

What about Bob?

The Inadequacy of CPA Security for Proxy Reencryption *

Aloni Cohen

MIT, Cambridge MA, USA,
aloni@mit.edu

Abstract. In the simplest setting of proxy reencryption, there are three parties: Alice, Bob, and Polly (the proxy). Alice keeps some encrypted data that she can decrypt with a secret key known only to her. She wants to communicate the data to Bob, but not to Polly (nor anybody else). Using proxy reencryption, Alice can create a *reencryption key* that will enable Polly to reencrypt the data for Bob's use, but which will not help Polly learn anything about the data.

There are two well-studied notions of security for proxy reencryption schemes: security under chosen-plaintext attacks (CPA) and security under chosen-ciphertext attacks (CCA). Both definitions aim to formalize the security that Alice enjoys against both Polly and Bob.

In this work, we demonstrate that CPA security guarantees much less security against Bob than was previously understood. In particular, CPA security does not prevent Bob from learning Alice's secret key after receiving a single honestly reencrypted ciphertext. As a result, CPA security provides scant guarantees in common applications.

We propose security under honest reencryption attacks (HRA), a strengthening of CPA security that better captures the goals of proxy reencryption. In applications, HRA security provides much more robust security. We identify a property of proxy reencryption schemes that suffices to amplify CPA security to HRA security and show that two existing proxy reencryption schemes are in fact HRA secure.

Keywords: proxy reencryption · definitions · public-key cryptography

1 Introduction

Consider three parties: Alice, Bob, and Polly Proxy. Alice keeps encrypted data (created with a public key) that she can decrypt with a secret key known only

* Supported by Facebook Fellowship 2018, NSF GRFP, NSF MACS CNS-1413920, DARPA IBM W911NF-15-C-0236, and Simons Investigator Award Agreement Dated 6-5-12. We would like to thank Rio LaVigne, Akshay Degwekar, Shafi Goldwasser, Vinod Vaikuntanathan, and anonymous reviewers for their helpful feedback.

to her. She wants to communicate some of the data to Bob, but not to Polly (nor anybody else). Assuming Alice and Polly know Bob’s public key, how can she communicate the data to him?

If she is willing to entrust Bob with all her secrets, past and future, Alice might try to tell Bob her secret decryption key by encrypting it using Bob’s public key. We call this the *Trivial Scheme*. If she does not have such trust in Bob, Alice can instead decrypt the data, and reencrypt it using Bob’s public key. But what if Alice does not want to do the work of decrypting and reencrypting large amounts of data?

Proxy reencryption (PRE) provides an elegant solution: Alice creates and gives to Polly a *reencryption key* that will enable Polly to reencrypt her data under Bob’s public key for his use, but that will not reveal the data to Polly. Proxy reencryption guarantees that Bob can recover the data uncorrupted (correctness) and that Polly cannot learn anything about Alice’s data (security). The most widely-studied model of security for proxy reencryption is called *CPA security*, named after the corresponding notion from standard encryption on which it is based.

But what about Bob? As already observed, if we do not require any security against Bob, proxy reencryption is trivial: Alice uses the Trivial Scheme, simply sending Bob her encrypted secret key. This is undesirable, unsatisfying, and insufficient for a number of supposed applications of proxy reencryption (Section 2).

Surprisingly, *the Trivial Scheme is a CPA secure proxy reencryption scheme* when the public key encryption scheme used is circularly secure [6]! Bob completely learns Alice’s secret key, and circular security is used only to prove security against a malicious Polly. Furthermore, the CPA-security of any proxy reencryption scheme remains uncompromised if Polly attaches the reencryption key to every reencrypted ciphertext sent to Bob, even though this would enable Bob to decrypt messages encrypted under Alice’s public key (Section 3.1).

These “constructions” of CPA-secure proxy reencryption—original to this work—demonstrate the inadequacy of CPA security for proxy reencryption. If they had been observed previously, perhaps CPA security would not have gained the traction that it has.

Throughout this work, we use CPA (respectively, CCA and HRA) to refer to the security notion for proxy reencryption, and Enc-CPA (resp., Enc-CCA) to refer to the security notion for standard encryption. We restrict our attention to *unidirectional* proxy reencryption, where the reencryption key allows Alice’s ciphertexts to be reencrypted to Bob’s key, but not the reverse. In a bidirectional scheme, Bob—using his own secret key and Alice’s public key—is able compute the Alice-to-Bob reencryption key himself; thus a lack of security against Bob is inherent.

1.1 CPA and CCA Security of Proxy Reencryption

First considered by Blaze, Bleumer, and Strauss [5], proxy reencryption has received significant and continuous attention in the last decade, including defi-

nitions [22, 4, 10, 28], number-theoretical constructions [3, 26, 12], lattice-based constructions [19, 2, 30, 17], implementations [25, 20, 31, 8], and early success in program obfuscation [21, 11].

Adapting notions from standard encryption, this literature considers two main indistinguishability-based security notions for proxy reencryption: security under *chosen plaintext attacks (CPA)* [3] and *chosen ciphertext attacks (CCA)* [10]. While CCA security is considered the gold-standard, CPA security has received significant attention [4, 3, 21], especially in lattice-based constructions [19, 2, 30, 31]. CPA security has been used as a testing ground for new techniques for proxy reencryption and in settings where efficiency concerns make the added security of CCA undesirable.

We now briefly describe the definitions of CPA and CCA security for proxy reencryption, with the goal of communicating the underlying intuition. For this informal description, we restrict our attention to the limited three party setting of Alice, Bob, and Polly and strip away many of the complexities of the full definition. For a full definitions of CPA and CCA security, see Definitions 3 and the full version of this article [13], respectively.

Both notions are typically defined using a security game between an adversary and a challenger in which the adversary’s task is to distinguish between encryptions of two messages. Both notions allow the adversary to corrupt either Bob (learning sk_{bob}) or Polly (learning the reencryption key rk). CCA and CPA security differ in the additional power granted to the adversary.

CCA security grants the adversary access to two oracles:

- \mathcal{O}_{Dec} : The decryption oracle takes as input a ciphertext along with the public key of either Alice or Bob, and outputs the decryption of the ciphertext using the corresponding secret key.
- $\mathcal{O}_{\text{ReEnc}}$: The reencryption oracle takes as input a ciphertext ct_{alice} and outputs a reencrypted ciphertext ct_{bob} .

These oracles come with restrictions to prevent the adversary from simply reencrypting or decrypting the challenge ciphertext, adapting replayable chosen-ciphertext security of standard encryption (Enc-CCA) in the natural way.

CPA security of proxy reencryption, however, removes both oracles.¹ Why? First, to adapt chosen-plaintext security from standard encryption (Enc-CPA) to proxy reencryption, we must of course do away with \mathcal{O}_{Dec} . Secondly, it seems we must also remove $\mathcal{O}_{\text{ReEnc}}$: otherwise, by corrupting Bob it seems that the adversary can use the combination of $\mathcal{O}_{\text{ReEnc}}$ and sk_{bob} to simulate \mathcal{O}_{Dec} . Removing both decryption and reencryption oracles, according to [3], naturally extends the Enc-CPA security to proxy reencryption, yielding CPA security.

As we observe in this work, a natural definition is not always a good definition. Not only is the above intuition for removing $\mathcal{O}_{\text{ReEnc}}$ false (see full version

¹ This description is an oversimplification. In the many party setting, the adversary has access to a reencryption oracle which will reencrypt ciphertexts between two uncorrupted parties or between two corrupted parties, but not from an honest party to a corrupted party.

[13]), but CPA security as defined above guarantees little against a honest-but-curious Bob, even under *normal operation*. The definition only requires that the adversary will not win the game as long as it never sees any reencrypted ciphertexts. It guarantees nothing if Bob sees even a single reencrypted ciphertext. This vulnerability is not purely theoretical: in the CPA secure scheme of [31], Bob can recover Alice’s secret key with significant probability from a single reencrypted ciphertext (Theorem 4).

This makes CPA security ill-suited for the most commonly cited applications of proxy reencryption, including forwarding of encrypted email and single-writer, many-reader encrypted storage (Section 2). CPA security is inadequate for proxy reencryption and must be replaced.

1.2 Security Against Honest Reencryption Attacks

What minimal guarantees should proxy reencryption provide? First, it should offer security against a dishonest proxy Polly when Alice and Bob are honest and using the proxy reencryption as intended. Polly’s knowledge of a reencryption key from Alice to Bob (or vice versa) should not help her learn anything about the messages underlying ciphertexts encrypted under pk_{alice} or pk_{bob} . Such security against the corrupted proxy is guaranteed by CPA.

Second, proxy reencryption should offer security against a dishonest receiver Bob when Alice and Polly are honest and using the proxy reencryption as intended. Bob’s knowledge of *honestly reencrypted ciphertexts* (that were honestly generated to begin with) should not help him learn anything about the messages underlying other ciphertexts encrypted under pk_{alice} that have not been reencrypted. As we show in this work, such security against the corrupted receiver is not guaranteed by CPA.

Generalizing these dual guarantees to many possibly colluding parties, we want security as long as the adversary only sees honestly reencrypted ciphertexts. In Section 4, we formalize this notion as proxy reencryption security against *honest reencryption attacks (HRA)*. Recall that CCA security provides the adversary with both \mathcal{O}_{Dec} and $\mathcal{O}_{\text{ReEnc}}$ while CPA provides neither oracle. In contrast, HRA security provides the adversary with a restricted reencryption oracle which will only reencrypt honestly generated ciphertexts.

By guaranteeing security of both kinds described above, HRA is a strengthening of CPA security that better captures our intuitions for security of proxy reencryption. Furthermore, HRA guarantees more meaningful security in the most common applications of proxy reencryption (Section 4.1). HRA security is an appropriate goal when developing new techniques for proxy reencryption and in settings where full CCA security is undesirable or out of reach.

Security of Existing Schemes. Can we construct a proxy reencryption scheme that is HRA secure? HRA security is a strict strengthening of CPA security, so it is not immediately clear that any existing constructions are HRA secure without redoing the proofs from scratch. Indeed, the CPA secure scheme of [31] is not HRA secure (Theorem 4).

In Section 5, we identify a property—*reencryption simulatability*—which is sufficient to boost CPA security to HRA security. Very roughly, reencryption simulatability means that reencrypted ciphertexts resulting from computing $\text{ReEnc}(\text{rk}_{\text{alice} \rightarrow \text{bob}}, \text{ct}_{\text{alice}})$ can be simulated without knowledge of the secret key sk_{alice} (but with knowledge of the plaintext message \mathbf{m}). Reencryption simulatability allows a reduction with access to the CPA oracles to efficiently implement the honest reencryption oracle, thereby reducing HRA security to CPA security.

We then examine the simple construction of proxy reencryption from any fully-homomorphic encryption [19], and the pairing-based construction of [4]. In the first case, if the fully-homomorphic encryption secure is circuit private, then the resulting proxy reencryption scheme is reencryption simulatable. In the second case, rerandomizing reencrypted ciphertexts suffices for reencryption simulation.²

1.3 Related Work

The below mentioned works are just the most directly relevant. There is by now an extensive research literature on proxy reencryption, presenting a zoo of definitions. There have been three main approaches to defining security: CPA, CCA, and (to a much lesser extent) obfuscation-based. The CPA notion, in one form or another, is by far the most well studied. In this work, we make the deliberate choice to address the core CPA definition, not to present an ultimate definition of security for proxy reencryption nor to address the vast array of different criticisms or strengthenings of CPA security that have been or may be considered. We hope that doing so will make these ideas more understandable and adaptable.

RIND-CPA Security In concurrent and independent work defining and constructing forward-secure proxy reencryption, Derler, Krenn, Lorünser, Ramacher, Slamanig, and Striecks identify the same problem with CPA security as discussed in this work [14, Definition 14]. As in our work, they address the problem with CPA security by defining a new security notion—RIND-CPA security—which expands the power of the adversary. They additionally separate RIND-CPA and CPA security with a construction that is essentially our Concatenation Scheme.

However, this is where the resemblance between [14] and our work ends. In the RIND-CPA game offered by [14], the adversary gets access to an reencryption oracle that works on all inputs (not just honestly generated ones), but only before the challenge ciphertext is generated.³ In contrast, HRA allows reencryption both before and after the challenge, but only for honestly generated ciphertexts.

RIND-IND is inadequate as a replacement for CPA security in the research literature: its usefulness in applications is unclear, and it appears too strong

² While we don't examine every pairing-based construction of proxy reencryption, we suspect that rerandomizing reencryption will suffice for reencryption simulation in many, if not all.

³ The full version [13] discusses the related definition of IND-CCA_{0,1} security from [28]

to provide a useful testing ground for the development of new techniques for constructing proxy reencryption.

In the course of normal operation of a proxy reencryption in applications, an adversarial party will typically see many reencrypted ciphertexts. These ciphertexts may come at any time—both before and after other ciphertexts whose contents should remain secret. HRA is meaningful in many such applications—many more than CPA security. But because RIND-CPA restricts the reencryption oracle to the period before the challenge ciphertext, its usefulness in applications is not clear.

The challenge of proving CCA security for encryption (proxy or otherwise) is demonstrating that an adversary cannot use dishonestly generated, malformed ciphertexts to win in the security game. In this respect, RIND-CPA security is much more akin to CCA security than to CPA security. HRA, on the other hand, makes minimal assumptions about the distribution of plaintext messages by allowing the adversary to choose messages itself, just as in Enc-CPA for standard encryption.

Appendix B discusses RIND-CPA security in more depth, expanding on the arguments above and proving that RIND-CPA and HRA security are incomparable.

Subsequent Work Two subsequent works continue the study of HRA secure proxy reencryption. Fuchsbauer, Kamath, Klein, and Pietrzak study CPA and HRA secure proxy reencryption in an adaptive corruption model [18]. As in our work, they prove the HRA security of their construction by first proving CPA security and then lifting it to full HRA security using a version of reencryption simulatability.

More recently, Dottling and Nishimaki study the problem of converting ciphertexts between different public-key encryption schemes, a problem they call universal proxy reencryption [15]. They define security by extending HRA security to the universal setting. [15] extends Theorem 5 to show that a computational version of reencryption simulatability suffices to lift CPA to HRA security. However, they prove HRA security directly rather, finding that proving computational reencryption simulatability is not much more simple than proving HRA security itself.

Other Related Work Our dual-guarantee conception of proxy reencryption security mirrors the security requirements of what Ivan and Dodis call CPA security [22]. Their notion differs substantially from what is now referred to by that name. The [22] conception of CPA security is only defined in a proof in the appendix of that work and seems to have been completely overlooked by the later works proposing the modern notion of CPA security (beginning with [4] and then in its present form in [3]). If, however, Ivan and Dodis had undertaken to revisit proxy reencryption after [3], they might have proposed a security definition similar to HRA.

In [28], Nuñez, Agudo, and Lopez provide a parameterized family of CCA-type attack models for proxy reencryption. Their weakest model corresponds to CPA security and their strongest to full CCA security. That work is partially a response to a claimed construction of CCA-1 secure proxy reencryption in a security model that does not allow reencryption queries. [28] provide an attack on the construction in the presence of a reencryption oracle consisting of carefully constructed, dishonestly generated queries which leak the reencryption key. They do not consider restricting the reencryption oracle in the security game to honestly generated ciphertexts. We discuss [28] further in the full version [13].

Finally, a parallel line of work initiated by Hohenberger, Rothblum, Shelat, Vaikuntanathan which studies proxy reencryption using an obfuscation-based definition that is incomparable to CPA security [21]. Their definition requires that the functionality of the obfuscated reencryption circuit be statistically close to that of the ideal reencryption functionality: namely, that $\text{ReEnc}(rk_{i \rightarrow j}, \text{Enc}(\mathbf{pk}_i, \mathbf{m})) \approx_s \text{Enc}(\mathbf{pk}_j, \mathbf{m})$. Thus the definition of [21] (and even the relaxed correctness found in [11]) imply *reencryption simulatability* defined in Section 5.

1.4 Organization

We begin by discussing applications of proxy reencryption and identifying the weaknesses of CPA security in those applications (Section 2). Then we present the existing CPA security definition and further demonstrate its weaknesses with two new schemes: the Trivial Scheme and Concatentation Scheme (Section 3). We propose a new security notion to overcome those weaknesses: security against honest reencryption attacks (HRA) (Section 4). We examine the relationship between CPA and HRA security and the HRA security (or insecurity) of existing reencryption schemes (Section 5). The appendix provides additional discussion of the Trivial Scheme (Appendix A) and comparison between HRA and RIND-CPA security (Appendix B). The full version provides additional discussion of CCA security [13].

2 Insufficiency of CPA Security for Applications

In Section 3, we recall the definition of CPA security of proxy reencryption from [3] and formalize the Trivial Scheme from the introduction satisfying the notion. In the Trivial Scheme, Bob learns Alice’s secret key after receiving a single reencrypted ciphertext.

We are faced with a choice: accept the existing definition of CPA security, or reject it and seek a definition that better captures our intuitions. In support of the latter, we describe a number of applications of proxy reencryption proposed in the literature for which CPA security (as implemented by the Trivial Scheme) is potentially unsatisfactory, but for which full CCA security may not always be necessary.⁴ We revisit these applications in Section 4.1 after proposing our new security notion.

⁴ We might also appeal for support to [22], the only paper in the proxy reencryption literature of which we are aware adopting a security definition providing a reencryp-

Encrypted Email Forwarding [5, 23, 4]. Forwarding of encrypted email without requiring the sender’s participation might be desirable for temporary delegation during a vacation [23] or for spam filtering [4]. Does the Trivial Scheme suffice? The Trivial Scheme enables Bob, the receiver of Alice’s forwarded (and reencrypted) email, to recover Alice’s secret key. If Alice trusts Bob enough to use the Trivial Scheme, she could instead reveal her secret key. The Trivial Scheme might be preferable in very specific trust or interaction models, but it does not offer meaningful security against Bob if Alice only wishes to forward a subset of emails (for example, from particular senders or during a specific time period).

Key Escrow [22]. Similar to email forwarding, Ivan and Dodis describe the application of key escrow as follows: “The problem is to allow the law enforcement agency to read messages encrypted for a set of users, for a limited period of time, without knowing the users secrets. The solution is to locate a key escrow agent between the users and the law enforcement agency, such that it controls which messages are read by the law enforcement agencies.” As in email forwarding, the “for a limited period of time” requirement suggests that Ivan and Dodis would not have been satisfied with the Trivial Scheme.⁵

Single-Writer, Many-Reader Encrypted Storage [4, 24, 25, 31]. Under different monikers (including DRM and publish/subscribe systems), these works describe systems in which a single privileged writer encrypts data and determines an access control policy for readers. A semi-honest proxy server is entrusted with reencryption keys and is tasked with enforcing the access control policy. Whether the Trivial Scheme suffices for these applications depends on what sort of access control policies are envisioned. If the access is *all or nothing* (i.e., all readers may access all data), the Trivial Scheme suffices; if the access is *fine grained* (i.e., each reader may access only a specific subset of the data), then it does not. Existing works are often unclear on which is envisioned.

For completeness, we mention that CPA security does suffice for two important applications of proxy reencryption: namely, key rotation for encrypted cloud storage [7, 16] and fully homomorphic encryption [19].

tion oracle without a decryption oracle. One could look to the originators of proxy reencryption for guidance, but the shortcoming we identify does not manifest in the original setting of [5] (there is only Alice and Bob; there is no Proxy). It is therefore of little help.

⁵ Note that Ivan and Dodis do not adopt the CPA definition used elsewhere, but a definition much closer to our own. There is no gap between their security guarantees and the requirements of their briefly-described application.

Though primarily focused on the setting where the key escrow agent enforces the limited time requirement by eventually refusing to reencrypt, [22] considers the possibility of dividing time into epochs and enforcing the time limitation technically. Such a proxy reencryption is called *temporary* in [4]. We do not discuss temporary proxy reencryption further.

3 Security Against Chosen Plaintext Attacks

In this section, we recall the definition of CPA security for proxy reencryption and illustrate its shortcomings. We begin with the definitions of syntax, correctness, and CPA security from [3, Definition 2.2] (with minor changes in notation and presentation, and the change noted in Remark 1 at the end of this subsection). The syntax and correctness requirements are common to CPA, HRA, and CCA security.

For the sake of concreteness, simplicity, and brevity, we restrict the discussion to unidirectional, single-hop schemes. In multi-hop schemes, reencryption keys $\text{rk}_{A \rightarrow B}$ and $\text{rk}_{B \rightarrow C}$ can be used to reencrypt a ciphertext ct_A from pk_A to pk_C . In single-hop schemes, they cannot. Single-hop or multi-hop schemes each have their benefits and drawbacks, and works typically focus on one or the other notion.⁶ To the best of our knowledge, our observations and results can all be adapted to the multi-hop setting.

Definition 1 (Proxy Reencryption: Syntax [3]). *A proxy reencryption scheme for a message space \mathcal{M} is a tuple of algorithms $\text{PRE} = (\text{Setup}, \text{KeyGen}, \text{ReKeyGen}, \text{Enc}, \text{ReEnc}, \text{Dec})$:*

$\text{Setup}(1^\lambda) \rightarrow \text{pp}$. *On input security parameter 1^λ , the setup algorithm outputs the public parameters pp .*

$\text{KeyGen}(\text{pp}) \rightarrow (\text{pk}, \text{sk})$. *On input public parameters, the key generation algorithm outputs a public key pk and a secret key sk . For ease of notation, we assume that both pk and sk include pp and refrain from including pp as input to other algorithms.*

$\text{ReKeyGen}(\text{sk}_i, \text{pk}_j) \rightarrow \text{rk}_{i \rightarrow j}$. *On input a secret key sk_i and a public key pk_j , where $i \neq j$, the reencryption key generation algorithm outputs a reencryption key $\text{rk}_{i \rightarrow j}$.*

$\text{Enc}(\text{pk}_i, \mathbf{m}) \rightarrow \text{ct}_i$. *On input a public key pk_i and a message $\mathbf{m} \in \mathcal{M}$, the encryption algorithm outputs a ciphertext ct_i .*

$\text{ReEnc}(\text{rk}_{i \rightarrow j}, \text{ct}_i) \rightarrow \text{ct}_j$. *On input a reencryption key from i to j $\text{rk}_{i \rightarrow j}$ and a ciphertext ct_i , the reencryption algorithm outputs a ciphertext ct_j or the error symbol \perp .*

⁶ The literature is divided about whether “single-hop” is merely a correctness property (i.e., able to reencrypt at least once, but agnostic about whether reencrypting more than once is possible) or if it is also a security property (i.e., a ciphertext can be reencrypted once, but never twice). This distinction manifests in the security definition. In works that consider only single-hop correctness [4, 3, 21, 28], the oracle $\mathcal{O}_{\text{ReKeyGen}}$ in the security game will not accept queries from honest users to corrupted users (i.e., (i, j) such that $i \in \text{Hon}$ and $j \in \text{Cor}$). We adopt this formalism in Definitions 3 and 5 for simplicity of presentation only.

In works that consider single-hop security [26, 12, 17], the oracle will answer such queries, but the challenge ciphertext must be encrypted under the key of an honest user i^* for which no such reencryption key was generated (which can be formalized in a number of ways).

This work adopts the simplest model, requiring only one hope of correctness, but neither requiring nor forbidding additional functionality.

$\text{Dec}(\text{sk}_j, \text{ct}_j) \rightarrow \mathbf{m}$. Given a secret key sk_j and a ciphertext ct_j , the decryption algorithm outputs a message $\mathbf{m} \in \mathcal{M}$ or the error symbol \perp .

Definition 2 (Proxy Reencryption: Correctness [3]). A proxy reencryption scheme PRE is correct with respect to message space \mathcal{M} if for all $\lambda \in \mathbb{N}$, $\text{pp} \leftarrow \text{Setup}(1^\lambda)$, and $\mathbf{m} \in \mathcal{M}$:

1. for all $(\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(\text{pp})$:

$$\text{Dec}(\text{sk}, \text{Enc}(\text{pk}, \mathbf{m})) = \mathbf{m}.$$

2. for all $(\text{pk}_i, \text{sk}_i), (\text{pk}_j, \text{sk}_j) \leftarrow \text{KeyGen}(\text{pp})$, $\text{rk}_{i \rightarrow j} \leftarrow \text{ReKeyGen}(\text{sk}_i, \text{pk}_j)$:

$$\text{Dec}(\text{sk}_j, \text{ReEnc}(\text{rk}_{i \rightarrow j}, \text{Enc}(\text{pk}_i, \mathbf{m}))) = \mathbf{m}.$$

Security is modeled by a game played by an adversary \mathcal{A} . \mathcal{A} has the power to corrupt a set of users Cor (learning their secret keys) while learning only the public keys for a set of honest users Hon . Additionally, \mathcal{A} may reencrypt ciphertexts (either by getting a reencryption key or calling a reencryption oracle) between pairs of users in Hon or in Cor , or from Cor to Hon , but not from Hon to Cor . This is in contrast to the simplified three-party setting discussed in the introduction, where the \mathcal{A} could not reencrypt whatsoever.

Definition 3 (Proxy Reencryption: Security Game for Chosen Plaintext Attacks (CPA) [3]). Let λ be the security parameter and \mathcal{A} be an oracle Turing machine. The CPA game consists of an execution of \mathcal{A} with the following oracles. The game consists of three phases, which are executed in order. Within each phase, each oracle can be executed in any order, $\text{poly}(\lambda)$ times, unless otherwise specified.

Phase 1:

Setup: The public parameters are generated and given to \mathcal{A} . A counter numKeys is initialized to 0, and the sets Hon (of honest, uncorrupted indices) and Cor (of corrupted indices) are initialized to be empty. This oracle is executed first and only once.

Uncorrupted Key Generation: Obtain a new key pair $(\text{pk}_{\text{numKeys}}, \text{sk}_{\text{numKeys}}) \leftarrow \text{KeyGen}(\text{pp})$ and give $\text{pk}_{\text{numKeys}}$ to \mathcal{A} . The current value of numKeys is added to Hon and numKeys is incremented.

Corrupted Key Generation: Obtain a new key pair $(\text{pk}_{\text{numKeys}}, \text{sk}_{\text{numKeys}}) \leftarrow \text{KeyGen}(\text{pp})$ and given to \mathcal{A} . The current value of numKeys is added to Cor and numKeys is incremented.

Phase 2: For each pair $i, j \leq \text{numKeys}$, compute the reencryption key $\text{rk}_{i \rightarrow j} \leftarrow \text{ReKeyGen}(\text{sk}_i, \text{pk}_j)$.

Reencryption Key Generation $\mathcal{O}_{\text{ReKeyGen}}$: On input (i, j) where $i, j \leq \text{numKeys}$, if $i = j$ or if $i \in \text{Hon}$ and $j \in \text{Cor}$, output \perp . Otherwise return the reencryption key $\text{rk}_{i \rightarrow j}$.

Reencryption $\mathcal{O}_{\text{ReEnc}}$: On input (i, j, ct_i) where $i, j \leq \text{numKeys}$, if $i = j$ or if $i \in \text{Hon}$ and $j \in \text{Cor}$, output \perp . Otherwise return the reencrypted ciphertext $\text{ReEnc}(\text{rk}_{i \rightarrow j}, \text{ct}_i)$.

Challenge Oracle: On input $(i, \mathbf{m}_0, \mathbf{m}_1)$ where $i \in \text{Hon}$ and $\mathbf{m}_0, \mathbf{m}_1 \in \mathcal{M}$, sample a bit $b \leftarrow \{0, 1\}$ uniformly at random, and return the challenge ciphertext $\text{ct}^* \leftarrow \text{Enc}(\text{pk}_i, \mathbf{m}_b)$. This oracle can only be queried once.

Phase 3:

Decision: On input a bit $b' \in \{0, 1\}$, return win if $b = b'$.

The CPA advantage of \mathcal{A} is defined as

$$\text{Adv}_{\text{cpa}}^{\mathcal{A}}(\lambda) = \Pr[\text{win}],$$

where the probability is over the randomness of \mathcal{A} and the oracles in the CPA game.

Definition 4 (Proxy Reencryption: CPA Security [3]). Given a security parameter 1^λ , a proxy reencryption scheme is CPA secure if for all probabilistic polynomial-time adversaries \mathcal{A} , there exists a negligible function negl such that

$$\text{Adv}_{\text{cpa}}^{\mathcal{A}}(\lambda) < \frac{1}{2} + \text{negl}(\lambda)$$

Remark 1. Another definitional subtlety of proxy reencryption worth discussing affects not just CPA security, but HRA and CCA security as well. Consider the specification of $\mathcal{O}_{\text{ReKeyGen}}$ and $\mathcal{O}_{\text{ReEnc}}$ in Definition 3. As defined, the reencryption keys $\text{rk}_{i \rightarrow j}$ are *persistent*: the same key is used each time a pair (i, j) is queried. This follows [3, Definition 2.5] and [2, 17], though we find our formalization somewhat simpler.

Contrast this with [3, Definition 2.2] in which reencryption keys are *ephemeral*: they are generated afresh each time either oracle is invoked on the same pair (i, j) . [7, 30, 10] similarly use ephemeral keys in their definitions. In the remaining papers we examined, it was less clear whether reencryption keys were ephemeral or persistent.

Neither definition implies the other; any scheme secure with persistent keys can be modified into one that is insecure with ephemeral keys, and vice-versa. The definitions, however, are not in serious tension; any scheme secure with persistent keys and having deterministic ReKeyGen is also secure with ephemeral keys, and ReKeyGen can in general be derandomized using pseudorandom functions. Of course, one can easily define a hybrid notion stronger than both by allowing the adversary to specify for each query whether it wants to use reencryption keys that are new or old.

We adopt the persistent formalization as it better captures ‘typical’ use. To the best of our knowledge, all claims in this work can be adapted to the ephemeral setting.

Remark 2. The power of the adversary above can be strengthened by allowing *adaptive* corruptions instead of dividing the game into phases. Our definitions of CPA and HRA security follow the convention of [3] primarily because it is most common in the research literature. For an examination of CPA and HRA security in the adaptive setting, see the subsequent work of Fuchsbauer, Kamath, Klein, and Pietrzak [18]. Adaptive-secure, bidirectional, CCA secure proxy reencryption has been studied in [10, 28].

3.1 Concatenation Scheme and Trivial Scheme

The weakness of CPA security lies in the specification of $\mathcal{O}_{\text{ReEnc}}$, which does not reencrypt any ciphertexts from honest to corrupt users. Said another way, $\mathcal{O}_{\text{ReEnc}}$ reencrypts between only those pairs keys for which $\mathcal{O}_{\text{ReKeyGen}}$ outputs a reencryption key (rather than returning \perp). We now describe two schemes that are CPA secure, but are insecure against a dishonest receiver of reencrypted ciphertexts. In both schemes, a single ciphertext reencrypted from an honest index to a corrupted index enables the decryption of messages encrypted under the honest index's public key.

Concatenation Scheme Let $\text{PRE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec}, \text{ReKeyGen}, \text{ReEnc})$ be a CPA-secure proxy reencryption scheme. Define a new scheme by modifying only reencryption and decryption:

$$\begin{aligned} \text{ReEnc}'(\text{rk}, \text{ct}) &:= \text{ReEnc}(\text{rk}, \text{ct}) \parallel \text{rk} \\ \text{Dec}'(\text{sk}, \text{ct}) &:= \begin{cases} \text{Dec}(\text{sk}, \text{ct}^1) & \text{if } \text{ct} = \text{ct}^1 \parallel \text{ct}^2 \\ \text{Dec}(\text{sk}, \text{ct}) & \text{otherwise} \end{cases} \end{aligned}$$

Theorem 1. *Let $\text{PRE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec}, \text{ReKeyGen}, \text{ReEnc})$ be a CPA-secure proxy reencryption scheme. The corresponding Concatenation Scheme $\text{PRE}' = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec}', \text{ReKeyGen}, \text{ReEnc}')$ is a CPA-secure proxy reencryption scheme.*

Proof. For any probabilistic, polynomial-time algorithm \mathcal{A}' (the CPA adversary against PRE'), we construct an efficient algorithm \mathcal{A} such that $\text{Adv}_{\text{cpa}}^{\mathcal{A}} = \text{Adv}_{\text{cpa}}^{\mathcal{A}'}$. By the CPA security of PRE , this advantage is negligible, completing the proof.

\mathcal{A} runs \mathcal{A}' and simulates the CPA security game for PRE' (if \mathcal{A}' does not follow the specification of the game, \mathcal{A} simply aborts). Except for calls to $\mathcal{O}_{\text{ReEnc}}$, all oracle calls by \mathcal{A}' are passed along unaltered by \mathcal{A} , along with their responses.

\mathcal{A} begins Phase 2 by requesting all admissible reencryption keys $\text{rk}_{i \rightarrow j}$ from its own reencryption key generation oracle. To answer oracle calls from \mathcal{A}' to $\mathcal{O}_{\text{ReEnc}}$, \mathcal{A} first queries its own reencryption oracle, which returns ct^1 . If $\text{ct}^1 = \perp$, then \mathcal{A}' returns \perp . Otherwise, \mathcal{A}' concatenates the appropriate reencryption key rk to form the new ciphertext $\text{ct} = \text{ct}^1 \parallel \text{rk}$. This is possible because if $\text{ct}^1 \neq \perp$, then \mathcal{A} is able to get the corresponding reencryption key at the beginning of Phase 2.

\mathcal{A} perfectly implements the CPA security game for PRE' , and \mathcal{A}' wins that game if and only if \mathcal{A} wins the corresponding game for PRE . Therefore, $\text{Adv}_{\text{cpa}}^{\mathcal{A}} = \text{Adv}_{\text{cpa}}^{\mathcal{A}'}$. Finally, the running time of \mathcal{A} is polynomially related to that of \mathcal{A}' .

While the Concatenation Scheme builds upon any CPA-secure proxy re-encryption scheme, the Trivial Scheme presented next makes use of public-key encryption enjoying *circular security*. Informally, circular security guarantees that encryptions of messages that are a function of the secret key(s) are as secure as encryptions of messages that are independent of the secret key(s), a security property that does not follow from standard semantic security.

In the Trivial Scheme, the reencryption key from party i to j is simply $\text{rk}_{i \rightarrow j} = \text{Enc}(\text{pk}_j, \text{sk}_i)$. CPA security of the resulting proxy reencryption scheme requires security against an adversary who has both $\text{rk}_{i \rightarrow j}$ and $\text{rk}_{j \rightarrow i}$. This requires that the underlying encryption scheme is circular secure.

Because existing definitions and constructions of circular secure encryption schemes based on standard assumptions (e.g., [6] from DDH) require a bound on the total number of public keys n , the corresponding Trivial Scheme will only satisfy a bounded-key variant of CPA security. Any circular secure encryption scheme without this limitation would yield a Trivial Scheme secure according to Definition 4. We defer the definitions of circular security, bounded-key CPA security, and the proof of Theorem 2 to Appendix A.

Trivial Scheme Let $(\text{KeyGen}_{\text{circ}}, \text{Enc}_{\text{circ}}, \text{Dec}_{\text{circ}})$ be an n -way circular-secure encryption scheme. Let $\text{Setup} \equiv \perp$, $\text{KeyGen} \equiv \text{KeyGen}_{\text{circ}}$; $\text{Enc} \equiv \text{Enc}_{\text{circ}}$;

$$\begin{aligned} \text{ReKeyGen}(\text{sk}_i, \text{pk}_j) &:= \text{Enc}_{\text{circ}}(\text{pk}_j, \text{sk}_i) \\ \text{ReEnc}(\text{rk}_{i \rightarrow j}, \text{ct}_i) &:= \text{ct}_i \parallel \text{rk}_{i \rightarrow j} \\ \text{Dec}(\text{sk}, \text{ct}) &:= \begin{cases} \text{Dec}_{\text{circ}}(\text{Dec}_{\text{circ}}(\text{sk}, \text{ct}^2), \text{ct}^1) & \text{if } \text{ct} = \text{ct}^1 \parallel \text{ct}^2 \\ \text{Dec}_{\text{circ}}(\text{sk}, \text{ct}) & \text{otherwise} \end{cases} \end{aligned}$$

Theorem 2. *Let $(\text{KeyGen}_{\text{circ}}, \text{Enc}_{\text{circ}}, \text{Dec}_{\text{circ}})$ be an n -way circular-secure encryption scheme. The corresponding Trivial Scheme PRE is an n -way CPA secure proxy reencryption scheme.*

4 Security Against Honest Reencryption Attacks

We seek a definition of security which holds as long as the adversary only sees honestly reencrypted ciphertexts, unless the set of corrupt parties can trivially violate security (by some chain of reencryption keys from an uncorrupted public key to a corrupted public key).

We term this notion security against *honest reencryption attacks (HRA)*. To formalize it, we model the ability of an adversary to see honest reencryptions by granting it access to a modified reencryption oracle $\mathcal{O}_{\text{ReEnc}}$. Instead of taking a ciphertext as input, the modified $\mathcal{O}_{\text{ReEnc}}$ takes as input a reference to a previously generated ciphertext (either as the output of \mathcal{O}_{Enc} or $\mathcal{O}_{\text{ReEnc}}$ itself).

To implement such an oracle, we introduce to the security game a key-value store \mathcal{C} as additional state: the values are ciphertexts ct which are keyed by a pair of integers (i, k) , where i denotes the index of the key pair $(\text{pk}_i, \text{sk}_i)$ under which ct was (re)encrypted, and k is a unique index assigned to ct .

As described, this new $\mathcal{O}_{\text{ReEnc}}$ admits a trivial attack: simply reencrypt the challenge ciphertext to a corrupted key and decrypt. To address this issue, we adapt an idea from [10]’s definition of CCA security: we rule out the trivial attack by storing a set Deriv of ciphertexts derived from the challenge by reencrypting, and rejecting queries to $\mathcal{O}_{\text{ReEnc}}$ for ciphertexts in Deriv and a corrupted target key. We might have instead chosen to forbid any reencryptions of the challenge ciphertext, even between uncorrupted keys. Though this would have simplified the definition, it would have been unsatisfactory. For example, in the single-writer, many-reader encrypted storage application the contents of a message \mathbf{m} that gets reencrypted from Alice to Charlie should be hidden from Bob.

We now present the honest reencryption attacks security game. The game is similar to the CPA security game with some modifications to Setup, Challenge, and $\mathcal{O}_{\text{ReEnc}}$, and the addition of an Enc oracle \mathcal{O}_{Enc} to Phase 2. \mathcal{O}_{Enc} may be executed $\text{poly}(\lambda)$ times and in any order relative to the other oracles in Phase 2. For the sake of clarity we reproduce the full game below and mark the modified oracles with a \star .

Definition 5 (Proxy Reencryption: Security Game for Honest Reencryption Attacks (HRA)). *Let λ be the security parameter and \mathcal{A} be an oracle Turing machine. The HRA game consists of an execution of \mathcal{A} with the following oracles.*

Phase 1:

\star Setup: *The public parameters are generated and given to \mathcal{A} . A counter numKeys is initialized to 0, and the sets Hon (of honest, uncorrupted indices) and Cor (of corrupted indices) are initialized to be empty.*

Additionally the following are initialized: a counter numCt to 0, a key-value store \mathcal{C} to empty, and a set Deriv to be empty. This oracle is executed first and only once.

Uncorrupted Key Generation: *Obtain a new key pair $(\text{pk}_{\text{numKeys}}, \text{sk}_{\text{numKeys}}) \leftarrow \text{KeyGen}(\text{pp})$ and give $\text{pk}_{\text{numKeys}}$ to \mathcal{A} . The current value of numKeys is added to Hon and numKeys is incremented.*

Corrupted Key Generation: *Obtain a new key pair $(\text{pk}_{\text{numKeys}}, \text{sk}_{\text{numKeys}}) \leftarrow \text{KeyGen}(\text{pp})$ and given to \mathcal{A} . The current value of numKeys is added to Cor and numKeys is incremented.*

Phase 2: For each pair $i, j \leq \text{numKeys}$, compute the reencryption key $\text{rk}_{i \rightarrow j} \leftarrow \text{ReKeyGen}(\text{sk}_i, \text{pk}_j)$.

Reencryption Key Generation $\mathcal{O}_{\text{ReKeyGen}}$: *On input (i, j) where $i, j \leq \text{numKeys}$, if $i = j$ or if $i \in \text{Hon}$ and $j \in \text{Cor}$, output \perp . Otherwise return the reencryption key $\text{rk}_{i \rightarrow j}$.*

- ★ Encryption \mathcal{O}_{Enc} : On input (i, \mathbf{m}) , where $i \leq \text{numKeys}$, compute $\text{ct} \leftarrow \text{Enc}(\text{pk}_i, \mathbf{m})$ and increment numCt . Store the value ct in \mathcal{C} with key (i, numCt) . Return $(\text{numCt}, \text{ct})$.
- ★ Challenge Oracle: On input $(i, \mathbf{m}_0, \mathbf{m}_1)$ where $i \in \text{Hon}$ and $\mathbf{m}_0, \mathbf{m}_1 \in \mathcal{M}$, sample a bit $b \leftarrow \{0, 1\}$ uniformly at random, compute the challenge ciphertext $\text{ct}^* \leftarrow \text{Enc}(\text{pk}_i, \mathbf{m}_b)$, and increment numCt . Add numCt to the set Deriv . Store the value ct^* in \mathcal{C} with key (i, numCt) . Return $(\text{numCt}, \text{ct}^*)$. This oracle can only be queried once.
- ★ Reencryption $\mathcal{O}_{\text{ReEnc}}$: On input (i, j, k) where $i, j \leq \text{numKeys}$ and $k \leq \text{numCt}$, if $j \in \text{Cor}$ and $k \in \text{Deriv}$ return \perp . If there is no value in \mathcal{C} with key (i, k) , return \perp .
Otherwise, let ct_i be that value in \mathcal{C} , let $\text{ct}_j \leftarrow \text{ReEnc}(\text{rk}_{i \rightarrow j}, \text{ct}_i)$, and increment numCt . Store the value ct_j in \mathcal{C} with key (j, numCt) . If $k \in \text{Deriv}$, add numCt to the set Deriv .
Return $(\text{numCt}, \text{ct}_j)$.

Phase 3:

Decision: On input a bit $b' \in \{0, 1\}$, return win if $b = b'$.

The HRA advantage of \mathcal{A} is defined as

$$\text{Adv}_{\text{hra}}^{\mathcal{A}}(\lambda) = \Pr[\text{win}],$$

where the probability is over the randomness of \mathcal{A} and the oracles in HRA game.

Definition 6 (Proxy Reencryption: HRA Security). Given a security parameter 1^λ , a proxy reencryption scheme is HRA secure if for all probabilistic polynomial-time adversaries \mathcal{A} , there exists a negligible function negl such that

$$\text{Adv}_{\text{hra}}^{\mathcal{A}}(\lambda) < \frac{1}{2} + \text{negl}(\lambda)$$

The Concatenation Scheme demonstrates that CPA security does not necessarily imply HRA security. Together with following theorem, we see that HRA security is a strict strengthening of CPA security.

Theorem 3. Let PRE be an HRA secure proxy reencryption scheme. Then PRE is CPA secure.

Proof. From any probabilistic, polynomial-time algorithm \mathcal{A} (the CPA adversary), we construct an efficient algorithm \mathcal{A}' such that $\text{Adv}_{\text{hra}}^{\mathcal{A}'} = \text{Adv}_{\text{CPA}}^{\mathcal{A}}$. By the HRA security of PRE this advantage is negligible, completing the proof.

\mathcal{A}' runs \mathcal{A} and simulates the CPA security game (if \mathcal{A} does not follow the specification of the CPA security game, \mathcal{A}' simply aborts). Except for calls to $\mathcal{O}_{\text{ReEnc}}$, all oracle calls by \mathcal{A}' are passed along unaltered by \mathcal{A} to the corresponding HRA oracles, along with their responses.

\mathcal{A}' begins Phase 2 by requesting all (admissible) reencryption keys $\text{rk}_{i \rightarrow j}$ from $\mathcal{O}_{\text{ReKeyGen}}$. Oracle calls from \mathcal{A} to $\mathcal{O}_{\text{ReEnc}}$ are answered by \mathcal{A}' by computing

the reencryption using the appropriate reencryption key; this is possible because $\mathcal{O}_{\text{ReEnc}}$ returns \perp if and only if \mathcal{A}' is unable to get the corresponding reencryption key.

\mathcal{A}' perfectly implements the CPA security game, and \mathcal{A} wins that game if and only if \mathcal{A}' wins the HRA security game. Therefore $\text{Adv}_{\text{hra}}^{\mathcal{A}'} = \text{Adv}_{\text{CPA}}^{\mathcal{A}}$. Finally, the running time of \mathcal{A}' is polynomially related to the that of \mathcal{A} .

4.1 Sufficiency of HRA Security for Applications

Returning to the applications of proxy reencryption described in Section 2, we observe that HRA security provides substantially stronger security guarantees.

Encrypted email forwarding Using proxy reencryption with HRA security, Alice can forward encrypted email to Bob for a short period of time (during a vacation, say) and be sure that Bob cannot read her email after she returns.

Key escrow Similar to encrypted email forwarding, proxy reencryption with HRA can be used to enable law enforcement to read certain encrypted messages without compromising the secrecy of documents outside the scope of a search warrant or subpoena, for example.

Single-writer, many-reader encrypted storage Whereas proxy reencryption with CPA security can only support all or nothing access (i.e., all readers may access all data), HRA security can support fine grained access control (i.e., each reader may access only a specific subset of the data).

There is no question that HRA does not provide as much security as CCA, and that CCA-secure proxy reencryption would yield more robust applications. HRA security, however, can provide meaningful guarantees in the above applications.

Encrypted email forwarding If Alice is forwarding all emails to Bob, then Bob could certainly mount an attack outside the honest reencryption model. On the other hand, if Alice is forwarding only those emails from a third-party sender Charlie, then such an attack is impossible without the involvement of Charlie (supposing of course that the sender of an email can be authenticated).

Key escrow The hypothetical legal regime that establishes the government's power of exceptional access by way of key escrow could additionally prohibit the government from mounting chosen-ciphertext attacks. In the United States, a Constitutional argument could perhaps be made that law-enforcement use of chosen-ciphertext attacks must be limited.

Single-writer, many-reader encrypted storage The only ciphertexts being reencrypted are those uploaded by the single-writer to the proxy server (hence the name). It is by no means a stretch to require that the proxy server does not allow writes by unauthorized parties (i.e., the readers). If the honest writer only uploads honestly generated ciphertexts, HRA is appropriate.

5 Security of Existing Proxy Reencryption Schemes

Can we construct HRA-secure proxy reencryption? The most natural place to begin is with existing schemes.

We begin by demonstrating that the CPA secure scheme of [31] is not HRA secure. Although CPA security is strictly weaker than HRA security, we might hope that the existing CPA secure schemes already satisfy the more stringent definition. To this end, we identify a natural property—*reencryption simulatability*—sufficient to boost CPA security to HRA security.⁷

We examine the simple construction of CPA secure proxy reencryption from any fully-homomorphic encryption (FHE) presented in [19]. While the resulting proxy reencryption may not be HRA secure in general, if the FHE is *circuit private*—a property Gentry imbues into his FHE by rerandomization—the same construction will be HRA secure. We then examine pairing-based schemes, finding there too that rerandomization provides a direct path to HRA security.⁸

Remark 3. It may seem that CCA security should imply HRA security, but unfortunately the situation is not so clear. Intuitively, CCA security allows the adversary to make relatively unrestricted queries to both $\mathcal{O}_{\text{ReEnc}}$ and \mathcal{O}_{Dec} , whereas HRA restricts the adversary to making only honest reencryption queries to $\mathcal{O}_{\text{ReEnc}}$.

However the oracles in the CCA definition are restricted in a way that stymies a naive attempt at a reduction. The CCA definition prevents reencryptions or decryptions of all ciphertexts that could in principle be derived from the challenge (including by rerandomization). On the other hand, the HRA security game restricts reencryption queries only when the ciphertext is actually a derivative of the challenge. The adversary may reencrypt other encryptions of the challenge messages, so long as those encryptions were honestly generated independently from the challenge ciphertext.

We do not resolve the question of whether CCA security implies HRA security. By Theorem 5, any CCA secure proxy reencryption scheme satisfying reencryption simulatability is also HRA secure. See the full version of this article [13] for further discussion.

5.1 HRA Insecurity of [31]

Though it is easy to construct CPA secure encryption schemes that are not HRA secure, the question remains whether any previously proposed schemes

⁷ Some existing notions in the proxy reencryption literature seem powerful enough to elevate CPA security to HRA security, including proxy invisibility [4], unlinkability [17], and punctured security [1]. However, these notions are not sufficiently well defined to draw any concrete conclusions. The notion of key-privacy [3] does not in general suffice for HRA security.

⁸ While we do not examine every pairing-based construction of proxy reencryption, we suspect that rerandomizing reencryption will suffice for reencryption simulation in many, if not all.

fail to satisfy HRA security. In this section, we show that the construction of Polyakov, Rohloff, Sahu, and Vaikuntanathan [31, Section 5] is one such scheme. Their construction is based on a public key encryption scheme of Brakerski and Vaikuntanathan [9] and is CPA secure assuming the hardness of Ring Learning With Errors (RLWE).

As with the Trivial and Concatenation schemes, the HRA attack is simple yet severe: any single honestly generated ciphertext enables the recipient Bob to recover the sender Alice’s secret key with significant probability.

Theorem 4. *The proxy reencryption scheme of [31, Section 5] is not HRA secure.*

Proof. Except where noted, the notation used below is consistent with [31]; we restrict our description to those facts necessary to describe the HRA attack.

For n a power of 2 and prime $q \equiv 1 \pmod{2n}$, let $R_q = \mathbb{Z}_q[x]/(x^n + 1)$ be a ring of degree $(n - 1)$ polynomials with coefficients in \mathbb{Z}_q . The sender Alice’s secret key is s , and the recipient Bob’s secret key is s^* . Bob’s public key includes $O(\log q) = \text{poly}(n)$ RLWE samples $\theta_i^* = \beta_i \cdot s^* + pe_i$, where p is a public prime and the β_i and e_i are ring elements sampled by Bob.⁹ Ciphertexts are pairs of ring elements $(c_0, c_1) \in R_q^2$. By [9, Lemma 4], the distribution of c_1 is statistically close to uniform over R_q . By [27, Lemma 2.25], c_1 is invertible with probability at least $e^{-1} - \text{negl}(n)$. The result of reencrypting (c_0, c_1) is a pair (c'_0, c'_1) such that $c'_0 - s^* \cdot c'_1 = c_0 - s \cdot c_1 + pE_1$, where E_i is computable given c_1 and the e_i . This fact is used by [31] to prove the correctness of their scheme. Rearranging the above, we see that

$$s \cdot c_1 = c_0 + pE_1 - c'_0 + s^* \cdot c'_1.$$

Given any ciphertext (c_0, c_1) and its reencryption (c'_0, c'_1) , Bob can evaluate the above and compute $s \cdot c_1$. With probability at least $e^{-1} - \text{negl}(n)$, c_1 is invertible and Bob can recover the secret s .

5.2 Reencryption Simulatability.

While HRA is a strictly stronger security notion than CPA, we now show that if a CPA secure proxy reencryption scheme has an additional property we call *reencryption simulatability*, then it must also be HRA secure. Very roughly, reencryption simulatability means that ciphertexts resulting from computing $\text{ReEnc}(\text{rk}_{i \rightarrow j}, \text{ct}_i)$ can be simulated without knowledge of the sender’s secret key sk_i (but with knowledge of the plaintext message \mathbf{m} and the recipient’s secret key sk_j). Note that ReEnc uses $\text{rk}_{i \rightarrow j}$ which is a function of sk_i .

Reencryption simulatability allows an algorithm with access to the CPA oracles to efficiently implement the honest reencryption oracle. For intuition, consider the following approach to reducing HRA security to CPA security. Suppose

⁹ [31] separate the computation of θ_i^* from Bob’s public key, but this is only a matter of presentation.

there existed an adversary \mathcal{A}_{hra} violating the HRA security of a scheme; the reduction plays the roles of both the CPA adversary and the challenger in the HRA security game, and attempts to violate CPA security. To succeed, the reduction must be able to answer honest reencryption queries from uncorrupted keys to corrupted keys. Though the reduction knows the messages being reencrypted, it does not know the appropriate reencryption key. However, if it could indistinguishably simulate these reencryptions then it could indeed leverage \mathcal{A}_{hra} to violate CPA security.

We emphasize that the goal of Definition 7 is to capture a large swath of possible schemes while still enabling very simple proofs of simulatability (and thus of HRA security for existing CPA secure schemes). It is not intended to be the only avenue for proving HRA security of new or existing constructions. Reencryption simulatability is not necessary for HRA security of proxy reencryption. In particular, analogous versions of Theorem 5 hold with computational simulatability guarantees, but are more complicated [15, Footnote 7 and Appendix A].

Definition 7 (Reencryption Simulatability). *A proxy reencryption scheme PRE is reencryption simulatable if there exists a probabilistic, polynomial-time algorithm ReEncSim such that with high probability over aux , for all $\mathbf{m} \in \mathcal{M}$:*

$$(\text{ReEncSim}(\text{aux}), \text{aux}) \approx_s (\text{ReEnc}(\text{rk}_{a \rightarrow b}, \text{ct}_a), \text{aux}),$$

where \approx_s denotes statistical indistinguishability, and ct_a and aux are sampled according to

$$\begin{aligned} \text{pp} &\leftarrow \text{Setup}(1^\lambda), \\ (\text{pk}_a, \text{sk}_a) &\leftarrow \text{KeyGen}(\text{pp}), \\ (\text{pk}_b, \text{sk}_b) &\leftarrow \text{KeyGen}(\text{pp}), \\ \text{rk}_{a \rightarrow b} &\leftarrow \text{ReKeyGen}(\text{sk}_a, \text{pk}_b), \\ \text{ct}_a &\leftarrow \text{Enc}(\text{pk}_a, \mathbf{m}), \\ \text{aux} &= (\text{pp}, \text{pk}_a, \text{pk}_b, \text{sk}_b, \text{ct}_a, \mathbf{m}). \end{aligned}$$

A special case of the above is when $\text{ReEncSim}(\text{aux}) = \text{Enc}(\text{pk}_b, \mathbf{m})$ simply computes a fresh encryption of the message. That is, if reencrypted ciphertexts are distributed like fresh ciphertexts, then the scheme is reencryption simulatable.

Theorem 5. *Let PRE be an CPA secure, reencryption simulatable, proxy reencryption scheme. Then PRE is HRA secure.*

Proof (Outline). The proof proceeds according to the intuition above. From any probabilistic, polynomial-time algorithm $\mathcal{A} = \mathcal{A}_{\text{hra}}$ (the HRA adversary), we construct an algorithm $\mathcal{A}' = \mathcal{A}_{\text{cpa}}$ such that $\text{Adv}_{\text{cpa}}^{\mathcal{A}'}(\lambda) \geq \text{Adv}_{\text{hra}}^{\mathcal{A}}(\lambda) - \text{negl}(\lambda)$; by the CPA security of PRE this advantage is negligible, completing the proof.

\mathcal{A}_{cpa} runs \mathcal{A}_{hra} and simulates the HRA security game (if \mathcal{A}_{hra} does not follow the specification of the HRA security game, \mathcal{A}_{cpa} simply aborts). To answer

oracle calls by \mathcal{A}_{hra} to any oracle other than $\mathcal{O}_{\text{ReEnc}}$, \mathcal{A}_{cpa} simply passes the calls and answers unaltered to the corresponding CPA oracles.

To answer oracle calls to $\mathcal{O}_{\text{ReEnc}}$ between two uncorrupted keys or two corrupted keys, \mathcal{A}_{cpa} uses the corresponding reencryption key. On the other hand, for calls to $\mathcal{O}_{\text{ReEnc}}$ from an uncorrupted key to a corrupted key, \mathcal{A}_{cpa} simulates the reencryption using ReEncSim . This is possible because \mathcal{A}_{cpa} knows the underlying \mathbf{m} (along with the other information in aux).

Reencryption simulatability implies that the views of \mathcal{A}_{hra} in the real security game (using the real $\mathcal{O}_{\text{ReEnc}}$) and the simulated security game (using ReEncSim) are statistically close. \mathcal{A}_{cpa} wins the CPA security game if and only if \mathcal{A}_{hra} wins in the simulated HRA game described above.

5.3 Fully Homomorphic Encryption and Proxy Reencryption

There is an intimate connection between FHE and proxy reencryption: a sufficiently powerful somewhat homomorphic encryption scheme implies CPA secure proxy reencryption, which can be used to “bootstrap” the scheme to achieve fully homomorphic encryption [19]. For the relevant FHE definitions, see [19, Section 2].

Let $\text{FHE} = (\text{Setup}_{\text{FHE}}, \text{KeyGen}_{\text{FHE}}, \text{Enc}_{\text{FHE}}, \text{Dec}_{\text{FHE}}, \text{Eval}_{\text{FHE}})$ be an FHE scheme. Proxy reencryption can be constructed as follows (compare to the Trivial Scheme):

$\text{KeyGen}_{\text{PRE}}$, Enc_{PRE} and Dec_{PRE} are identical to their FHE counterparts.

$\text{ReKeyGen}_{\text{PRE}}(\text{sk}_i, \text{pk}_j) = \text{Enc}_{\text{FHE}}(\text{pk}_j, \text{sk}_i) \parallel \text{pk}_j$. The reencryption key $\text{rk}_{i \rightarrow j}$ is an encryption under pk_j of sk_i , along with the target public key pk_j .

$\text{ReEnc}_{\text{PRE}}(\text{rk}_{i \rightarrow j}, \text{ct}_i)$: Let $\text{ct}_{i \rightarrow j} \leftarrow \text{Enc}_{\text{FHE}}(\text{pk}_j, \text{ct}_i)$. Homomorphically compute the FHE decryption function $\text{Dec}_{\text{FHE}}(\text{sk}_i, \text{ct}_i)$ using the corresponding ciphertexts $\text{rk}_{i \rightarrow j}$ and $\text{ct}_{i \rightarrow j}$ (under pk_j). Namely, $\text{ct}_j = \text{Eval}_{\text{FHE}}(\text{pk}_j, \text{Dec}_{\text{FHE}}, \text{rk}_{i \rightarrow j}, \text{ct}_{i \rightarrow j})$.

The correctness of the FHE implies the correctness of the resulting proxy reencryption:

$$\text{Dec}_{\text{PRE}}(\text{sk}_j, \text{ct}_j) = \text{Dec}_{\text{FHE}}(\text{sk}_j, \text{ct}_j) = \text{Dec}_{\text{FHE}}(\text{sk}_i, \text{ct}_i) = \text{Dec}_{\text{PRE}}(\text{sk}_i, \text{ct}_i).$$

Furthermore, the proxy reencryption scheme is CPA secure.

To demonstrate that the construction might not be HRA secure, consider the following fully homomorphic encryption scheme FHE' constructed from any existing scheme FHE. The only modification made in FHE' is to $\text{Eval}_{\text{FHE}'}$:

$$\text{Eval}_{\text{FHE}'}(\text{pk}_j, C, \text{ct}_1, \text{ct}_2) := \text{Eval}_{\text{FHE}}(\text{pk}_j, C, \text{ct}_1, \text{ct}_2) \parallel \text{ct}_1.$$

Note that FHE' does not violate FHE compactness if ct_1 (in the proxy reencryption construction, $\text{rk}_{i \rightarrow j}$) is always of some size bounded by a polynomial in the security parameter of the FHE scheme; this suffices for our purpose. Instantiating the proxy reencryption construction with FHE' essentially results in the Concatenation Scheme, which is not HRA secure.

Circuit Privacy. An FHE scheme is *circuit private* if ciphertexts resulting from FHE evaluations are statistically indistinguishable from fresh ciphertexts [19]. Namely, if for any circuit C and any ciphertexts $\text{ct}_1, \dots, \text{ct}_t$:

$$\text{Enc}_{\text{FHE}}(\text{pk}, C(\text{ct}_1, \dots, \text{ct}_t)) \approx_s \text{Eval}_{\text{FHE}}(\text{pk}, C, \text{ct}_1, \dots, \text{ct}_t).$$

In [19], an FHE scheme without circuit privacy is modified to be circuit private by rerandomizing the ciphertexts resulting from Eval_{FHE} .

While our proxy reencryption construction above is not in general HRA secure, it is easy to see that if the underlying FHE is circuit private, then the proxy reencryption is reencryption simulatable (Definition 7). By Theorem 5, the resulting scheme is therefore HRA secure.

5.4 Pairing-Based Proxy Reencryption

Many constructions of proxy reencryption are based on the hardness of Diffie-Hellman-type problems over certain bilinear groups, including [4, 10, 26, 3, 21].

A prototypical construction is that of [4], which itself is based on the original scheme of [5]. For every λ , let G_1 and G_2 be groups of prime order $q = \Theta(2^\lambda)$, and let g be a generator of G_1 . Let e be a non-degenerate bilinear map $e : G_1 \times G_1 \rightarrow G_2$ (i.e., for all $h \in G_1$ and $a, b \in \mathbb{Z}_q$, $e(h^a, h^b) = e(h, h)^{ab}$, and for all generators g of G_1 , $e(g, g) \neq 1$). Let $Z = e(g, g)$. The message-space of the scheme is G_2 .

$\text{Setup}(1^\lambda)$: Output $\text{pp} = (q, g, G_1, G_2, e)$.

$\text{KeyGen}(\text{pp})$: Sample $a \leftarrow \mathbb{Z}_q$ uniformly at random. Output $\text{sk} = a$ and $\text{pk} = g^a$.

$\text{Enc}(\text{pk}, \mathbf{m})$: Sample $k \leftarrow \mathbb{Z}_q$ uniformly at random. Output $\text{ct} = (\text{pk}^k, \mathbf{m}Z^k) = (g^{ak}, \mathbf{m}Z^k)$.

$\text{ReKeyGen}(\text{sk}_A = a, \text{pk}_B = g^b)$: Output $\text{rk}_{A \rightarrow B} = g^{b/a}$.

$\text{ReEnc}(\text{rk}_{A \rightarrow B}, \text{ct}_A)$: Let $\text{ct}_A = (\alpha_1, \alpha_2)$. Output

$$\text{ct}_B = (e(\alpha_1, \text{rk}_{A \rightarrow B}), \alpha_2) = (e(g^{ak}, g^{b/a}), \mathbf{m}Z^k) = (Z^{bk}, \mathbf{m}Z^k).$$

$\text{Dec}(\text{sk}, \text{ct})$: Let $\text{ct} = (\alpha_1, \alpha_2)$. If $\alpha_1 \in G_2$ (i.e., if it is the result of ReEnc), then output $\alpha_2/\alpha_1^{1/a} = \mathbf{m}Z^k/Z^k = \mathbf{m}$. Otherwise $\alpha_1 \in G_1$ (i.e., it is a fresh ciphertext); output $\alpha_2/e(\alpha_1, g)^{1/a} = \mathbf{m}Z^k/e(g^{ak}, g)^{1/a} = \mathbf{m}Z^k/Z^k = \mathbf{m}$.

Is this scheme HRA secure? It is tempting to say that the scheme is reencryption simulatable; after all, given a message \mathbf{m} it is indeed straightforward to sample $(Z^{bk}, \mathbf{m}Z^k)$ for random $k \leftarrow \mathbb{Z}_q$. However $\text{ct}_A = (g^{ak}, \mathbf{m}Z^k)$ and $\text{ct}_B = \text{ReEnc}(\text{rk}_{A \rightarrow B}, \text{ct}_A) = (Z^{bk}, \mathbf{m}Z^k)$ share the randomness k . Given $\text{ct}_A = (g^{ak}, \mathbf{m}Z^k)$ and \mathbf{m} , it is not clear how to output $(g^{bk}, \mathbf{m}Z^k)$.

Rerandomization A minor modification to the scheme above guarantees reencryption simulatability and therefore HRA security. Replace ReEnc above with ReEnc' :

$\text{ReEnc}'(\text{rk}_{A \rightarrow B}, \text{ct}_A)$: Compute $(Z^{bk}, \mathbf{m}Z^k) = \text{ReEnc}(\text{rk}_{A \rightarrow B}, \text{ct}_A)$. Sample $k' \leftarrow \mathbb{Z}$ uniformly at random, and output $(Z^{bk} \cdot e(g^b, g^{k'}), \mathbf{m}Z^k \cdot e(g, g^{k'})) = (Z^{b(k+k')}, \mathbf{m}Z^{k+k'})$.

The resulting proxy reencryption scheme maintains the CPA security of the original, as the only modification is the rerandomization of reencrypted ciphertexts (which can be done by anyone with knowledge of the public parameters).

Furthermore, the scheme is now reencryption simulatable. To see why, observe that the resulting reencrypted ciphertexts are uniformly distributed in $\{(\text{ct}_1, \text{ct}_2) \in G_2 \times G_2 : \text{ct}_2 / \text{ct}_1^{1/b} = \mathbf{m}\}$, independent of all other keys and ciphertexts. Such ciphertexts are easily sampled with knowledge of pp , $\text{pk}_B = g^b$ and \mathbf{m} as follows.

$\text{ReEncSim}(\text{pp}, \text{pk}_B, \mathbf{m})$: Sample $k' \leftarrow \mathbb{Z}_q$ uniformly at random, and output $(e(\text{pk}_B, g^{k'}), \mathbf{m} \cdot e(g, g^{k'})) = (Z^{bk'}, \mathbf{m}Z^{bk'})$.

Thus, by Theorem 5, the modified scheme is HRA secure. Observe that rerandomization was the key to achieving circuit privacy (and thereby HRA security) in the FHE-based proxy reencryption construction as well.

The pairing-based schemes of [3] and [21] already incorporate rerandomization during reencryption. In the former case, it is used to achieve “key privacy;” in the latter, to achieve obfuscation of the reencryption functionality. In each, it is straightforward to show that the schemes are also reencryption simulatable and therefore HRA secure.

References

1. Ananth, P., Cohen, A., Jain, A.: Cryptography with updates. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 445–472. Springer (2017)
2. Aono, Y., Boyen, X., Wang, L., et al.: Key-private proxy re-encryption under lwe. In: International Conference on Cryptology in India. pp. 1–18. Springer (2013)
3. Ateniese, G., Benson, K., Hohenberger, S.: Key-private proxy re-encryption. In: CT-RSA. vol. 5473, pp. 279–294. Springer (2009)
4. Ateniese, G., Fu, K., Green, M., Hohenberger, S.: Improved proxy re-encryption schemes with applications to secure distributed storage. ACM Transactions on Information and System Security (TISSEC) 9(1), 1–30 (2006)
5. Blaze, M., Bleumer, G., Strauss, M.: Divertible protocols and atomic proxy cryptography. In: Advances in Cryptology EUROCRYPT’98, pp. 127–144. Springer (1998)
6. Boneh, D., Halevi, S., Hamburg, M., Ostrovsky, R.: Circular-secure encryption from decision diffie-hellman. In: Annual International Cryptology Conference. pp. 108–125. Springer (2008)
7. Boneh, D., Lewi, K., Montgomery, H., Raghunathan, A.: Key homomorphic prfs and their applications. In: Advances in Cryptology–CRYPTO 2013, pp. 410–428. Springer (2013)
8. Borcea, C., Polyakov, Y., Rohloff, K., Ryan, G., et al.: Picador: End-to-end encrypted publish–subscribe information distribution with proxy re-encryption. Future Generation Computer Systems 71, 177–191 (2017)

9. Brakerski, Z., Vaikuntanathan, V.: Fully homomorphic encryption from ring-lwe and security for key dependent messages. In: Annual cryptology conference. pp. 505–524. Springer (2011)
10. Canetti, R., Hohenberger, S.: Chosen-ciphertext secure proxy re-encryption. In: Proceedings of the 14th ACM conference on Computer and communications security. pp. 185–194. ACM (2007)
11. Chandran, N., Chase, M., Liu, F.H., Nishimaki, R., Xagawa, K.: Re-encryption, functional re-encryption, and multi-hop re-encryption: A framework for achieving obfuscation-based security and instantiations from lattices. In: PKC. pp. 95–112. Springer (2014)
12. Chow, S.S., Weng, J., Yang, Y., Deng, R.H.: Efficient unidirectional proxy re-encryption. In: International Conference on Cryptology in Africa. pp. 316–332. Springer (2010)
13. Cohen, A.: What about bob? the inadequacy of cpa security for proxy reencryption. Cryptology ePrint Archive, Report 2017/785 (2017), <https://eprint.iacr.org/2017/785>
14. Derler, D., Krenn, S., Lorünser, T., Ramacher, S., Slamanig, D., Striecks, C.: Revisiting proxy re-encryption: Forward secrecy, improved security, and applications. In: IACR International Workshop on Public Key Cryptography. pp. 219–250. Springer (2018)
15. Dttling, N., Nishimaki, R.: Universal proxy re-encryption. Cryptology ePrint Archive, Report 2018/840 (2018), <https://eprint.iacr.org/2018/840>
16. Everspaugh, A., Paterson, K., Ristenpart, T., Scott, S.: Key rotation for authenticated encryption. In: Annual International Cryptology Conference. pp. 98–129. Springer (2017)
17. Fan, X., Liu, F.H.: Proxy re-encryption and re-signatures from lattices (2017)
18. Fuchsbauer, G., Kamath, C., Klein, K., Pietrzak, K.: Adaptively secure proxy re-encryption. Cryptology ePrint Archive, Report 2018/426 (2018), <https://eprint.iacr.org/2018/426>
19. Gentry, C.: A fully homomorphic encryption scheme. Stanford University (2009)
20. He, Y.J., Hui, L.C., Yiu, S.M.: Avoid illegal encrypted drm content sharing with non-transferable re-encryption. In: Communication Technology (ICCT), 2011 IEEE 13th International Conference on. pp. 703–708. IEEE (2011)
21. Hohenberger, S., Rothblum, G., Shelat, A., Vaikuntanathan, V.: Securely obfuscating re-encryption. Theory of Cryptography pp. 233–252 (2007)
22. Ivan, A.A., Dodis, Y.: Proxy cryptography revisited. In: NDSS (2003)
23. Jakobsson, M.: On quorum controlled asymmetric proxy re-encryption. In: International Workshop on Public Key Cryptography. pp. 112–121. Springer (1999)
24. Khurana, H., Heo, J., Pant, M.: From proxy encryption primitives to a deployable secure-mailing-list solution. In: International Conference on Information and Communications Security. pp. 260–281. Springer (2006)
25. Lee, S., Park, H., Kim, J.: A secure and mutual-profitable drm interoperability scheme. In: Computers and Communications (ISCC), 2010 IEEE Symposium on. pp. 75–80. IEEE (2010)
26. Libert, B., Vergnaud, D.: Unidirectional chosen-ciphertext secure proxy re-encryption. In: International Workshop on Public Key Cryptography. pp. 360–379. Springer (2008)
27. Lyubashevsky, V., Peikert, C., Regev, O.: A toolkit for ring-lwe cryptography. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 35–54. Springer (2013)

28. Nunez, D., Agudo, I., Lopez, J.: A parametric family of attack models for proxy re-encryption. In: Computer Security Foundations Symposium (CSF), 2015 IEEE 28th. pp. 290–301. IEEE (2015)
29. Oz, F., Murray, B., Dreyfuss, R.: What About Bob. Touchstone Pictures (1991)
30. Phong, L., Wang, L., Aono, Y., Nguyen, M., Boyen, X.: Proxy re-encryption schemes with key privacy from lwe. Tech. rep., Cryptology ePrint Archive, Report 2016/327, 2016. <http://eprint.iacr.org/2016/327> (2016)
31. Polyakov, Y., Rohloff, K., Sahu, G., Vaikuntanathan, V.: Fast proxy re-encryption for publish/subscribe systems. ACM Transactions on Privacy and Security (TOPS) 20(4), 14 (2017)

A The Trivial Scheme

The following description and definition of circular security is adapted with slight modification from [6].

Let $(\text{KeyGen}, \text{Enc}, \text{Dec})$ be a public-key encryption scheme with key space \mathcal{K} and message space \mathcal{M} such that $\mathcal{K} \subseteq \mathcal{M}$. Let $n > 0$ be an integer and let \mathcal{C} be the set of functions $\mathcal{C} = \{f : \mathcal{K}^n \rightarrow \mathcal{M}\}$ consisting of

- all $|\mathcal{M}|$ constant functions $f_m(z) = m$ for all $z \in \mathcal{K}^n$, and
- all n selector functions $f_i(x_1, \dots, x_n) = x_i$ for $1 \leq i \leq n$.

We define circular security using the following game between a challenger and an adversary \mathcal{A} . For an integer $n > 0$ and a security parameter λ , the game proceeds as follows:

Initialization: The challenger chooses a random bit $b \leftarrow \{0, 1\}$. It generates $(\text{pk}_1, \text{sk}_1), \dots, (\text{pk}_n, \text{sk}_n)$ by running $\text{KeyGen}(1^\lambda)$ n times, and sends $(\text{pk}_1, \dots, \text{pk}_n)$ to \mathcal{A} .

Queries: The adversary repeatedly issues queries where each query is of the form (i, f) with $1 \leq i \leq n$ and $f \in \mathcal{C}$. The challenger responds by setting $y = f(\text{sk}_1, \dots, \text{sk}_n)$ and

$$\text{ct} \leftarrow \begin{cases} \text{Enc}(\text{pk}_i, y) & \text{if } b = 0 \\ \text{Enc}(\text{pk}_i, 0^{|y|}) & \text{if } b = 1 \end{cases}$$

and sends ct to \mathcal{A} .

Finish: Finally, the adversary outputs a bit $b' \in \{0, 1\}$.

We say that \mathcal{A} wins the game if $b = b'$. Let win be the event that \mathcal{A} wins the game and define \mathcal{A} 's advantage as

$$\text{Adv}_{\text{circ},n}^{\mathcal{A}}(\lambda) = \Pr[\text{win}].$$

Definition 8 (n -Circular Security). *We say that a public-key encryption scheme $(\text{KeyGen}, \text{Enc}, \text{Dec})$ is n -way circular secure if for all probabilistic polynomial time adversaries \mathcal{A} , there exists a negligible function negl such that*

$$\text{Adv}_{\text{circ},n}^{\mathcal{A}}(\lambda) < \frac{1}{2} + \text{negl}(\lambda).$$

Because existing constructions of circularly secure encryption schemes based on standard assumptions require a bound on the total number of public keys n , the corresponding Trivial Scheme will only satisfy a bounded-key variant of CPA security, defined next.

Definition 9 (Proxy Reencryption: n -CPA Security). For $n \in \mathbb{N}$, the n -CPA security game is identical to the CPA security game in Definition 3 except for the following underlined modifications. Recall that `numKeys` is initialized to 0 and is incremented after every key generation call in the security game.

Uncorrupted Key Generation: If `numKeys` = n , return \perp . Otherwise, obtain a new key pair $(pk_i, sk_i) \leftarrow \text{KeyGen}(\text{pp})$. \mathcal{A} is given pk_i . The current value of `numKeys` is added to `Hon` and `numKeys` is incremented.

Corrupted Key Generation: If `numKeys` = n , return \perp . Otherwise, obtain a new key pair $(pk_i, sk_i) \leftarrow \text{KeyGen}(\text{pp})$. \mathcal{A} is given (pk_i, sk_i) . The current value of `numKeys` is added to `Cor` and `numKeys` is incremented.

The corresponding n -CPA advantage of \mathcal{A} is denoted $\text{Adv}_{\text{cpa},n}^{\mathcal{A}}(\lambda)$. A proxy reencryption scheme is n -CPA secure if for all probabilistic polynomial-time adversaries \mathcal{A} , there exists a negligible function negl such that

$$\text{Adv}_{\text{cpa},n}^{\mathcal{A}}(\lambda) < \frac{1}{2} + \text{negl}(\lambda)$$

Trivial Scheme Let $(\text{KeyGen}_{\text{circ}}, \text{Enc}_{\text{circ}}, \text{Dec}_{\text{circ}})$ be an n -way circular secure encryption scheme. Let `Setup` $\equiv \perp$, `KeyGen` $\equiv \text{KeyGen}_{\text{circ}}$; `Enc` $\equiv \text{Enc}_{\text{circ}}$;

$$\text{ReKeyGen}(sk_i, pk_j) := \text{Enc}_{\text{circ}}(pk_j, sk_i)$$

$$\text{ReEnc}(rk_{i \rightarrow j}, ct_i) := ct_i \parallel rk_{i \rightarrow j}$$

$$\text{Dec}(sk, ct) := \begin{cases} \text{Dec}_{\text{circ}}(\text{Dec}_{\text{circ}}(sk, ct^2), ct^1) & \text{if } ct = ct^1 \parallel ct^2 \\ \text{Dec}_{\text{circ}}(sk, ct) & \text{otherwise} \end{cases}.$$

Theorem 2 states that if $(\text{KeyGen}_{\text{circ}}, \text{Enc}_{\text{circ}}, \text{Dec}_{\text{circ}})$ is an n -way circular secure encryption scheme, then the corresponding Trivial Scheme PRE is an n -CPA secure proxy reencryption scheme. In fact, the proof below extends the case when there are n uncorrupted keys and any number of corrupted keys.

Proof (of Theorem 2). For all $n \in \mathbb{N}$ and any probabilistic, polynomial-time algorithm \mathcal{A} (the adversary against the trivial scheme), we construct an efficient algorithm $\mathcal{A}_{\text{circ}}$ such that $\text{Adv}_{\text{circ},n}^{\mathcal{A}_{\text{circ}}} = \frac{1}{2} \cdot \text{Adv}_{\text{cpa},n}^{\mathcal{A}}$. By the hypothesis, this advantage is negligible, completing the proof.

At the beginning of the game, the circular security challenger picks a random bit b . If $b = 0$, then the Queries oracle encrypts all messages correctly; if $b = 1$, then the Queries oracle encrypts only the message 0. $\mathcal{A}_{\text{circ}}$ runs \mathcal{A} and simulates the CPA security game for PRE (if \mathcal{A} does not follow the specification of the game, $\mathcal{A}_{\text{circ}}$ simply aborts).

At the start of Phase 1, $\mathcal{A}_{\text{circ}}$ calls its Initialization oracle in the circular security game. In return it receives the public keys $(pk_1^{\text{circ}}, \dots, pk_n^{\text{circ}})$. To answer

an Uncorrupted Key Generation query, $\mathcal{A}_{\text{circ}}$ gives to \mathcal{A} the first unused public key $\text{pk}_i^{\text{circ}}$ from this list. To answer a Corrupted Key Generation query, $\mathcal{A}_{\text{circ}}$ generates a new key pair $(\text{pk}, \text{sk}) \leftarrow \text{KeyGen}$ and gives (pk, sk) to the adversary.

\mathcal{A} begins Phase 2 by using its Queries oracle to learn the reencryption keys for all pairs of uncorrupted keys generated. Using its knowledge of the corrupted secret keys, it also computes reencryption keys for all the pairs of corrupted keys generated. Oracle calls from \mathcal{A} to $\mathcal{O}_{\text{ReKeyGen}}$ are answered with the corresponding reencryption key (or with \perp). To answer oracle calls from \mathcal{A} to $\mathcal{O}_{\text{ReEnc}}$, computes the appropriate response; namely, it concatenates the reencryption key to the ciphertext (or returns \perp).

At some point, \mathcal{A} queries the Challenge oracle with an honest key index i and a pair of messages $(\mathbf{m}_0, \mathbf{m}_1)$. $\mathcal{A}_{\text{circ}}$ chooses a random one of the messages \mathbf{m} and queries its own Queries oracle with the pair (i, \mathbf{m}) , returning the resulting ciphertext to \mathcal{A} .

Finally, \mathcal{A} guesses whether $\mathbf{m} = \mathbf{m}_0$ or \mathbf{m}_1 . If \mathcal{A} guesses correctly, $\mathcal{A}_{\text{circ}}$ guesses the bit $b' = 0$. Otherwise, $\mathcal{A}_{\text{circ}}$ guesses a random $b' \leftarrow \{0, 1\}$. Conditioned on $b = 0$, $\mathcal{A}_{\text{circ}}$ perfectly simulates the PRE security game, and guesses $b' = 0$ with probability $\text{Adv}_{\text{cpa}, n}^{\mathcal{A}}$. It follows that $\text{Adv}_{\text{circ}, n}^{\mathcal{A}_{\text{circ}}} = \frac{1}{2} \cdot \text{Adv}_{\text{cpa}, n}^{\mathcal{A}}$.

B Comparison to RIND-CPA

The concurrent work of Derler, Krenn, Lorünser, Ramacher, Slamanig, and Striecks identify the same problem with CPA security as discussed [14]. They define a new security notion—RIND-CPA security—as an additional property that proxy reencryptions schemes should guarantee.

The key feature of RIND-CPA security is that the adversary gets access to an unrestricted ReEnc oracle, but only before seeing the challenge ciphertext. The definition is similar to IND-CCA_{0,1} of [28]. The definition of the RIND-CPA security experiment is from [14, Experiment 8].

Definition 10 (RIND-CPA Security Experiment).

$\text{pp} \leftarrow \text{Setup}(1^\lambda), (\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(\text{pp}), b \leftarrow \{0, 1\}$
 $(\text{pk}^*, \text{st}) \leftarrow \mathcal{A}(\text{pp}, \text{pk})$
 $\text{rk} \leftarrow \text{ReKeyGen}(\text{sk}, \text{pk}^*)$
 $(\mathbf{m}_0, \mathbf{m}_1, \text{st}) \leftarrow \mathcal{A}^{\{\text{ReEnc}(\text{rk}, \cdot)\}}(\text{st})$
 $b^* \leftarrow \mathcal{A}(\text{st}, \text{Enc}(\text{pk}, \mathbf{m}_b))$
if $b = b^$ return 1, else return 0.*

RIND-CPA security requires that for all efficient adversaries, the probability of outputting 1 in the experiment is $\frac{1}{2} \pm \text{negl}(\lambda)$.

In this section, we compare the approach of [14] with ours. We begin by describing RIND-CPA security as defined by [14]. Next, we compare RIND-CPA with HRA security informally, arguing that HRA provides the better generalization of Enc-CPA security to the PRE setting. Finally, we show that HRA and RIND-CPA security are incomparable notions.

B.1 Informal comparison

RIND-CPA is less suitable than HRA as a replacement for CPA security of proxy reencryption. First and most importantly, HRA better captures the intuitive guarantees of Enc-CPA security for standard encryption. Second, access to an unrestricted ReEnc oracle makes it a more useful as a testing ground for the development of new techniques. Finally, two idiosyncrasies of the [14] formulation of RIND-CPA security present additional issues.

Capturing Enc-CPA security. In Enc-CPA security for standard encryption, the adversary is able to arbitrarily affect the distribution of plaintext messages. One way of viewing this aspect of the definition is that Enc-CPA requires security while being agnostic as to the true distribution over messages (except that it is efficiently sampleable). Other than choosing the distribution over messages, the adversary is only allowed to see publicly-available information (i.e. public keys and parameters) and honestly encrypted ciphertexts. Informally, the Enc-CPA guarantee is that security should hold under normal operating conditions against eavesdropping parties without making distributional assumptions on plaintext messages. However, Enc-CPA makes no guarantees about dishonestly generated or malformed ciphertexts.

HRA security captures this intuitive guarantee better than RIND-CPA. In the course of normal operation of a proxy reencryption, an adversarial party will see reencrypted ciphertexts. These ciphertexts may come at any time—both before and after other ciphertexts whose contents should remain secret. While HRA allows reencryption both before and after the challenge, RIND-CPA restricts the reencryption oracle to the period before the challenge.

HRA makes minimal assumptions about the distribution of plaintext messages by allowing the adversary to choose messages itself, just as in Enc-CPA. RIND-CPA goes further by making requirements in the face of malformed or dishonestly generated ciphertexts.

A testing ground for new techniques. For classical encryption, Enc-CCA security is strictly stronger than Enc-CPA security. In fact, there are many settings where Enc-CPA security is demonstrably insufficient. Why then does the cryptography community continue to study it? There are many answers to this question, but we mention only two. First, although insufficient for some applications, Enc-CPA is useful in others. Second, it is useful as an intermediate goal because it seems to capture a sort of hard core of the general problem of encryption and spurs the development of new techniques.

HRA security enjoys these same features; RIND-CPA does not. As for usefulness for applications, HRA is meaningful in many of the envisioned applications of proxy reencryption—many more than CPA security. Because RIND-CPA restricts the reencryption oracle to the period before the challenge ciphertext, its usefulness in applications is less clear.

The challenge of constructing CCA secure proxy reencryption is the same as the challenge of Enc-CCA secure encryption: namely, dealing with dishonestly

generated, possibly malformed ciphertexts. RIND-CPA, by allowing malformed ciphertexts, presents similar challenges as full CCA security.

As for the usefulness of HRA as an intermediate goal towards CCA security, the historical development of proxy reencryption is proof itself. This sounds paradoxical: how can this be true if the notion has only just been introduced in this work? Many of cryptographers that were targeting CPA security developed schemes that achieve HRA security with only minimal modification. The techniques developed in these constructions were later adapted to achieve CCA security. This suggests that cryptographers’ intuitions for the hard core of reencryption were not flawed, only the formalization of these intuitions as CPA security. HRA security is a better formalization for these intuitions and thus an appropriate intermediate goal for reencryption research.

Idiosyncrasies of the RIND-CPA definition. We mention two unusual properties of the [14] definition. Unlike the adversary’s access to a ReEnc oracle, these properties are not inherent in the RIND-CPA concept. It would be easy to propose a modified RIND-CPA definition that did not have these properties (e.g., IND-CCA_{0,1} in [28]).

First, the definition only considers the two party setting. Much like the informal description of proxy reencryption in Section 1, there is only a single uncorrupted key and a single corrupted key. It is easy to show that security in the two party setting does not imply security in a many party setting.

Second, not only are inputs to ReEnc allowed to be malformed, but the corrupted public key pk^* can be malformed as well. The adversary outputs pk^* itself and it needs not be honestly generated. This makes RIND-CPA security as defined in [14] formally incomparable to all other definitions of proxy reencryption security we know of, including the IND-CCA_{0,1} of [28].

These drawbacks of the [14] definition do not affect the proof of Theorem 7, but neither does the proof depend on them.

B.2 Separating RIND-CPA and HRA security

The following pair of theorems support the conclusion that HRA security and RIND-CPA security are incomparable.

Theorem 6. *If there exists an HRA secure PRE scheme, then there exists a PRE scheme that is HRA secure but not RIND-CPA secure.*

Proof. Suppose PRE is HRA secure, and let \top be a special symbol that is not a valid ciphertext. Define a new scheme PRE' by modifying reencryption as follows:

$$\text{ReEnc}'(rk, ct) := \begin{cases} \text{ReEnc}(rk, ct) & \text{if } ct \neq \top \\ rk & \text{if } ct = \top \end{cases} .$$

PRE' is still HRA secure: $\mathcal{O}_{\text{ReEnc}'}$ is functionally equivalent to $\mathcal{O}_{\text{ReEnc}}$ when restricted to honestly generated ciphertexts..

PRE' is not RIND-CPA secure: a single call to $\mathcal{O}_{\text{ReEnc}}(i, j, \top)$ (made before the challenge) allows the adversary to learn the reencryption key $\text{rk}_{i \rightarrow j}$ and thereby decrypt the challenge ciphertext.

Theorem 7. *Under the assumptions stated below, there exists a PRE scheme that is RIND-CPA secure but not HRA secure.*

The claim assumes the existence of pair of encryption schemes, PRE and FHE with the following properties. PRE is a RIND-CPA secure proxy reencryption scheme with a ciphertext space $\mathcal{C}_{\text{inner}}$. FHE is a circuit private fully homomorphic encryption scheme with message space $\mathcal{M}_{\text{fhe}} = \mathcal{C}_{\text{inner}}$. The message spaces and ciphertext spaces of the two schemes are all disjoint and efficiently decidable. Finally, the additional proxy reencryption scheme PRE_{FHE} corresponding to the FHE scheme (see Section 5.3) is RIND-CPA secure.¹⁰ For simplicity, we also assume perfect correctness of reencryption (for both schemes) and of homomorphic evaluation.

Below we present an intuition for the proof of Theorem 7. The proof is included in the full version [13].

Proof Intuition for Theorem 7 Recall that RIND-CPA security allows the adversary access to an unrestricted ReEnc oracle, but only before the challenge ciphertext is generated. The main difficulty in separating RIND-CPA and HRA security is the restriction in the HRA reencryption oracle to honestly generated ciphertexts.

The first idea in our construction is the observation that separating RIND-CPA and HRA security would be easy if it were possible to use Enc oracle to generate a fresh, honest encryption of the challenge plaintext. This fresh encryption could be reencrypted by the HRA reencryption oracle to a corrupted key, revealing the challenge plaintext.

The second idea is to have two layers of encryption, where the message space of the outer layer is equal to the ciphertext space of the inner layer. If the challenge ciphertext comes from the inner layer, then it can be used as input to the Enc oracle to generate a new *outer ciphertext* containing information about the challenge plaintext—namely, an encryption of the challenge ciphertext. The outer ciphertext is honestly generated and can be reencrypted to a corrupt party and decrypted. But it seems we are no better off; decrypting the outer ciphertext only returns the challenge ciphertext still encrypted under the key of an honest party.

The third idea is to modify ReEnc —using fully homomorphic encryption—to reencrypt both the outer ciphertext and the inner challenge ciphertext. In addition to the usual reencrypted ciphertext, we augment ReEnc to output an

¹⁰ The proof requires that an encryption scheme be both fully homomorphic and support proxy reencryption with RIND-CPA security. For concreteness, we have chosen to assume that there exists an FHE scheme whose corresponding PRE is RIND-CPA secure, but a different construction would suffice. We do not further explore the underlying cryptographic assumptions needed to instantiate this encryption scheme.

additional, *doubly reencrypted* ciphertext, where both the outer and inner ciphertexts have been reencrypted. If the recipient of the resulting ciphertext is corrupt, the adversary can decrypt both layers and recover the challenge plaintext, violating HRA security.

We now describe the intuition for how to perform double reencryption. Suppose the proxy reencryption scheme used for the outer layer of encryption is also fully homomorphic. Such a scheme can be constructed from any FHE scheme (see Section 5.3). Given input an outer layer ciphertext $\text{ct}_{\text{outer}} = \text{Enc}(\text{ct}_{\text{inner}})$, ReEnc will homomorphically evaluate $\text{Eval}_{\text{fhe}}(\text{ReEnc}, \text{ct}_{\text{outer}})$. The result is an (non-reencrypted) outer ciphertext containing a reencrypted inner ciphertext. Then, ReEnc reencrypts that outer ciphertext. This produces a reencrypted outer ciphertext containing a reencrypted inner ciphertext.

Violating HRA security is simple: the adversary encrypts the challenge ciphertexts, reencrypts it to a corrupted key, then decrypts the doubly-reencrypted component twice to recover the challenge message.

It remains to prove that the constructed PRE scheme is RIND-CPA secure. The homomorphic double reencryption functionality can be simulated by a sequence of calls to Enc , Dec and $\mathcal{O}_{\text{ReEnc}}$, allowing us to analyze the two-layered scheme without the double-reencryption modification to ReEnc . The RIND-CPA security of that scheme follows directly from the RIND-CPA security of the PRE scheme underlying the two layers.