

# Improved Security Notions for Proxy Re-Encryption to Enforce Access Control

Ela Berners-Lee <sup>†</sup>  
Ela.Berners-Lee.2010@live.rhul.ac.uk  
Royal Holloway, University of London

**Abstract** Proxy Re-Encryption (PRE) allows a ciphertext encrypted under Alice’s public key to be transformed to an encryption under Bob’s public key without revealing either the plaintext or the decryption keys. PRE schemes have clear applications to cryptographic access control by allowing outsourced data to be selectively shared to users via re-encryption to appropriate keys. One concern for this application is that the server should not be able to perform unauthorised re-encryptions. We argue that current security notions do not adequately address this concern. We revisit existing definitions for PRE, starting by challenging the concept of unidirectionality, which states that re-encryption tokens from  $A$  to  $B$  cannot be used to re-encrypt from  $B$  to  $A$ . We strengthen this definition to reflect realistic scenarios in which adversaries may try to reverse a re-encryption by retaining information about prior ciphertexts and re-encryption tokens. We then strengthen the adversarial model to consider malicious adversaries that may collude with corrupt users and attempt to perform unauthorised re-encryptions; this models a malicious cloud service provider aiming to subvert the re-encryption process to leak sensitive data. Finally we revisit the notion of authenticated encryption for PRE. This currently assumes the same party who created the message also encrypted it, which is not necessarily the case in re-encryption. We thus introduce the notion of *ciphertext origin authentication* to determine who encrypted the message (initiated a re-encryption) and show how to fulfil this requirement in practice.

**Keywords:** Proxy re-encryption, applied cryptography, unidirectional, multi-hop, malicious model, access control

## 1 Introduction

There are many practical situations in which a ciphertext encrypted under one key must be re-encrypted such that it represents an encryption of the same message under a different key. This is trivial when data is stored locally, but is less straightforward when data is stored remotely by an untrusted server such as a cloud service provider. Proxy Re-Encryption (PRE) [4] enables a third party to re-encrypt a ciphertext using an *update token* generated by the client, in such a way that neither the decryption keys nor plaintext are revealed.

A common motivation cited for PRE is *email forwarding* [2–4, 6, 7, 17], where Alice wants to forward her emails to Bob and have him read them on her behalf, without revealing her secret key. With PRE, she can generate an update token which the email server uses to re-encrypt ciphertexts under Alice’s key to ciphertexts under Bob’s key, without the server reading her emails.

Another motivation, which we focus on in this paper, is enforcing cryptographic access control over remotely stored files [3, 16]. If data is given a classification level, and keys are shared with

---

<sup>†</sup>The author was supported by the EPSRC and the UK government as part of the Centre for Doctoral Training in Cyber Security at Royal Holloway, University of London (EP/K035584/1)

users according to access control policy, then re-encryption is used to enforce changes to the policy. In particular, re-encryption can signify a change in user access rights, revocation or key expiry.

In this paper, we revisit the security notions for PRE with a particular focus on enforcing access control as an application. We show that, in many cases, existing notions are insufficient to ensure the necessary security in this setting.

The main issue not addressed by existing literature is that a malicious server should not be able to perform an unauthorised re-encryption. We break this down into two main security notions: the inability to create a valid update token even given having seen a number of valid tokens, and a stronger notion of unidirectionality which considers reversal attacks.

*Malicious adversaries.* Most previous work considers *honest-but-curious* adversaries that follow the protocol honestly but try to learn the underlying plaintexts.

*Token robustness.* Existing work [10] tackles the issue of controlling which ciphertexts are re-encrypted by defining *ciphertext dependence*, where tokens are created to only be valid for specific ciphertexts. We strengthen this definition to create a security notion which states that an adversary cannot generate a valid update token which re-encrypts to a previously unused key, even having seen a number of legitimate tokens. We then give an example which shows that this notion is stronger than ciphertext dependence. We call this *token robustness*.

*Unidirectionality.* To tackle re-encryptions to keys which a ciphertext has previously been encrypted under, we revisit the existing notion of unidirectionality, which states that a re-encryption token can only be used to transform a ciphertext under  $pk_i$  to  $pk_j$  and not from  $pk_j$  to  $pk_i$  (otherwise the scheme is *bidirectional*). The ability to re-encrypt back to the old key can grant access back to an unauthorised user or re-encrypt to an expired key. Current notions do not consider reversal attacks where a server may retain some limited information about an old ciphertext and update token to reverse the re-encryption. This consideration is particularly important for token robust schemes where the token used to perform the update is crucial in reverting a ciphertext back to an expired key. We formally define reversal attacks with respect to the size of the state the server must retain in order to reverse a re-encryption. We use this to form an upper bound on security definitions for directionality. This is stronger than existing notions for unidirectionality as it gives the adversary more control over the information they have access to than traditional notions, which only consider tokens given to the adversary. We then use this together with token robustness to create a notion for *best-achievable unidirectionality*. Overall, our security model covers a wider range of attacks than prior definitions. We show that these definitions can be met by introducing a simple adaptation of ElGamal-based PRE in Appendix B.

*Ciphertext Origin Authentication.* Finally we revisit the notion of data origin authentication for PRE. Typically, data origin authentication assumes that the same party who created the message also encrypted it, hence tying the data owner's identity to the message within the ciphertext is sufficient. Whilst this is a valid assumption for many encryption scenarios, for access control where more than one party shares a key, the same assumption cannot be made. We create a new notion of *ciphertext origin authentication* where the encryptor / re-encryption initiator's identity is tied to the ciphertext as opposed to the message, and re-encryption updates this accordingly. We offer an extension to our unidirectional token robust scheme, and show how to develop similar extensions for other schemes.

The structure of this paper is as follows: In Section 2 and ?? we formally discuss existing work and current notions of security for PRE. We define *token robustness* in Section 4 and *maximal irreversibility* in Section 5. We then , and build on these to define *best-achievable unidirectionality*. In Section 6 we define the requirements for a PRE scheme to be secure for a malicious server and define *token robustness*. In Section 5 we critique existing notions of unidirectionality, and present the first security definition for that considers reversibility. A formal

security definition for unidirectionality as it is currently considered we include in Section 5.2. We then use this to create a stronger definition for *best-achievable unidirectionality*. In Section 7 we define ciphertext origin authentication to provide authenticated PRE and discuss how to achieve this.

## 2 Preliminaries

We will always assume a PKE scheme which is IND-CPA secure (see Appendix A), and therefore has a randomised encryption algorithm.

**Definition 1 (multi-hop PRE scheme).** A (multi-hop) proxy re-encryption (PRE) scheme consists of the following algorithms:

- $\text{Setup}(1^\lambda) \rightarrow \text{param}$ : Takes the security parameter  $1^\lambda$  and outputs the set of parameters  $\text{param}$ , including a description of the message space  $\mathcal{M}$  and token space  $\mathcal{D}$ . We note that  $\text{param}$  is an input for all subsequent algorithms but we leave it out for compactness of notation.
- $\text{KeyGen}(1^\lambda) \xrightarrow{\$} (\text{pk}, \text{sk})$ : Generates a public-private key pair.
- $\text{Enc}(\text{pk}, m) \xrightarrow{\$} C$ : Given a public key  $\text{pk}$  and message  $m \in \mathcal{M}$ , returning the ciphertext  $C$ .
- $\text{Dec}(\text{sk}, C) \rightarrow m$  or  $\perp$ : Given a secret key  $\text{sk}$  and a ciphertext  $C$ , returns either a message  $m$  or the error symbol  $\perp$ .
- $\text{ReKeyGen}(\text{sk}_1, \text{pk}_2) \rightarrow \Delta_{1,2}$ : For two keypairs  $(\text{pk}_1, \text{sk}_1), (\text{pk}_2, \text{sk}_2)$ , outputs an update token  $\Delta_{1,2}$ .
- $\text{ReEnc}(\Delta_{1,2}, C_1) \rightarrow C_2$ : Takes a ciphertext  $C_1$  and a token  $\Delta_{1,2}$  and translates the ciphertext to output a new ciphertext  $C_2$ .

A PRE scheme is correct if for every message  $m \in \mathcal{M}$ , any sequence of  $\kappa$  keypairs  $(\text{pk}_1, \text{sk}_1), \dots, (\text{pk}_\kappa, \text{sk}_\kappa) \leftarrow \text{KeyGen}(1^\lambda)$ , all ciphertexts  $C_1 \xleftarrow{\$} \text{Enc}(\text{pk}_1, m)$ , and all re-encrypted ciphertexts  $\{C_i \leftarrow \text{ReEnc}(\Delta_{i-1,i}, C_{i-1})\}_{i=1}^\kappa$  where  $\Delta_{i-1,i} \leftarrow \text{ReKeyGen}(\text{sk}_{i-1}, \text{pk}_i)$ :

$$\Pr[\text{Dec}(\text{sk}_i, C_i) \neq m] \leq \text{negl}(\lambda),$$

for some negligible function  $\text{negl}(\lambda)$ .

In other words, ciphertexts decrypt correctly, including ciphertexts which are re-encryptions of other correct ciphertexts. Definition 1 can be intuitively adapted to the symmetric key setting, where token generation requires knowledge of both secret keys.

For PRE to be practically useful, it must be cheaper either computationally or in terms of the bandwidth used. This could refer either to size of messages exchanged or to price. For example, Amazon Web Services charges users for downloads but not uploads, giving some justification for schemes which upload data to cloud storage but require no download [10].

### 2.1 Additional properties

There are some additional properties which a PRE scheme can have. Whether or not these properties are required depends on the security model and the application. We define some of these properties below.

*Directionality:*

A PRE scheme is *bidirectional* if a re-encryption token from  $\text{pk}_i$  to  $\text{pk}_j$  can be used to re-encrypt from  $\text{pk}_j$  to  $\text{pk}_i$  and we write  $\Delta_{i \leftrightarrow j}$ . Otherwise, it is *unidirectional* and we write  $\Delta_{i \rightarrow j}$ . We reserve the notation  $\Delta_{i,j}$  for the general case where directionality is not specified.

If a PRE scheme is not bidirectional then it is *unidirectional* and we write  $\Delta_{i \rightarrow j}$ . We reserve the notation  $\Delta_{i,j}$  for the general case where directionality is not specified.

We take this description from existing work [3, 6, 10, 13] (see Section 2.2). To date this is not given as a separate definition, rather unidirectional/bidirectional PRE schemes are explicitly defined. We observe that it lacks the necessary formalism to be considered a security definition as needed for access control and key expiry. This motivates our formulating the current notion formally in Section 5.2, before giving a stronger definition in Section 5.3. However, we note that there are some applications where bidirectionality may be considered a necessary feature.

*Single/Multi-hop*: Some PRE schemes are *single-hop* meaning ciphertexts can only be re-encrypted once. Typically in these cases, ReEnc changes the format of the ciphertext and there are two decryption algorithms  $\text{Dec}_1, \text{Dec}_2$ , one for ciphertexts which can be re-encrypted and the other for ciphertexts which cannot be re-encrypted. In contrast a *multi-hop* scheme can be re-encrypted multiple times. Single-hop schemes only have limited use and are mainly considered for unidirectionality purposes. Since we focus on the practical application of access control as a motivation, we will assume multi-hop as a necessary requirement of a PRE scheme in the remainder of this work.

*Ciphertext dependence*: Informally ReKeyGen takes some additional information as input about the ciphertext that is to be re-encrypted under a new key. Let ReKeyGen in a PRE scheme be redefined to take additional information  $\tilde{C}$  about ciphertext  $C$  as input:  $\text{ReKeyGen}(\text{sk}_i, \text{pk}_j, \tilde{C}) \rightarrow \Delta_{i,j,C}$ . This PRE scheme is *ciphertext dependent* if for all  $C_i^1 \xleftarrow{\$} \text{Enc}(\text{pk}_i, m_1)$  and  $C_i^2 \xleftarrow{\$} \text{Enc}(\text{pk}_i, m_2)$  such that  $C_i^1 \neq C_i^2$ , and all re-encryption tokens  $\Delta_{i,j,C_i^1} \xleftarrow{\$} \text{ReKeyGen}(\text{sk}_i, \text{pk}_j, \tilde{C}_i^1)$ , then:

$$\Pr \left[ \text{Dec}(\text{sk}_j, \text{ReEnc}(\Delta_{i,j,C_i^1}, C_i^2)) = m_2 \right] \leq \text{negl}(\lambda). \quad (1)$$

In existing work [10] and in our scheme,  $\tilde{C}$  is the header of the ciphertext. For simplicity we will assume this is the case in the remainder of this paper. We note that not all applications require ciphertext dependence. For example in key expiry, all ciphertexts under the old key need to be re-encrypted. In subsequent definitions we will assume ciphertext dependence, but we note that these definitions can easily be extended to ciphertext independent schemes by setting  $\tilde{C} = \emptyset$ .

## 2.2 Existing Work

PRE was first introduced in [4]. The common approach to PRE in practice is the key encapsulation approach, where the ciphertext header contains the data encryption key  $k_D$ , encrypted with a key encryption key  $k_K$ . Typically, such schemes perform re-encryption by replacing  $k_K$  — so the ciphertext header now contains the same  $k_D$  encrypted with the new key encryption key  $k'_K$  and the body of the ciphertext remains unchanged:

$$C = ([k_D]_{k_K}, [m]_{k_D}) \xrightarrow{\text{ReEnc}(\Delta, C)} C' = ([k_D]_{k'_K}, [m]_{k_D})$$

The appeal of this approach is that it is efficient and can use hybrid encryption. It is used widely, for example in Amazon’s Key Management Service [1]. Whilst these schemes are simple and easy to implement, they do not completely re-randomise a ciphertext during re-encryption. A particular concern with the key encapsulation approach is that a malicious user can simply retain  $k_D$  and be able to decrypt the message, regardless of how many times it is re-encrypted.

Indistinguishability notions for PRE are described in [7] and [10]. A similar definition for CCA-security is used in [3, 6, 15]. Existing notions which imply complete re-randomisation for public-key

PRE include *unlinkability* in [6], and for symmetric key PRE include *ciphertext independence*<sup>1</sup> in [5] and UP-REENC security in [10].

Ciphertext dependence was first introduced in [10] for symmetric PRE, but has not yet been picked up by subsequent work. However their work does not explicitly consider unauthorised re-encryptions or unidirectionality.

One attempt to formalise the definition of directionality is given in [12], but they do not view directionality as a security definition, rather a classification of PRE schemes. They therefore define unidirectional PRE schemes and bidirectional PRE schemes as opposed to defining directionality separately. Furthermore the definition of unidirectionality in [12] assumes that a unidirectional scheme is single-hop, which is not necessarily the case. Other more recent work which informally describes unidirectionality say no PPT algorithm can output a token that can re-encrypt to the old key [10], whereas other works [16] say no PPT algorithm can output an equivalent encryption of the old ciphertext.

For a long time it was an open problem to create a scheme which is both unidirectional and multi-hop and, as such, there exist a number of PRE schemes which are unidirectional and single-hop [7, 15, 17]. They achieve unidirectionality by having two distinct levels of ciphertext which have different formats — level 2 for ciphertexts which can be re-encrypted, and level 1 for ciphertexts which cannot be re-encrypted. It is this format change which prevents a ciphertext from being re-encrypted more than once. This approach is undesirable, since it does not allow for multiple re-encryptions and therefore has limited practical application. Furthermore it does not convey how easy it is for a malicious server to reverse the re-encryption process. We discuss this in Section 5.

In many multi-hop schemes the number of re-encryptions is fixed, and the size of the ciphertext grows linearly with each re-encryption [6, 13]. The related problem of multi-hop unidirectional proxy re-signatures has a solution given in [14], however the message must be provided with the signature and thus such a scheme cannot be easily adapted to re-encryption. The first PRE scheme which is both unidirectional and multi-hop was given in [16], but does not address ciphertext dependence and current methods for achieving this cannot be applied to their scheme.

There does not appear to be existing work in either symmetric or public-key PRE that considers a malicious server which may perform unauthorised re-encryptions.

### 3 Indistinguishability

The most common notion of indistinguishability for PRE is that a re-encryption of a ciphertext should *preserve* the indistinguishability given by the underlying encryption scheme, which we call pres-IND-CPA (more details are given in Appendix A). This means that it is often not considered a requirement that re-encryption fully re-randomises a ciphertext. However full re-randomisation must be considered a necessary security property for applications such as access control (revocation) and key expiry. Definition 2 addresses this, based on UP-REENC in [10] adapted to the public-key setting. Definition 2 models a revoked user trying to distinguish a re-encrypted ciphertext from two potential original ciphertexts. In this game, the adversary needs to distinguish which of two possible ciphertexts is re-encrypted by the LR-ReEnc oracle. For stronger security, the adversary has access to a token generation oracle. OReKeyGen. The adversary is also given  $t$  secret keys, to model revocation scenarios where a user knows the old key and oracles have the restriction that they will not return tokens to a compromised key.

---

<sup>1</sup> We reserve this terminology for PRE schemes for which token generation is not specific to a given ciphertext as in Section 2.1.

ReEnc-IND-CPA $^A_{\mathcal{P}\mathcal{R}\mathcal{E}}(\lambda)$	LR-ReEnc( $i, j, C_0, C_1$ )
$param \leftarrow \text{Setup}(\lambda)$ $b \xleftarrow{\$} \{0, 1\}$ $(pk_1, sk_1), \dots, (pk_\kappa, sk_\kappa) \leftarrow \text{KeyGen}(1^\lambda)$ $\text{compromised} = \{sk_1, \dots, sk_\ell\}$ $b' \leftarrow \mathcal{A}^{\text{LR-ReEnc, OReKeyGen}}(1^\lambda, \text{compromised}, pk_1, \dots, pk_\kappa)$ <b>return</b> $b' = b$	<b>if</b> $ C_0  \neq  C_1 $ <b>or</b> $sk_j \in \text{compromised}$ <b>return</b> $\perp$ $\Delta_{i,j,C_b} \xleftarrow{\$} \text{ReKeyGen}(sk_i, pk_j, \tilde{C}_b)$ $C' \leftarrow \text{ReEnc}(\Delta_{i,j,C}, C_b)$ <b>return</b> $(C')$
	<hr style="width: 50%; margin-left: 0;"/> <b>if</b> $sk_j \in \text{compromised}$ <b>return</b> $\perp$ $\Delta_{i,j,C} \xleftarrow{\$} \text{ReKeyGen}(sk_i, pk_j, \tilde{C})$ <b>return</b> $\Delta_{i,j,C}$

Figure 1: ReEnc-IND-CPA game. Schemes that meet this definition must fully re-randomise a ciphertext upon re-encryption.

**Definition 2.** A PRE scheme  $\mathcal{P}\mathcal{R}\mathcal{E}$  is re-encryption indistinguishable against chosen plaintext attack (ReEnc-IND-CPA) if for all PPT adversaries  $\mathcal{A}$  there exists a negligible function  $\text{negl}(\lambda)$  such that:

$$\Pr \left[ \text{ReEnc-IND-CPA}^A_{\mathcal{P}\mathcal{R}\mathcal{E}}(\lambda) = 1 \right] \leq \frac{1}{2} + \text{negl}(\lambda), \quad (2)$$

where ReEnc-IND-CPA is given in Figure 1.

This definition can be easily extended to symmetric PRE by providing the adversary with encryption oracles for both keys, see [10]. If ReKeyGen is deterministic, the adversary can win by calling OReKeyGen( $i, j, C_0$ ) to obtain  $\Delta$ , compute ReEnc( $\Delta, C_1$ ) and compare this to LR-ReEnc( $i, j, C_0, C_1$ ). Since we consider re-randomisation a necessary property, from now on we will assume that ReKeyGen is randomised.

## 4 Token Robustness

This section defines *token robustness* - a stronger notion than ciphertext dependence. Informally, token robustness states that even with access to a token generation oracle, an adversary cannot create a new valid token which re-encrypts a ciphertext to a key it was never previously encrypted under. We cover re-encryption to keys a ciphertext was previously encrypted under in Section 5. Before we define token robustness, we need to define token validity.

Before we define token robustness, we need the following definition.

**Definition 3 (Token validity).** Let  $\delta \in \mathcal{D}$ . We say that  $\delta$  is a valid update token if there exist keys  $(pk_1, sk_1), (pk_1, sk_2) \leftarrow \text{KeyGen}(1^\lambda)$  and a ciphertext  $C \xleftarrow{\$} \text{Enc}(pk_1, m)$  for some  $m \in \mathcal{M}$  such that:

$$\Pr[\text{Dec}(sk_2, \text{ReEnc}(\delta, C)) \neq m] \leq \text{negl}(\lambda). \quad (3)$$

**Definition 4.** We say that a PRE scheme  $\mathcal{P}\mathcal{R}\mathcal{E}$  has token robustness if for all PPT adversaries  $\mathcal{A}$ , there exists a negligible function  $\text{negl}(\lambda)$  such that

$$\Pr \left[ \text{Tok-Rob}^A_{\mathcal{P}\mathcal{R}\mathcal{E}}(\lambda) = 1 \right] \leq \text{negl}(\lambda), \quad (4)$$

Tok- $\text{Rob}_A^{\text{PRE}}(\lambda)$	OEnc( $i, m$ )	OReKeyGen( $i, j, C$ )
$param \leftarrow \text{Setup}(1^\lambda)$ $(pk_1, sk_1), \dots, (pk_\kappa, sk_\kappa) \leftarrow \text{KeyGen}(1^\lambda)$ $generated = \emptyset$ <b>for</b> all $m \in \mathcal{M}$ <b>let</b> $chain(m) = \emptyset$ $(C_i^A, i, j, \delta^A) \leftarrow \mathcal{A}^{\text{OEnc, OReKeyGen}}(1^\lambda, pk_1, \dots, pk_\kappa)$ <b>if</b> $C_i^A \notin generated$ <b>return</b> $\perp$ $C_j^A \leftarrow \text{ReEnc}(\delta^A, C_i^A)$ <b>let</b> $m = \text{Dec}(sk_i, C_i^A), m' = \text{Dec}(sk_j, C_j^A)$ <b>if</b> $m = m'$ <b>and</b> $j \notin chain(m)$ <b>return</b> 1 <b>else return</b> 0	$C \xleftarrow{\$} \text{Enc}(pk_i, m)$ $chain(m).add\ i$ $generated.add\ C$ <b>return</b> $C$	$\Delta_{i,j,C} \xleftarrow{\$} \text{ReKeyGen}(sk_i, pk_j, \tilde{C})$ $chain(\text{Dec}(sk_i, C)).add\ j$ $C' \leftarrow \text{ReEnc}(\Delta_{i,j,C}, C)$ <b>if</b> $C \in generated$ $generated.add\ C'$ <b>return</b> $\Delta_{i,j,C}$

Figure 2: The token robustness game Tok- $\text{Rob}$

where Tok- $\text{Rob}$  is given in Figure 2.

In the Tok- $\text{Rob}$  game, the adversary has access to a token generation oracle OReKeyGen and an encryption oracle OEnc, and attempts to output a valid token  $\delta^A$  which re-encrypts a target ciphertext  $C_i^A$  from being under  $pk_i$  to being under  $pk_j$ . The list  $chain$  records which messages have been encrypted under which keys by adding the appropriate keys to  $chain(m)$  whenever OEnc or OReKeyGen are called. This means the adversary cannot trivially win by submitting a token which was the output of an OReKeyGen query. The problem of re-encrypting to a key the ciphertext was previously encrypted under is covered in our notion of *irreversibility* in Section 5. Another condition is the adversary's target ciphertext  $C_i^A$  must be an output of the OEnc oracle (or a re-encryption of such a ciphertext) to ensure that the adversary has no additional advantage from storing information created when encrypting the ciphertext. For example, in our scheme in Appendix B, encryption selects a random  $y$  and sets  $\tilde{C} = g^y$ . If the adversary encrypts the message for themselves then they learn  $y$ , which the server would not know in the cloud storage application. The function  $generated$  is used to keep track of these ciphertexts.

**Theorem 1.** *No ciphertext independent PRE scheme has token robustness.*

*Proof.* If a PRE scheme is not ciphertext dependent, then the same update token can be used to re-encrypt more than one ciphertext. In the Tok- $\text{Rob}$  game, let  $C_1 \leftarrow \text{OEnc}(i, m_1), C_2 \leftarrow \text{OEnc}(i, m_2)$ . Then the adversary can submit  $(i, j, C_1)$  to OReKeyGen to obtain  $\Delta_{i,j}$  and then submit  $(C_2, i, j, \Delta_{i,j})$  (in other words, set  $C^A = C_2, \delta^A = \Delta_{i,j}$ ). Since  $j \neq chain(m_2)$ , the adversary wins the game with probability 1.  $\square$

However, ciphertext dependence alone does not imply token robustness. For example, suppose that  $\text{ReKeyGen}(pk_i, sk_j, \tilde{C}_1)$  outputs  $\Delta_{i,j,C_1} = (\Delta_0 = \tilde{C}_1, \Delta_1)$ , and that  $\text{ReEnc}(\Delta, C)$  incorporates ciphertext dependence by verifying that  $\Delta_0 = \tilde{C}$ , with re-encryption only proceeding if this is true. Then an adversary can trivially craft a valid token for a different ciphertext by setting  $\delta^A = (\tilde{C}_2, \Delta_1)$ . We see that token robustness is a stronger notion than ciphertext dependence.

## 5 Directionality revisited

Recall that the existing definition of unidirectionality states that an update token  $\Delta_{1 \rightarrow 2}$  cannot be used to re-encrypt a ciphertext under  $pk_2$  to  $pk_1$ . We argue that this notion is not sufficient

for access control as it is not a security definition and does not consider reversal attacks where the adversary may have retained information on the old ciphertext. It also does not couple well with ciphertext dependence or token robustness. In this section we will elaborate on this claim before offering a stronger security definition, which we call  $\bar{\lambda}$ -irreversible.

### 5.1 Problems with traditional directionality

Unidirectionality is required in a number of applications for security reasons, despite the fact that it is not currently defined as a security property. In Section 5.2, we give a security definition for the current understanding of unidirectionality. However, this notion does not consider reversal attacks where the server has the new ciphertext and the information used to perform the re-encryption. However this is a possible means of a malicious server performing an unauthorised re-encryption and so warrants further consideration. Whilst this is not of any concern for email forwarding, it is an important consideration for access control as reversing a re-encryption can mean regranteeing access to a revoked user. Particularly for ciphertext dependent schemes where update tokens are randomised, reversal attacks become more significant as each re-encryption token should only be able to reverse one specific ciphertext. Therefore existing notions of unidirectionality do not couple well with ciphertext dependence.

There are schemes for which storage of the update token alone cannot reverse the re-encryption, but retaining elements of the update token and elements of the original ciphertext can reverse a re-encryption. We provide more explicit examples in Section 5.4. If storing some component of the original ciphertext makes a reversal attack successful, then we should assume that the adversary will do so, especially if the amount of storage needed is at most the size of the update token (which they are presumed to retain in existing notions of unidirectionality). This shows that unidirectional schemes do not prevent an adversary from reversing a re-encryption. For practical reasons, header values and update tokens are often designed to be small, and thus they can easily be retained. Current models of unidirectionality permit an adversary to retain the update token in order to attempt to re-encrypt; we argue that there is no reason to restrict the information an adversary may store to just update tokens when adversarial success may be greater when considering other information that is already available to attackers, particularly in the case of a malicious server. First we give a formal definition for the traditional notion of unidirectionality Section 5.2 before giving a definition which considers reversibility Section 5.3.

### 5.2 Traditional unidirectionality as a security notion

Some applications for PRE such as key rotation require unidirectionality as a security definition, even though it is not currently considered a security definition. In particular, schemes which have additional security properties such as ReEnc-IND-CPA where tokens are randomised, or schemes which ciphertext dependent, the intuition for unidirectionality becomes less intuitive in terms of what the adversary has access to and what they are trying to achieve. Here we give a security definition for unidirectionality as it is currently understood.

In the unidirectionality game, the challenge oracle takes as input a number  $N$  which indicates the number of times the challenge ciphertext is to be re-encrypted before being given to the adversary. In other words,  $N$  indicates the level of the ciphertext. This is to accommodate both single and multi-hop PRE schemes.

**Definition 5.** A PRE scheme  $\mathcal{PRE}$  is unidirectional if for all PPT algorithms  $\mathcal{A}$  there exists a negligible function  $\text{negl}(\lambda)$  such that

$$\Pr \left[ \text{UNI}_{\mathcal{PRE}}^{\mathcal{A}}(\lambda) = 1 \right] \leq \text{negl}(\lambda). \quad (5)$$



$\text{UNI}_{\mathcal{A}, \lambda}^{\mathcal{PRE}}(\lambda)$	$\text{OReKeyGen}(C)$
$param \leftarrow \text{Setup}(\lambda)$	$\Delta_{i,j,C} \leftarrow \text{ReKeyGen}(sk_i, pk_j, \tilde{C})$
$(pk_1, sk_1), \dots, (pk_\kappa, sk_\kappa) \xleftarrow{\$} \text{KeyGen}(1^\lambda)$	<b>return</b> $\Delta_{i,j,C}$
$(i, j) \leftarrow \mathcal{A}_0(1^\lambda, pk_1, \dots, pk_\kappa)$	
<b>challenge</b> = $\{\emptyset\}$	
$C_j, C_i^A \leftarrow \mathcal{A}_1^{\text{ChalGen}, \text{OReKeyGen}}(1^\lambda)$	
<b>if</b> $C_j \notin \text{challenge}$ <b>return</b> $\perp$	
<b>return</b> $\text{Dec}(sk_i, C_i^A) = \text{Dec}(sk_j, C_j)$	
$\text{ChalGen}(N)$	
$m \xleftarrow{\$} \mathcal{M}$	
<b>if</b> $N = 0$	
$C_j \leftarrow \text{Enc}(pk_j, m)$	
<b>else</b>	
$k_0, \dots, k_{N-1} \xleftarrow{\$} \{0, \dots, \kappa\}$	
$k_N = j$	
$C_{k_0} \leftarrow \text{Enc}(pk_{k_0}, m)$	
<b>for</b> $n = 1, \dots, N - 1$	
$\Delta_{k_{n-1}, k_n, C_{k_{n-1}}} \leftarrow \text{ReKeyGen}(sk_{k_{n-1}}, pk_{k_n}, \tilde{C}_{k_{n-1}})$	
$C_{k_n} \leftarrow \text{ReEnc}(\Delta_{k_{n-1}, k_n, \tilde{C}_{k_{n-1}}}, C_{k_{n-1}})$	
<b>challenge.add</b> $C_j$	
<b>return</b> $C_j$	

Figure 3: The unidirectionality game UNI. The adversary must re-encrypt a ciphertext under  $pk_j$  to a ciphertext under  $pk_i$ , given an oracle which outputs tokens  $\Delta_{i,j}$ . The challenge oracle outputs a ciphertext under  $pk_j$ , which can have been re-encrypted a number of times  $N$  as specified by the adversary. In a single-hop game, this is limited to  $N \in \{0, 1\}$ .

A PRE scheme  $\mathcal{PRE}$  is bidirectional if there exists a PPT algorithm  $\mathcal{A}$  that wins the unidirectionality game UNI with overwhelming probability.

We believe that this definition formally captures the intuition behind unidirectionality as it is currently understood in the literature. We hope that for applications where unidirectionality is a security requirement, this definition can be used to assess the suitability of potential PRE schemes.

### 5.3 Directionality reconsidered

Now that we have argued that current notions of unidirectionality are not suitable for access control, we present a security definition for reversal attacks. Before we define this, we explain the key principles behind the motivation of the definition.

**Principle 1:** *Malicious storage cannot be prevented.* It is impossible for one part to prevent

another from storing extra information without additional assumptions. In particular we cannot prevent the server from retaining the old ciphertext.

**Principle 2:** *The amount of storage needed to reverse a process has a lower bound.* Whilst we cannot prevent a malicious server from retaining an old ciphertext, we can ensure there is no ‘easier’ way for them to obtain the old ciphertext. By ‘easier’, we mean that the server needs significantly less storage than keeping the components of the original ciphertext that were updated. This is similar to the motivation behind an *economically rational server* considered in [9].

This definition is important in that if a scheme is token robust and we can prove that the only way of reversing a re-encryption is by storing a state the size of the original ciphertext then this is the best notion of unidirectionality that can be achieved without assuming that the adversary honestly deletes old ciphertext, which cannot be relied upon. In Section 6 we use this definition to define *best-achievable unidirectionality*. By considering unidirectionality in this way, the problem of creating a unidirectional multi-hop PRE scheme may be solved more easily using token robustness and could therefore lead to more unidirectional multi-hop schemes that are practically implementable.

An assumption often made in the access control literature is that once a message (or ciphertext) is known, it is known forever and so there is no point in considering leakage from the old ciphertext. However, this assumption does not always make sense in cloud storage applications when revoked users are considered. One of the most common reasons for outsourcing storage is limited local space. Therefore, a user’s ability to retrieve a file at one point does not necessarily mean they have the capacity to retain all files they have ever had access to. Revoked users may try to decrypt a file previously encrypted under a key known to them. In many existing schemes, if a revoked user is able to retain limited information about the old ciphertext, they have an advantage in creating a re-encryption under the old key.

We now define a reversal attack game which takes into account the amount of information the adversarial server may have retained during the re-encryption process using a storage parameter  $\bar{\lambda}$ . The following game has an adversary in three stages. All three adversaries receive the security parameter  $\lambda$ , storage parameter  $\bar{\lambda}$ , public keys and system parameters as input. The first stage adversary  $\mathcal{A}_0$  receives a randomly chosen message and decides which keys should be used for encryption and re-encryption. The second stage adversary  $\mathcal{A}_1$  receives the ciphertext and update token and determines what should be retained in the state  $st_{\mathcal{A}}$ , which is bounded by a storage parameter  $\bar{\lambda}$ . Note this adversary never receives the message. Since this adversary knows the storage bound, it can compute many potential states before selecting which one will be passed on to  $\mathcal{A}_2$ . The final adversary  $\mathcal{A}_2$  receives the re-encrypted ciphertext  $C'$  and state  $st_{\mathcal{A}}$ , and uses this to try to output a ciphertext which is an encryption of the same message under the original key. This adversary never receives the message, original ciphertext or the update token — they only receive the information retained by  $\mathcal{A}_1$ . Note that this does not need to be the original ciphertext, it can be any ciphertext equivalent to the original. This emulates the scenario where the server must decide how much information to retain about the old ciphertext and update token, before later attempting to reverse the re-encryption (or revert to an equivalent ciphertext).

**Definition 6.** *Given a PRE scheme  $\mathcal{PRE}$ , let  $s$  be the size of the components in ciphertexts as established by the scheme and let  $c$  be the number of ciphertext components updated by  $\text{ReEnc}$ . Then for  $\bar{\lambda} \in \{0, s, \dots, cs\}$ , we define the advantage of an adversary  $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1, \mathcal{A}_2)$  in winning the  $\text{Rev-ReEnc}$  game given in Figure 4 as:*

$$\text{adv}_{\mathcal{A}, \bar{\lambda}}^{\text{Rev-ReEnc}}(\lambda) = \left| \Pr \left[ \text{Rev-ReEnc}_{\mathcal{A}, \bar{\lambda}}^{\mathcal{PRE}}(\lambda) = 1 \right] - \frac{1}{2^{cs - \bar{\lambda}}} \right|, \quad (6)$$

where  $\frac{1}{2^{cs - \bar{\lambda}}}$  is the probability that an adversary who has retained  $\bar{\lambda}$  bits of (the updatable components of)  $C$  can correctly guess the remaining bits.

$\text{Rev-ReEnc}_{\mathcal{A}, \bar{\lambda}}^{\text{PRE}}(\lambda)$
$param \leftarrow \text{Setup}(\lambda)$
$(\text{pk}_1, \text{sk}_1), \dots, (\text{pk}_\kappa, \text{sk}_\kappa) \xleftarrow{\$} \text{KeyGen}(1^\lambda)$
$m \xleftarrow{\$} \mathcal{M}$
$(i, j) \leftarrow \mathcal{A}_0(1^\lambda, 1^{\bar{\lambda}}, \text{pk}_1, \dots, \text{pk}_\kappa, m)$
$C \xleftarrow{\$} \text{Enc}(\text{pk}_i, m)$
$\Delta_{i,j,C} \xleftarrow{\$} \text{ReKeyGen}(\text{sk}_i, \text{pk}_j, \tilde{C})$
$st_{\mathcal{A}} \leftarrow \mathcal{A}_1(1^\lambda, 1^{\bar{\lambda}}, \text{pk}_1, \dots, \text{pk}_\kappa, \Delta, C)$
<b>if</b> $ st_{\mathcal{A}}  > \bar{\lambda}$ <b>return</b> $\perp$
$C' \leftarrow \text{ReEnc}(\Delta_{i,j,C}, C)$
$C_{\mathcal{A}} \leftarrow \mathcal{A}_2(1^\lambda, 1^{\bar{\lambda}}, \text{pk}_1, \dots, \text{pk}_\kappa, st_{\mathcal{A}}, C')$
<b>return</b> $\text{Dec}(\text{sk}_i, C_{\mathcal{A}}) = \text{Dec}(\text{sk}_j, C')$

Figure 4: The reversal game Rev-ReEnc.

We say that a proxy re-encryption scheme  $\text{PRE}$  is  $\bar{\lambda}$ -irreversible if for all PPT adversaries  $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1, \mathcal{A}_2)$ , the advantage of winning the game is negligible:

$$\text{adv}_{\mathcal{A}, \bar{\lambda}}^{\text{Rev-ReEnc}}(\lambda) \leq \text{negl}(\lambda). \quad (7)$$

Conversely, a PRE scheme is  $\bar{\lambda}$ -reversible if there exists a PPT algorithm  $\mathcal{A}$  that can win the  $\text{Rev-ReEnc}_{\mathcal{A}, \bar{\lambda}}^{\text{PRE}}(\lambda)$  game with state  $st_{\mathcal{A}}$  of size at most  $\bar{\lambda}$ , with non-negligible probability.

We briefly note that constructing this notion as an indistinguishability game is difficult since the state which the adversary outputs is not fixed. For example, if the adversary stores a truncated hash of the original ciphertext then the game cannot compute another ciphertext which makes indistinguishability difficult. Since this definition aims to convey that an adversary should not be able to re-encrypt back to the old key, we do not consider the lack of indistinguishability as hindering this.

We now formulate a definition for *maximum irreversibility*. Informally, the amount of storage needed to reverse a re-encryption is at least the size of the old ciphertext components that were updated.

**Definition 7.** A PRE scheme is maximally irreversible if it is  $cs$ -reversible, where  $c$  is the number of the number of ciphertext components updated by  $\text{ReEnc}$  and  $s$  is the component size.

Clearly, that *maximum irreversibility* is stronger than traditional notions of unidirectionality because it covers directionality attacks for ciphertext dependent schemes in addition to ciphertext independent schemes. We later use this notion together with token robustness (which is stronger than ciphertext dependence) to form a definition for maximum unidirectionality.

## Observations

1. The storage bound  $\bar{\lambda}$  can be considered similarly to the security parameter  $\lambda$  in that the larger  $\bar{\lambda}$  is, the more secure the scheme is. However, even small values for  $\bar{\lambda}$  are still meaningful since

they convey how easy it is to reverse a re-encryption and can therefore be used to compare different schemes.

2. The most useful values which  $\bar{\lambda}$  can take are  $\bar{\lambda} = |\Delta|$  as this is comparable to traditional bidirectionality or  $\bar{\lambda} = cs$  as this makes a scheme maximally irreversible. In general, useful values are in the range  $|\Delta| \leq \bar{\lambda} \leq |C|$ .
3. If a scheme is both ReEnc-IND-CPA and maximally irreversible then it is  $|C|$ -irreversible.
4. All traditionally bidirectional schemes can be shown to be  $|\Delta|$ -reversible, but there also exist  $|\Delta|$ -reversible schemes which are *not* traditionally bidirectional, as we shall see in Section 5.4. Since more attacks are covered, saying that a scheme is  $|\Delta|$ -reversible is stronger than saying it is bidirectional in the traditional sense.
5. Our definition also applies to schemes where the ciphertext is not fully re-randomised, since best-achievable unidirectionality only requires that the state is the size of the number of components which are updated as opposed to the entire ciphertext.
6. We assume that the ciphertext is as compact as it could be, so for all ciphertexts  $C$  there is no compression function which allows  $C$  to be stored as  $C^*$  with  $|C^*| < |C|$ . For example, an IND-CPA scheme will have this property, as all ciphertexts are indistinguishable from random strings, so only storing a subset of the bits means the adversary can do no better than guessing to obtain the remainder of the ciphertext. If the ciphertexts can be compressed, we can define a compressed version of the PRE scheme by having encryption compress the ciphertext, decryption decompress it and have the re-encryption function perform decompression and compression. Hence this definition can be used for all PRE schemes.

#### 5.4 Existing schemes under the new definition

*Some traditionally unidirectional schemes are  $|\Delta|$ -reversible.* In both [10] and [5], the update token consists of the new header and another value used to change the body of the ciphertext using an arithmetic operation. We can generalise this by saying  $\Delta = (\Delta_0, \Delta_1)$ , where  $\Delta_0 = \tilde{C}'$  and  $(\Delta_1)^{-1}$  is easily computable. To reverse the re-encryption, the adversary  $\mathcal{A}_1$  retains the old header  $\tilde{C}$ , and computes the inverse of  $\Delta_1$ , setting  $st_{\mathcal{A}} = (\tilde{C}, \Delta_1^{-1})$ . This is equivalent to  $st_{\mathcal{A}} = \Delta^{-1}$ . Then  $\mathcal{A}_2$  can recover  $C \leftarrow \text{ReEnc}(st_{\mathcal{A}}, C')$  to win the game. Note that  $\mathcal{A}$  does not need to retain  $\Delta_0$  as this is contained in the new ciphertext. The state  $st_{\mathcal{A}}$  is clearly the same size as  $\Delta = (\tilde{C}', \Delta_1)$ , and we can therefore consider such schemes to be  $|\Delta|$ -reversible. Since any adversary willing to store information of size up to  $|\Delta|$  will not restrict themselves to retaining  $\Delta$  alone, our definition better reflects the intuition behind bidirectionality. In particular, because [10] is ciphertext-dependent, it should be considered bidirectional under these realistic assumptions.

*Some existing bidirectional schemes are maximally irreversible.* In the multi-hop PRE scheme of [6], ReKeyGen takes two secret keys as input and the ciphertext includes a number of components including  $B = (g^a)^r$ , where  $\text{pk} = g^a, \text{sk} = a$  and  $r \xleftarrow{\$} \mathbb{Z}_q$ . The re-encryption token takes as input two secret keys  $a, b$  and outputs  $\Delta_{a,b} = b/a$ , which is then used to update  $B$  and no other part of the ciphertext. Since both the ciphertext element  $B$  and the re-keying token  $\Delta_{a,b}$  are integers modulo  $q$ , an adversary hoping to reverse the re-encryption by storing  $\Delta_{a,b}$  could have simply retained  $B$ . Particularly for applications where there is one message per keypair, the server would need to store one token per re-encrypted ciphertext in which case they could have retained every original ciphertext <sup>2</sup>. Similarly, the original symmetric proxy re-encryption scheme [4] may also be considered maximally irreversible.

<sup>2</sup>As the value  $B$  is unique for each ciphertext, retaining  $B$  for one ciphertext does not allow a different ciphertext to be re-encrypted whereas the update token can re-encrypt any ciphertext in either direction. This further demonstrates why token robustness is a necessary requirement.

## 6 Proxy Re-Encryption in the Malicious model

We now describe the requirements for a PRE scheme to be secure in the malicious model. We discuss some conditions which apply to re-encryption generally, before explaining the stronger conditions specific to our setting. Clearly, *correctness* is a necessary property of all PRE schemes. We consider ReEnc-IND-CPA as another necessary condition for revocation and key expiry, despite the fact that this is not the case in much existing work [3, 4, 6, 7, 14].

In the malicious model, we must ensure that giving the server the ability to perform some re-encryptions does not mean they can perform unauthorised re-encryptions. In particular, we want our setting to consider revoked users who are honest-but-curious in that they may try to decrypt re-encrypted ciphertexts, but not collude with the server directly. We thus require a means of ensuring that only authorised re-encryptions are possible. The inability to perform unauthorised re-encryptions breaks down to two necessary properties:

**Maximal irreversibility:** *The token used to perform a re-encryption cannot be used to reverse that re-encryption.* This is particularly necessary when considering revocation and key expiry. If re-encryption has been performed to revoke access, then reversing that re-encryption regrants access to the revoked user. Our definition of maximal irreversibility conveys this under realistic assumptions.

**Token Robustness:** *No matter how many re-encryption tokens the server sees, the server cannot use these to form a token which encrypts a ciphertext to a new key.* This means the adversary is unable to share messages with users who have not had access to them before.

Ciphertext dependence is reasonably trivial to achieve for ElGamal-based schemes by having the randomness used to encrypt the message input to ReKeyGen. We build on this existing technique [10] to create a token robust scheme. We combine these definitions to form the following definition which is a requirement for a PRE scheme used to enforce changes to access control policy on a malicious server.

**Definition 8.** *A PRE scheme is best-achievable unidirectional if it is both token robust and maximally irreversible.*

Token robustness implies that a token  $\Delta_{i,j,C_1}$  cannot be used to re-encrypt a ciphertext  $C_2 \stackrel{\$}{\leftarrow} \text{Enc}(pk_j, m)$  where  $C_2 \neq C_1$  (except with negligible probability), which covers the traditional notion of unidirectionality [full paper]. Coupled with maximal irreversibility, this means that given a re-encrypted ciphertext  $C_2$  under  $pk_j$ , the only way that the adversary can produce a ciphertext  $C_1$  such that  $\text{Dec}(sk_i, C_1) = \text{Dec}(sk_j, C_2)$  where  $pk_i$  is the original key is by retaining a state the size of the original ciphertext during the re-encryption process.

We show in Appendix B it is possible to have a scheme which is best-achievable unidirectional using a simple adaptation of ElGamal-based PRE, and prove its security under the CDH assumption.

## 7 Ciphertext Origin Authentication

We now revisit the traditional notion of data origin authentication and how this needs tweaking for PRE. Traditionally, data origin authentication is intended for settings where the party who created the message also encrypts it. However, for PRE this is not always the case. Particularly in applications where more than one party shares an encryption key, proof of having used this key is not sufficient to authenticate the encryptor / re-encryption initiator.

It may be beneficial in some applications to use individual signature keys to verify specific identities, in addition to the encryption keys for confidentiality. Both signatures and PRE could

be combined to create an authenticated PRE scheme. This is useful in auditing changes to access control policy, or enabling users to verify which user has revoked their access.

Now that we have outlined this distinction, we formulate the notion of *Ciphertext Origin Authentication* (COA).

## 7.1 Authentication with corrupted users

Many authenticated encryption schemes including Signcryption [18] implement data origin authentication by having a check during decryption which terminates the process if the check fails. Such a check is also made in [10] which to our knowledge is the only ReEnc-IND-CPA scheme to provide authentication. However, if corrupted users are being considered then honest termination of a process is not guaranteed. Therefore compulsory COA provides stronger security against users who want to change policy without being caught, as the message can *only* be derived if identity is verified correctly. We call this *correctness upon verification* and consider it a secondary goal.

## 7.2 Correctness upon verification

The most intuitive way of proving which entity encrypted a message is to show proof of knowledge of the secret information used to form the new ciphertext. In our scheme outlined in Figure 6 this is the value  $y$ . However, the COA check must not leak  $y$  as this will enable decryption using the public key. Therefore, we need the initiator to prove that they know  $y$  without revealing it.

Recall the basic ElGamal signature scheme:

1.  $\text{Sign}(\mathbf{x}, m) \rightarrow \sigma = (r, s)$ :  $r = g^k$  for  $k \xleftarrow{\$} \mathbb{Z}_p$  and  $s = (h(m) - \mathbf{x}r)k^{-1} \pmod{p-1}$ , for hash function  $h()$ .
2.  $\text{Verify}(\mathbf{X}, \sigma, m) \rightarrow \{0, 1\}$ : if  $g^{h(m)} = \mathbf{X}^r \cdot r^s$  return 1, else return 0.

We can obtain a non-optional COA check by replacing  $g^y$  in our original scheme (Figure 6) with an ElGamal signature on  $y$  signed using the initiator's public key. We also adapt the Verify algorithm so that it derives a specific value  $a$ . We call the resulting algorithm VerRetrieve.

$\text{sigKeyGen}(1^\lambda)$	$\text{Sign}(\mathbf{x}, y)$	$\text{VerRetrieve}(\mathbf{X}, \sigma)$
$\mathbf{x} \xleftarrow{\$} \mathbb{Z}_p^*$	$k \xleftarrow{\$} \mathbb{Z}_p, r = g^k$	$(r, s) = \sigma$
$\mathbf{X} = g^{\mathbf{x}}$	$s = (y - \mathbf{x}r)k^{-1} \pmod{p-1}$	$a = \mathbf{X}^r \cdot r^s$
$\text{ssk} = \mathbf{x}, \text{svk} = \mathbf{X}$	$\sigma = (r, s)$	<b>return</b> $a$
<b>return</b> $(\mathbf{x}, \mathbf{X})$	<b>return</b> $\sigma$	

Ciphertexts now have the form  $C = (\sigma, m \cdot g^{xy})$ . By the correctness of ElGamal signatures, VerRetrieve should return  $a = g^y$ , since  $\mathbf{X}^r \cdot r^s = g^{\mathbf{x}r} \cdot g^{k(y-\mathbf{x}r)k^{-1}} = g^y$ .

Since obtaining  $g^y$  is necessary for decryption, and this can only be learned by successfully computing the VerRetrieve operation described above using the encryptor's public key, so verification is not optional.

However, if the scheme is only adapted with the change outlined above then there is no confirmation that the obtained  $g^y$  was correct, as the decryptor has no means verifying the message. Therefore in order to have COA, we also need a message integrity check.

We propose adapting the encryption mechanism to replace  $\bar{C} = m \cdot g^{xy}$  with  $\bar{C} = (g^{x\hat{y}h(m)}, m \cdot g^{xy})$  where  $\hat{y} \xleftarrow{\$} \mathbb{Z}_p^*$  and adding the matching signature to the header. See Appendix B for full details.

### 7.3 COA in other schemes

Our notion of COA is not restricted to our scheme. It is sufficient to create a signature using the encryptor’s signing key and the randomness used to form the ciphertext and changing re-encryption and decryption accordingly. We now demonstrate how to extend existing schemes to offer COA. For other ElGamal-based schemes including [4], a similar adaptation can be made.

We now propose an extension to ReCrypt [10] (described in Appendix C). Whilst this scheme uses symmetric PRE for re-encryption, we believe that COA is still valid for symmetric encryption. In ReCrypt, the message integrity check is optional and therefore it would be arguable that an optional COA check suffices in their model. We describe a simple extension which creates an optional check here, and a non-optional check in Appendix C.

For a basic extension, the ciphertext should replace  $\chi$  with  $(\chi, \sigma = \text{Sign}(\mathbf{x}, \chi))$  and decryption should include the step  $\text{Verify}(\mathbf{X}, \sigma, \chi)$  and only return  $m$  if this outputs 1. We note that this ReCrypt may be preferable to our scheme from Figure 6 depending on the application, as it permits the re-encryption of longer messages, with the caveat that it is not best-achievable unidirectional.

## 8 Conclusions and Open Problems

We revisited the notion of unidirectionality in PRE schemes and provided a formal security definition that covers reversal attacks. We have shown how, under this new definition, existing PRE schemes which are considered traditionally bidirectional may be considered unidirectional and vice versa. We also outlined properties a PRE scheme needs to be considered secure in the malicious model, in particular defining token robustness — the inability of the server to forge update tokens. Finally, we introduced a new notion of ciphertext origin authentication for authenticated PRE and discuss how to implement this. Schemes meeting these definitions are given in the appendices.

A useful extension of this work is to create a best-achievable unidirectional token robust scheme which can be used for longer messages. This could be achieved trivially by having the update token be as long as the ciphertext, but this is an inefficient solution, going against the motivations of outsourcing re-encryption. Developing a best-achievable unidirectional PRE scheme with small update tokens has similar challenges to white-box cryptography and obfuscation [11]. Another related challenge is to create a best-achievable unidirectional token robust symmetric PRE scheme. We leave these as open problem.

We also leave to future work creating a CCA secure PRE scheme which is best-achievable unidirectional and token robust, as well as defining what it means for a PRE scheme to be post-compromise secure [8] and creating a post-compromise secure PRE scheme.

## 9 Acknowledgements

Many thanks to my supervisor Keith Martin for guiding my ideas into a worthwhile piece of work. Also many thanks to those who gave up their time to proofread my work, especially Christian Janson and James Alderman.

## References

1. Amazon Web Services: Protecting data using client-side encryption (2017), <http://docs.aws.amazon.com/AmazonS3/latest/dev/UsingClientSideEncryption.html>

2. Ateniese, G., Benson, K., Hohenberger, S.: Key-private proxy re-encryption. In: Fischlin, M. (ed.) Topics in Cryptology - CT-RSA 2009, The Cryptographers' Track at the RSA Conference 2009, San Francisco, CA, USA, April 20-24, 2009. Proceedings. Lecture Notes in Computer Science, vol. 5473, pp. 279–294. Springer (2009), [https://doi.org/10.1007/978-3-642-00862-7\\_19](https://doi.org/10.1007/978-3-642-00862-7_19)
3. Ateniese, G., Fu, K., Green, M., Hohenberger, S.: Improved proxy re-encryption schemes with applications to secure distributed storage. *ACM Trans. Inf. Syst. Secur.* 9(1), 1–30 (2006), <http://doi.acm.org/10.1145/1127345.1127346>
4. Blaze, M., Bleumer, G., Strauss, M.: Divertible protocols and atomic proxy cryptography. In: Nyberg, K. (ed.) Advances in Cryptology - EUROCRYPT '98, International Conference on the Theory and Application of Cryptographic Techniques, Espoo, Finland, May 31 - June 4, 1998, Proceeding. Lecture Notes in Computer Science, vol. 1403, pp. 127–144. Springer (1998), <https://doi.org/10.1007/BFb0054122>
5. Boneh, D., Lewi, K., Montgomery, H.W., Raghunathan, A.: Key homomorphic prfs and their applications. In: Canetti, R., Garay, J.A. (eds.) Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I. Lecture Notes in Computer Science, vol. 8042, pp. 410–428. Springer (2013), [https://doi.org/10.1007/978-3-642-40041-4\\_23](https://doi.org/10.1007/978-3-642-40041-4_23)
6. Canetti, R., Hohenberger, S.: Chosen-ciphertext secure proxy re-encryption. In: Ning, P., di Vimercati, S.D.C., Syverson, P.F. (eds.) Proceedings of the 2007 ACM Conference on Computer and Communications Security, CCS 2007, Alexandria, Virginia, USA, October 28-31, 2007. pp. 185–194. ACM (2007), <http://doi.acm.org/10.1145/1315245.1315269>
7. Chow, S.S.M., Weng, J., Yang, Y., Deng, R.H.: Efficient unidirectional proxy re-encryption. In: Bernstein, D.J., Lange, T. (eds.) Progress in Cryptology - AFRICACRYPT 2010, Third International Conference on Cryptology in Africa, Stellenbosch, South Africa, May 3-6, 2010. Proceedings. Lecture Notes in Computer Science, vol. 6055, pp. 316–332. Springer (2010), [https://doi.org/10.1007/978-3-642-12678-9\\_19](https://doi.org/10.1007/978-3-642-12678-9_19)
8. Cohn-Gordon, K., Cremers, C.J.F., Garratt, L.: On post-compromise security. In: IEEE 29th Computer Security Foundations Symposium, CSF 2016, Lisbon, Portugal, June 27 - July 1, 2016. pp. 164–178. IEEE Computer Society (2016), <https://doi.org/10.1109/CSF.2016.19>
9. van Dijk, M., Juels, A., Oprea, A., Rivest, R.L., Stefanov, E., Triandopoulos, N.: Hourglass schemes: how to prove that cloud files are encrypted. In: Yu, T., Danezis, G., Gligor, V.D. (eds.) the ACM Conference on Computer and Communications Security, CCS'12, Raleigh, NC, USA, October 16-18, 2012. pp. 265–280. ACM (2012), <http://doi.acm.org/10.1145/2382196.2382227>
10. Everspaugh, A., Paterson, K.G., Ristenpart, T., Scott, S.: Key rotation for authenticated encryption. *IACR Cryptology ePrint Archive* 2017, 527 (2017), <http://eprint.iacr.org/2017/527>
11. Hohenberger, S., Rothblum, G.N., Shelat, A., Vaikuntanathan, V.: Securely obfuscating re-encryption. *J. Cryptology* 24(4), 694–719 (2011), <https://doi.org/10.1007/s00145-010-9077-7>
12. Ivan, A., Dodis, Y.: Proxy cryptography revisited. In: Proceedings of the Network and Distributed System Security Symposium, NDSS 2003, San Diego, California, USA. The Internet Society (2003), <http://www.isoc.org/isoc/conferences/ndss/03/proceedings/papers/14.pdf>
13. Liang, X., Cao, Z., Lin, H., Shao, J.: Attribute based proxy re-encryption with delegating capabilities. In: Li, W., Susilo, W., Tupakula, U.K., Safavi-Naini, R., Varadharajan, V. (eds.) Proceedings of the 2009 ACM Symposium on Information, Computer and Communications Security, ASIACCS 2009, Sydney, Australia, March 10-12, 2009. pp. 276–286. ACM (2009), <http://doi.acm.org/10.1145/1533057.1533094>
14. Libert, B., Vergnaud, D.: Multi-use unidirectional proxy re-signatures. In: Ning, P., Syverson, P.F., Jha, S. (eds.) Proceedings of the 2008 ACM Conference on Computer and Communications Security, CCS 2008, Alexandria, Virginia, USA, October 27-31, 2008. pp. 511–520. ACM (2008), <http://doi.acm.org/10.1145/1455770.1455835>
15. Libert, B., Vergnaud, D.: Unidirectional chosen-ciphertext secure proxy re-encryption. In: Cramer, R. (ed.) Public Key Cryptography - PKC 2008, 11th International Workshop on Practice and Theory in Public-Key Cryptography, Barcelona, Spain, March 9-12, 2008. Proceedings. Lecture Notes in Computer Science, vol. 4939, pp. 360–379. Springer (2008), [https://doi.org/10.1007/978-3-540-78440-1\\_21](https://doi.org/10.1007/978-3-540-78440-1_21)



pres-IND-CPA $_{\mathcal{PRE}}^A(\lambda)$	LR( $i, m_0, m_1$ )	OReKeyGen( $i, j, C$ )
$param \leftarrow \text{Setup}(1^\lambda)$ $b \xleftarrow{\$} \{0, 1\}$ $(pk_1, sk_1), \dots, (pk_\kappa, sk_\kappa) \leftarrow \text{KeyGen}(1^\lambda)$ $b' \leftarrow \mathcal{A}^{\text{LR}, \text{OReKeyGen}, \text{OReEnc}}(1^\lambda, pk_1, \dots, pk_\kappa)$ <b>return</b> $b' = b$	<b>if</b> $ m_0  \neq  m_1 $ <b>return</b> $\perp$ $C \xleftarrow{\$} \text{Enc}(pk_i, m_b)$ <b>return</b> $C$	$\Delta_{i,j,C} \xleftarrow{\$} \text{ReKeyGen}(sk_i, pk_j, \tilde{C})$ <b>return</b> $\Delta_{i,j,C}$
	<hr/> OReEnc( $i, j, C$ )	
		<hr/> $\Delta_{i,j,C} \xleftarrow{\$} \text{ReKeyGen}(sk_i, pk_j, \tilde{C})$ $C' \leftarrow \text{ReEnc}(\Delta_{i,j,C}, C)$ <b>return</b> $C'$

Figure 5: The pres-IND-CPA game which reflects the most common notion of indistinguishability for PRE.

16. Polyakov, Y., Rohloff, K., Sahu, G., Vaikuntanthan, V.: Fast proxy re-encryption for publish/subscribe systems. IACR Cryptology ePrint Archive 2017, 410 (2017), <http://eprint.iacr.org/2017/410>
17. Shao, J., Cao, Z.: Cca-secure proxy re-encryption without pairings. In: Jarecki, S., Tsudik, G. (eds.) Public Key Cryptography - PKC 2009, 12th International Conference on Practice and Theory in Public Key Cryptography, Irvine, CA, USA, March 18-20, 2009. Proceedings. Lecture Notes in Computer Science, vol. 5443, pp. 357–376. Springer (2009), [https://doi.org/10.1007/978-3-642-00468-1\\_20](https://doi.org/10.1007/978-3-642-00468-1_20)
18. Zheng, Y.: Digital signcryption or how to achieve cost(signature & encryption)  $\ll$  cost(signature) + cost(encryption). In: Jr., B.S.K. (ed.) Advances in Cryptology - CRYPTO '97, 17th Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 1997, Proceedings. Lecture Notes in Computer Science, vol. 1294, pp. 165–179. Springer (1997), <https://doi.org/10.1007/BFb0052234>

## A Common definitions for confidentiality in PRE

We note that since [6], the main notion for security in PRE schemes is against chosen-ciphertext attacks (CCA) as opposed to chosen-plaintext (CPA) attacks. However since the focus of this paper is to revisit definitions with respect to unauthorised re-encryptions, for simplicity we restrict security to CPA and leave IND-CCA to future work. We further note that since recent practical schemes which are both unidirectional and multi-hop also only focus on CPA security [16] we do not consider this to be a significant weakening of security in comparison with existing practical schemes.

The following definition is a formalism of the preservation of indistinguishability introduced in [6] adapted to CPA security. We note that this definition does not consider compromised keys.

**Definition 9.** A PRE scheme  $\mathcal{PRE}$  preserves IND-CPA if for all PPT algorithms  $\mathcal{A}$  there exists a negligible function  $\text{negl}(\lambda)$  such that:

$$\Pr \left[ \text{pres-IND-CPA}_{\mathcal{PRE}}^A(\lambda) = 1 \right] \leq \frac{1}{2} + \text{negl}(\lambda),$$

where pres-IND-CPA is given in Figure 5.

$\text{Setup}(1^\lambda) \rightarrow \text{param}$	$\text{KeyGen}(1^\lambda) \xrightarrow{\$} (X, x)$	$\text{Enc}(X, m) \xrightarrow{\$} C$	$\text{Dec}(x, C) \rightarrow m$
$p$ large prime	$x \xleftarrow{\$} \mathbb{Z}_p^*$	$y \xleftarrow{\$} \mathbb{Z}_p^*$	$m = \tilde{C}^{-x} \cdot \bar{C}$
$g$ a generator of $\mathbb{Z}_p$	$X = g^x$	$\tilde{C} = g^y$	<b>return</b> $m$
$\mathcal{M} = \mathbb{Z}_p$	$\text{pk} = X, \text{sk} = x$	$\bar{C} = m \cdot X^y$	
$\mathcal{D} = \mathcal{K} = \mathbb{Z}_p^*$	<b>return</b> $(\text{pk}, \text{sk})$	<b>return</b> $C = (\tilde{C}, \bar{C})$	
<b>return</b> $\text{param} = (p, g, \mathcal{M}, \mathcal{K}, \mathcal{D})$			
	$\text{ReKeyGen}(x_i, X_j, \tilde{C}) \xrightarrow{\$} \Delta_{i,j,C}$	$\text{ReEnc}(\Delta_{i,j,C}, C) \rightarrow C'$	
	$y' \xleftarrow{\$} \mathbb{Z}_p^*$	$C = (\tilde{C}, \bar{C})$	
	$\Delta_{i,j,C}^0 = g^{y'}$	$\tilde{C}' = \Delta_{i,j,C}^0 = g^{y'}$	
	$\Delta_{i,j,C}^1 = X_j^{y'} \cdot \tilde{C}^{-x_i} = g^{x_j y' - x_i y}$	$\bar{C}' = \bar{C} \cdot \Delta_{i,j,C}^1 = m \cdot X_j^{y'}$	
	<b>return</b> $\Delta_{i,j,C} = (\Delta_{i,j,C}^0, \Delta_{i,j,C}^1)$	<b>return</b> $C' = (\tilde{C}', \bar{C}')$	

Figure 6: An ElGamal-based scheme similar to [4] which is best-achievable unidirectional and token robust.

Informally, the PRE scheme is still IND-CPA secure even when the adversary is given access to a re-encryption and token generation oracle. Clearly the underlying PKE scheme must be IND-CPA in order for the PRE scheme to be pres-IND-CPA.

Observe that the above definition applies whether or not the PRE scheme is ciphertext-dependent or unidirectional. It can be easily extended to symmetric PRE by providing the adversary with encryption oracles for both keys, see [10].

## B A Secure PRE scheme in the Malicious model

Recall that for PRE suitable for access control, we require a multi-hop scheme. For PRE in the malicious model we require a scheme which is unidirectional, ciphertext-dependent and token robust. A multi-hop, ciphertext dependent scheme is given in Figure 6. We use our definitions in Definition 6 and Definition 4 to assess the unidirectionality and token robustness.

**Correctness:** Let  $C_i = (g^y, m \cdot g^{x_i y})$  be an encryption of  $m$  under  $g^{x_i}$ . The update token resulting from  $\text{ReEnc}(x_i, g^{x_j}, g^y)$  has the form  $\Delta_{i,j,C} = (g^{y'}, X_j^{y'} \cdot (g^y)^{-x_i}) = (g^{y'}, g^{x_j y' - x_i y})$ . Then re-encryption derives a ciphertext of the form  $C_j = (g^{y'}, m \cdot g^{x_i y} \cdot g^{x_j y' - x_i y}) = (g^{y'}, m \cdot g^{x_j y'})$ .

To achieve best-achievable unidirectionality, we require that a new  $\hat{y}$  is selected uniformly at random. Then to verify the received message  $m'$ , the receiver derives  $\hat{a}$  from  $\sigma$  and confirms that  $a^{x^{h(m')}} = \bar{C}_0$ . If they match then we have both message integrity and COA. Therefore this mechanism provides a means of pairing  $g^y$  with  $m$  and associating this with the identity of the encryptor. The full adapted scheme is given in Figure 7.

### B.1 Security analysis

First we show that this scheme is ReEnc-IND-CPA, then best-achievable unidirectional.

**Theorem 2.** *The scheme described in Figure 6 is ReEnc-IND-CPA under the decisional Diffie-Hellman assumption.*

*Proof.* In this scheme, a re-encrypted ciphertext has the same form as a ciphertext that was encrypted directly under the new key, so the domain of the re-encryption to the new key is equal to the domain of encryption under the new key. In other words, re-encrypted ciphertexts under  $x_j$  are identically distributed to ciphertexts encrypted for the first time under  $x_j$ . Therefore the problem reduces to ElGamal being IND-CPA, so we can assume the scheme is ReEnc-IND-CPA.  $\square$

**Theorem 3.** *The scheme in Figure 6 is best-achievable unidirectional.*

We prove this through two lemmas.

For maximal irreversibility, observe that the update token alone cannot be used to reverse a ciphertext. If an adversary has both the re-encryption token used as well as the first component of the old ciphertext, then reversing the re-encryption is trivial. However, we observe that retaining both values would require the same amount of storage as retaining the original ciphertext, and therefore is maximally irreversible. Any reversal attacks are thus obsolete. We prove this formally in Lemma 1.

**Lemma 1.** *The scheme described in Figure 6 is maximally irreversible under the Computational Diffie-Hellman (CDH) assumption.*

In this proof we assume that  $\bar{\lambda} \in \{0, s, \dots, cs\}$  where  $s$  is the size of components in the ciphertext and  $c$  is the number of components updated during re-encryption, as in Definition 6. We do not consider the advantage of the adversary storing  $g^{x_i y - x_j y'}$  and only part of  $g^y$  (or vice versa). In Appendix D we provide a proof which shows the unidirectionality of the scheme when  $\bar{\lambda} \in \{0, 1, \dots, cs\}$ .

*Proof.* The CDH assumption states that, given  $(g^a, g^b)$ , it is computationally infeasible to compute  $g^{ab}$ . The component size in this scheme is  $s = p$ , as we are working modulo  $p$ .

First we observe that  $\tilde{C}$  and  $\Delta_{i,j,C}^1$  can be considered interchangeably as input to  $st_A$ , as bits retained of  $\Delta_{i,j,C}^1$  can be used with  $\tilde{C}'$  to calculate bits of  $\tilde{C}$ . For simplicity, we only consider  $\Delta_1$  as a potential input to  $st_A$ . There are only two values which can help reverse the re-encryption — the old ciphertext header  $\tilde{C}$  and  $\Delta_{i,j,C}^1$ . For the scheme to be best-achievable unidirectional, we need to show that both values must be retained for a successful reversal attack.

We begin by showing that an adversary who only retains  $\Delta_{i,j,C}^1 = g^{x_i y - x_j y'}$  cannot derive  $\tilde{C} = g^y$ . In addition to  $g^{x_i y - x_j y'}$ , we also assume that the server knows the public keys  $g^{x_i}, g^{x_j}$  and the header of the new file,  $g^{y'}$ . For simpler notation, let  $a = x_i, b = x_j, c = y, d = y'$ . The question becomes: given  $(a, g^b, g^d, g^{bd-ac})$ , find  $g^c$ . First we assume we have an oracle which solves the CDH problem:  $\text{OCDH}(g^a, g^b) \rightarrow g^{ab}$ . In this case, given  $(a, g^b, g^d, g^{bd-ac})$  we call  $\text{OCDH}(g^b, g^d) \rightarrow g^{db}$  and use this to retrieve  $g^{ac} = (g^{bd-ac} \cdot g^{-bd})^{-1}$ . We then calculate  $g^c = (g^{ac})^{1/a}$ .

For the other direction, assume we have an oracle  $\text{Og}^c(a, g^b, g^d, g^{bd-ac}) \rightarrow g^c$ . We need to show that we can use this oracle to find  $g^{bd}$  given  $(g^b, g^d)$ . First we observe that for all  $x \in \mathbb{Z}_p^*$ , there exists a  $c$  such that  $g^x = g^{db-c} \pmod p$ . Therefore, calling  $\text{Og}^c(1, g^b, g^d, g^x)$  for some  $x \in \mathbb{Z}_p^*$  returns  $g^c$  where  $g^x = g^{db-c}$ . We can then retrieve  $g^{bd} = g^x \cdot g^c$ . The problem statement is thus equivalent to the CDH problem. We conclude that an adversary who only retains  $g^{x_i y - x_j y'}$  cannot derive  $g^y$ .

Analogously, an adversary who only retains  $g^y$  cannot calculate  $g^{x_i y - x_j y'}$ . The adversary must thus retain both  $g^y$  and  $g^{x_i y - x_j y'}$ . Hence, the number of components needed to reverse the re-encryption is equal to the number of components transformed in the re-encryption process. We conclude the scheme is best-achievable unidirectional.

It remains to show that the adversary cannot output an equivalent ciphertext  $(g^{\hat{y}}, m \cdot g^{x_i \hat{y}})$  for some  $\hat{y} \in \mathbb{Z}_p^*$  given  $(g^{y'}, m \cdot g^{x_j y'}), x_i, g^{x_j}$ . To do this the adversary would need to be able to calculate  $g^{x_j y'}$ , which breaks the DDH assumption.  $\square$

This proof shows that the attacker needs to store just as much as if they were storing the original ciphertext, so they could have avoided needing to perform the attack by simply retaining the original ciphertext.

**Lemma 2.** *The scheme in Figure 6 has token robustness under the CDH assumption.*

*Proof.* To win the token robustness game, the adversary must output a token which re-encrypts an honestly-generated ciphertext so that it is under a key which it has not been encrypted under before.

Recall that the encryption oracle  $\text{OEnc}(i, m)$  outputs a ciphertext of the form  $(g^y, m \cdot g^{x_i y})$  and the token generation oracle  $\text{OReKeyGen}(i, j, C)$  outputs a re-encryption token of the form  $(g^{y'}, g^{x_j y' - x_i y})$ . Let the challenge ciphertext be denoted  $C_i^A = (g^y, m \cdot g^{x_i y})$ . Then the adversary must output a token of the form  $(g^{y'}, g^{x_j y' - x_i y})$ , where  $j \notin \text{chain}(m)$ . It may be possible that querying two completely different keys and a different ciphertext results in a token  $\Delta$ , where  $\Delta = g^{x_k y_k - x_l y_l} = g^{x_i y - x_j y'}$  for some  $y'$ , but this only occurs with negligible probability.

We note that it is trivial for the adversary to calculate  $g^{x_j y'}$  for some  $y' \xleftarrow{\$} \mathbb{Z}_p^*$  from  $\text{pk}_j$  by setting  $\text{pk}_j^{y'} = g^{x_j y'}$ . It remains for the adversary to calculate  $g^{x_i y}$ . Since the ciphertext is honestly generated ( $\text{generated}(C_i^A) = \text{true}$ ), the adversary does not know  $y$  and so cannot use this to compute  $\delta^A = X_j^{y'} \cdot (\text{pk}_i^y)^{-1} = g^{x_j y' - x_i y}$ . Since factoring is a difficult problem modulo  $p$ ,  $g^{x_i y}$  cannot be easily separated from tokens of the form  $\Delta_{i,k,C_i^A} = g^{x_k y'' - x_i y}$  for some  $x_k \leftarrow \text{KeyGen}(1^\lambda), y'' \in \mathbb{Z}_p^*$ . Finding the most common factor is also a difficult problem modulo  $p$ , so a string of tokens of this form also cannot be used to derive  $g^{x_i y}$ .

Alternatively, as each token output by  $\text{OReKeyGen}$  includes new randomisation  $y'$  which is not revealed to the adversary, these tokens blind the value  $g^{x_i y}$  that the adversary must retrieve to win the game.

We conclude that the adversary can only gain ciphertexts of the correct form for keys  $\text{pk}_i$  where  $i \in \text{chain}(\text{Dec}(\text{sk}_i, C_i^A))$ . Since the adversary must output a token for a new key to win the game, we conclude the scheme is token robust.  $\square$

We have shown that our scheme is suitable for the malicious model according to the goals we outlined in Section 1. This means a malicious server will be unable to perform unauthorised re-encryptions on stored files as much as can be guaranteed given realistic storage assumptions.

An adapted version of this scheme which provides Ciphertext Origin Authentication (COA) is given in Figure 7.

## C Non-optional COA extension for [10]

We give an extension to the ReCrypt scheme given in [10] which includes a compulsory COA check. A description of ReCrypt is given in Figure 8.

As with our scheme, we attach the identity of the encryptor / re-encryption initiator with the randomness used in the encryption. However, since ReCrypt is not an ElGamal-based scheme, the adaptation for a mandatory COA check is more complicated.

Essentially, we replace  $x, y \xleftarrow{\$} \mathcal{K}$  and  $\chi = x + y$  in the ciphertext header by having  $x, y \xleftarrow{\$} \{i, j, : i + j = g^r\}$  for  $r \xleftarrow{\$} \mathbb{Z}_p^*$  and setting  $\chi = g^r$ . Then by signing  $r$ ,  $\text{VerRetrieve}$  will return  $g^r = \chi$  and decryption is adjusted accordingly. Note that we move  $y$  from the ciphertext body into the ciphertext header, as it is now needed in generating update tokens. The full extension is given in Figure 9.

$\text{Setup}(1^\lambda) \rightarrow param$ $p$ large prime $g$ a generator of $\mathbb{Z}_p$ <b>return</b> $(p, g)$	$\text{encKeyGen}(1^\lambda) \xrightarrow{\S} (X, x)$ $x \xleftarrow{\S} \mathbb{Z}_p^*$ $X = g^x$ $pk = X, sk = x$ <b>return</b> $(X, x)$	$\text{sigKeyGen}(1^\lambda) \xrightarrow{\S} (\mathbf{X}, \mathbf{x})$ $\mathbf{x} \xleftarrow{\S} \mathbb{Z}_p^*$ $\mathbf{X} = g^{\mathbf{x}}$ $ssk = \mathbf{x}, svk = \mathbf{X}$ <b>return</b> $(\mathbf{X}, \mathbf{x})$	
$\text{Sign}(\mathbf{x}, y) \xrightarrow{\S} \sigma$ $k \xleftarrow{\S} \mathbb{Z}_p, r = g^k$ $s = (y - \mathbf{x}r)k^{-1} \pmod{p-1}$ $\sigma = (r, s)$ <b>return</b> $\sigma$	$\text{VerRetrieve}(\mathbf{X}, \sigma) \rightarrow a$ $(r, s) = \sigma$ $a = \mathbf{X}^r \cdot r^s$ <b>return</b> $a$	$\text{Enc}(\mathbf{x}, X, m) \xrightarrow{\S} C$ $y \xleftarrow{\S} \mathbb{Z}_p^*, \sigma \leftarrow \text{Sign}(\mathbf{x}, y)$ $\hat{y} \xleftarrow{\S} \mathbb{Z}_p^*, \hat{\sigma} \leftarrow \text{Sign}(\mathbf{x}, \hat{y})$ $\tilde{C} = (\sigma, \hat{\sigma})$ $\tilde{C} = (X^{\hat{y}h(m)}, m \cdot X^y)$ <b>return</b> $C = (\tilde{C}, \tilde{C})$	$\text{Dec}(\mathbf{X}, x, C) \rightarrow m$ $(\sigma, \hat{\sigma}) = \tilde{C}$ $a \leftarrow \text{VerRetrieve}(\mathbf{X}_A, \sigma)$ $\hat{a} \leftarrow \text{VerRetrieve}(\mathbf{X}_A, \hat{\sigma})$ $m' = a^{-x} \cdot \tilde{C}_1$ <b>if</b> $\hat{a}^{xh(m')} \neq \tilde{C}_0$ : <b>return</b> $\perp$ <b>else return</b> $m'$
$\text{ReKeyGen}(\mathbf{X}_A, \mathbf{x}_B, x_i, X_j, \tilde{C}) \xrightarrow{\S} \Delta_{i,j,C}$ $(\sigma, \hat{\sigma}) = \tilde{C}$ $a \leftarrow \text{VerRetrieve}(\mathbf{X}_A, \sigma)$ $\hat{a} \leftarrow \text{VerRetrieve}(\mathbf{X}_A, \hat{\sigma})$ $y' \xleftarrow{\S} \mathbb{Z}_p^*, \sigma' \leftarrow \text{Sign}(\mathbf{x}_B, y')$ $\hat{y}' \xleftarrow{\S} \mathbb{Z}_p^*, \hat{\sigma}' \leftarrow \text{Sign}(\mathbf{x}_B, \hat{y}')$ $\tilde{C}' = (\sigma', \hat{\sigma}')$ $\Delta_{i,j,C} = (\tilde{C}', X_j^{\hat{y}'} \cdot \hat{a}^{-x_i}, X_j^{y'} \cdot a^{-x_i})$ <b>return</b> $\Delta_{i,j,C}$		$\text{ReEnc}(\Delta_{i,j,C}, C) \rightarrow C$ $(\tilde{C}', \Delta_1, \Delta_2) = \Delta_{i,j,C}$ $\tilde{C}' = (\tilde{C}_0 \cdot \Delta_1, \tilde{C}_1 \cdot \Delta_2)$ <b>return</b> $(\tilde{C}', \tilde{C}')$	

Figure 7: Adapted scheme with ciphertext origin authentication and message integrity.

## D More fine-grained choices for $\bar{\lambda}$

Recall that in our definition of unidirectionality Definition 6, the security parameter  $\bar{\lambda}$  is determined in terms of components  $\bar{\lambda} \in \{0, s, \dots, cs\}$  where  $s$  is the component size and  $c$  is the number of ciphertext components updated during re-encryption. A more fine-grained definition would be to allow  $\bar{\lambda}$  to be any number of bits  $\bar{\lambda} \in \{0, 1, \dots, cs\}$ . We decided against this definition as whilst it be more fine-grained in terms of security, it makes proofs much harder for what we consider to not be a significant distinction in practice. The main difference in terms of proofs is that the more fine-grained options for  $\bar{\lambda}$  require additional proofs that an adversary who only retains part of a component or component has no real advantage in calculating the rest.

To demonstrate this as well as for general interest we include a proof here that shows that our scheme in Figure 6 is also best-achievable unidirectional in this stricter sense, as long as the prime  $p$  chosen by the Setup algorithm is a Mersenne prime.

To prove that our scheme is best-achievable unidirectional, we need the following lemma, after which best-achievable unidirectionality follows from a trivial adaptation of Lemma 1.

Setup( $1^\lambda$ )	KeyGen( $1^\lambda$ )	Enc( $k, m$ )	Dec( $k, (\tilde{C}, \bar{C})$ )
<p><math>p</math> large prime  <math>g</math> a generator of <math>\mathbb{Z}_p</math>  <b>return</b> <math>p, g</math></p>	<p><math>k \xleftarrow{\\$} \mathcal{KG}(1^\lambda)</math></p>	<p><math>x, y \xleftarrow{\\$} \mathcal{K}</math>  <math>\chi = x + y</math>  <math>\tau = h(m) + F(x, 0)</math>  <math>\tilde{C} = \mathcal{E}_k(\chi, \tau)</math>  <math>\bar{C}_0 = y</math>  <b>for</b> <math>0 \leq i \leq n</math> :  <math>\bar{C}_i = m_i + F(x, i)</math>  <b>return</b> <math>(\tilde{C}, \bar{C})</math></p>	<p><math>(\chi, \tau) \leftarrow \mathcal{D}_k(\tilde{C})</math>  <b>if</b> <math>(\chi, \tau) = \perp</math> <b>return</b> <math>\perp</math>  <math>y = \bar{C}_0</math>  <b>for</b> <math>0 \leq i \leq n</math> :  <math>m_i = \bar{C}_i - F(\chi - y, i)</math>  <b>if</b> <math>h(m) + F(\chi - y, 0) \neq \tau</math>  <b>return</b> <math>\perp</math>  <b>else return</b> <math>m</math></p>
	<p>ReKeyGen(<math>k_i, k_j, \tilde{C}</math>)</p>	<p>ReEnc(<math>\Delta_{i,j,C}, C</math>)</p>	
	<p><math>(\chi, \tau) \leftarrow \mathcal{D}_k(\tilde{C}_0)</math>  <b>if</b> <math>(\chi, \tau) = \perp</math> <b>return</b> <math>\perp</math>  <math>x', y' \xleftarrow{\\$} \mathcal{K}</math>  <math>\chi' = \chi + x' + y'</math>  <math>\tau' = \tau + F(x', 0)</math>  <math>\tilde{C}' = \mathcal{E}_{k_j}(\chi', \tau')</math>  <b>return</b> <math>\Delta_{i,j,C} = (\tilde{C}', x', y')</math></p>	<p><math>(\tilde{C}', x', y') \leftarrow \Delta_{i,j,C}</math>  <math>y = \bar{C}_0</math>  <math>\bar{C}'_0 = y' + y</math>  <b>for</b> <math>0 \leq i \leq n</math> :  <math>\bar{C}'_i = \bar{C}_i + F(x', i)</math>  <b>return</b> <math>(\tilde{C}', \bar{C}')</math></p>	

Figure 8: ReCrypt [10]

**Lemma 3.** Let  $p$  be a Mersenne prime of length  $n$ , and let  $a, b \xleftarrow{\$} \mathbb{Z}_p^*$ . Then, for all PPT algorithms  $\mathcal{B}$  there exists a negligible function  $\text{negl}(\lambda)$  such that:

$$\Pr[\mathcal{B}(g^a, g^b, [g^{ab}]_{\bar{\lambda}}) \rightarrow g^{ab}] \leq \frac{1}{2^{n-\bar{\lambda}}} + \text{negl}(\lambda) \quad (8)$$

where  $[g^{ab}]_{\bar{\lambda}}$  denotes the  $\bar{\lambda}$  known bits of  $g^{ab}$ .

*Proof.* Since  $p$  is a Mersenne prime it has the form  $2^n - 1$ , and so is  $n$  bits long and there is a bijection between integers in  $\mathbb{Z}_p$  and bitstrings of length  $n$ . In other words every integer in  $\mathbb{Z}_p$  can be represented as a bitstring. Therefore random elements of  $\mathbb{Z}_p$  of the form can be modelled as random  $n$ -bit strings. More specifically, elements  $g^c$  where  $c$  is chosen uniformly at random can be considered as random  $n$ -bit strings.

We proceed by induction. We note that an element of this  $\mathbb{Z}_p$  has size  $n$ .

**Base case:**  $\bar{\lambda} = 0$ . By the Computational Diffie-Hellman assumption (CDH), the adversary cannot compute  $g^{ab}$  with probability significantly greater than  $\frac{1}{2^n}$ .

**Assume for**  $\bar{\lambda} = i$ . An adversary who knows the first  $i$  bits of  $g^{ab}$  cannot calculate the remaining  $n - i$  bits with probability higher than  $\frac{1}{2^{n-i}}$ .

We now consider a variant of the DDH game which factors in the adversary's knowledge of  $i$  bits of  $g^{ab}$ . We denote by  $[g^{ab}]_{\bar{\lambda}}$  the  $\bar{\lambda}$  known bits of  $g^{ab}$ . Let  $a, b$  be selected uniformly at random and let  $c$  be selected at random with the condition that  $[g^c]_i = [g^{ab}]_i$ :  $a, b \xleftarrow{\$} \mathbb{Z}_p^*, c \xleftarrow{\$} \{x \in \mathbb{Z}_p^* : [g^x]_i = [g^{ab}]_i\}$ . A consequence of assuming our hypothesis is correct for  $i$  is that an adversary cannot distinguish between  $(g^a, g^b, g^{ab})$  and  $(g^a, g^b, g^c)$  with non-negligible probability. If our assumption were false, the adversary would have an advantage in this game with probability  $\frac{1}{2^i}$ . We will use this result moving forward.

**Show for**  $\bar{\lambda} = i + 1$ . For  $a, b \xleftarrow{\$} \mathbb{Z}_p^*, c \xleftarrow{\$} \{x \in \mathbb{Z}_p^* : [g^x]_i = [g^{ab}]_i\}$ , we have that  $[g^c]_{i+1} \neq [g^{ab}]_{i+1}$

Setup( $1^\lambda$ )	encKeyGen( $1^\lambda$ )	sigKeyGen( $1^\lambda$ )	Enc( $k, \mathbf{x}, m$ )
<p><math>p</math> large prime  <math>g</math> a generator of <math>\mathbb{Z}_p</math>  <b>return</b> <math>p, g</math></p>	<p><math>k \xleftarrow{\\$} \mathcal{KG}(1^\lambda)</math></p>	<p><math>\mathbf{x} \xleftarrow{\\$} \mathbb{Z}_p^*</math>  <math>\mathbf{X} = g^{\mathbf{x}}</math>  <b>return</b> <math>(\mathbf{x}, \mathbf{X})</math></p>	<p><math>r \xleftarrow{\\$} \mathbb{Z}_p^*</math>  <math>x, y \xleftarrow{\\$} \{i, j, : i + j = g^r\}</math>  <math>\sigma \leftarrow \text{Sign}(\mathbf{x}, r)</math>  <math>\tau = h(m) + F(x, 0)</math>  <math>\tilde{C} = (\mathcal{E}_k(\sigma, \tau), y)</math>  <b>for</b> <math>0 \leq i \leq n</math> :  <math>\quad \tilde{C}_i = m_i + F(x, i + 1)</math>  <b>return</b> <math>(\tilde{C}, \tilde{C})</math></p>
Dec( $x, \mathbf{X}, (\tilde{C}, \tilde{C})$ )	ReKeyGen( $\mathbf{X}_A, \mathbf{x}_B, k_i, k_j, \tilde{C}$ )	ReEnc( $\Delta_{i,j,C}, C$ )	
<p><math>(\sigma, \tau) \leftarrow \mathcal{D}_k(\tilde{C}_0)</math>  <b>if</b> <math>(\sigma, \tau) = \perp</math> <b>return</b> <math>\perp</math>  <math>\chi \leftarrow \text{VerRetrieve}(\mathbf{X}, \sigma)</math>  <b>for</b> <math>0 \leq i \leq n</math> :  <math>\quad m_i = \tilde{C}_i - F(\chi - y, i + 1)</math>  <b>if</b> <math>h(m) + F(\chi - y, 0) \neq \tau</math>  <b>return</b> <math>\perp</math>  <b>else return</b> <math>m</math></p>	<p><math>(\sigma, \tau) \leftarrow \mathcal{D}_{k_i}(\tilde{C}_0)</math>  <b>if</b> <math>(\sigma, \tau) = \perp</math> <b>return</b> <math>\perp</math>  <math>\chi \leftarrow \text{VerRetrieve}(\mathbf{X}_A, \sigma)</math>  <math>y = \tilde{C}_1, x = \chi - y</math>  <math>r' \xleftarrow{\\$} \mathbb{Z}_p^*</math>  <math>x', y' \xleftarrow{\\$} \{i, j, : i + j = g^{r'}\}</math>  <math>\sigma' \leftarrow \text{Sign}(\mathbf{x}_B, r')</math>  <math>a = x' - x</math>  <math>\tau' = \tau + F(a, 0)</math>  <math>\tilde{C}' = (\mathcal{E}_{k_j}(\sigma', \tau'), y')</math>  <b>return</b> <math>\Delta_{i,j,C} = (\tilde{C}', a)</math></p>	<p><math>(\tilde{C}', a) \leftarrow \Delta_{i,j,C}</math>  <b>for</b> <math>0 \leq i \leq n</math> :  <math>\quad \tilde{C}'_i = \tilde{C}_i + F(a, i + 1)</math>  <b>return</b> <math>(\tilde{C}', \tilde{C}')</math></p>	

Figure 9: ReCrypt [10] with compulsory COA

with probability  $\frac{1}{2}$ . In other words, the bit which the adversary is trying to predict is going to be different from  $g^{ab}$  in  $g^c$  with probability  $\frac{1}{2}$ . Suppose for a contradiction that the adversary has a non-negligible advantage in distinguishing  $(g^a, g^b, g^{ab})$  from  $(g^a, g^b, g^c)$ . This would mean an adversary has an advantage in winning the variant of DDH in case  $i$  half the time, whenever  $[g^c]_{i+1} \neq [g^{ab}]_{i+1}$ . This makes the assumption in case  $i$  false. Equivalently, the adversary would have an advantage in winning DDH with probability  $\frac{1}{2^i}$ . Therefore by contradiction, an adversary has no significant advantage calculating another bit of  $g^{ab}$  over guessing. Overall, the adversary cannot compute the remaining bits  $g^{ab}$  with probability significantly greater than  $\frac{1}{2^{n-i-1}}$ .

We conclude that knowledge of the first  $\lambda$  bits of  $g^{ab}$  in addition to  $(g^a, g^b)$  does not help the adversary compute the remaining bits of  $g^{ab}$  with probability higher than  $\frac{1}{2^{n-\lambda}}$  when  $p$  is a Mersenne prime.