chosen-plaintext attack. In T. Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 295–313, Santa Barbara, CA, USA, Aug. 15–19, 2010. Springer, Berlin, Germany.
20. J. B. Nielsen. Separating random oracle proofs from complexity theoretic proofs: The non-committing encryption case. In M. Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 111–126, Santa Barbara, CA, USA, Aug. 18–22, 2002. Springer, Berlin, Germany.
21. C. Peikert and B. Waters. Lossy trapdoor functions and their applications. In R. E. Ladner and C. Dwork, editors, *40th ACM STOC*, pages 187–196, Victoria, British Columbia, Canada, May 17–20, 2008. ACM Press.
22. Y. Seurin. On the lossiness of the Rabin trapdoor function. In H. Krawczyk, editor, *PKC 2014*, volume 8383 of *LNCS*, pages 380–398, Buenos Aires, Argentina, Mar. 26–28, 2014. Springer, Berlin, Germany.

23. S. Yilek. Resetable public-key encryption: How to encrypt on a virtual machine. In J. Pieprzyk, editor, *CT-RSA 2010*, volume 5985 of *LNCS*, pages 41–56, San Francisco, CA, USA, Mar. 1–5, 2010. Springer, Berlin, Germany.