

Chosen-Ciphertext Secure Proxy Re-Encryption

Ran Canetti* Susan Hohenberger†

October 23, 2007

Abstract

In a proxy re-encryption (PRE) scheme, a proxy is given special information that allows it to translate a ciphertext under one key into a ciphertext of the same message under a different key. The proxy cannot, however, learn anything about the messages encrypted under either key. PRE schemes have many practical applications, including distributed storage, email, and DRM. Previously proposed re-encryption schemes achieved only semantic security; in contrast, applications often require security against chosen ciphertext attacks. We propose a definition of security against chosen ciphertext attacks for PRE schemes, and present a scheme that satisfies the definition. Our construction is efficient and based only on the Decisional Bilinear Diffie-Hellman assumption in the standard model. We also formally capture CCA security for PRE schemes via both a game-based definition and simulation-based definitions that guarantee universally composable security. We note that, simultaneously with our work, Green and Ateniese proposed a CCA-secure PRE, discussed herein.

1 Introduction

Encryption is one of the most fundamental cryptographic functions, and yet its practical adoption is often hampered by key management problems. Suppose pk_1 and pk_2 are two independently chosen encryption keys. As pointed out by Mambo and Okamoto [23], it is a common situation in practice that data is encrypted under pk_1 and an application requires that it be encrypted under pk_2 . When the holder of sk_1 is online, this translation is easy: $E_1(m)$ is decrypted using sk_1 to obtain m , then m is encrypted under pk_2 , resulting in $E_2(m)$. Yet in many applications—such as encrypted email forwarding [6, 5], distributed file systems [1, 2], or the DRM of Apple’s iTunes [26]—this translation is being performed by an untrusted party! As a demonstration of the email forwarding scenario, imagine that Alice is going on vacation and wishes to have her mail server forward all of her encrypted email to Bob, *without giving her secret key to either the mail server or Bob*.

In 1998, Blaze, Bleumer and Strauss (BBS) proposed a solution to this widely-encountered key management problem [6, 5]. They introduced *proxy re-encryption*, where a (potentially untrusted) proxy is given a re-encryption key $rk_{1,2}$ that allows it to translate a message m encrypted under public key pk_1 into an encryption of the same message m under a different public key pk_2 – without

*IBM T.J. Watson Research Center, Hawthorne, NY, USA, canetti@watson.ibm.com. Supported by NSF grants CT-0430450 and CFF-0635297, and a US-Israel Binational Science Foundation Grant.

†Johns Hopkins University, Baltimore, MD, USA, susan@cs.jhu.edu. Research partially performed at IBM Zurich Research Laboratory and supported by NSF grant CT-0716142.

being able to learn anything about the encrypted messages! I.e., the CPA-security of encryptions under pk_1 and pk_2 hold even against an adversary possessing $rk_{1,2}$.

BBS categorize two types of re-encryption schemes. If the re-encryption key $rk_{1,2}$ necessarily allows the proxy to translate ciphertexts under pk_1 into ciphertexts under pk_2 and *vice versa*, then the scheme is called *bidirectional*. If the re-encryption key $rk_{1,2}$ allows the proxy to translate only from pk_1 to pk_2 , then the scheme is called *unidirectional*. (Any unidirectional scheme can be easily transformed to a bidirectional one, setting the bidirectional key $rk_{1,2}$ to be the two unidirectional keys between pk_1 and pk_2 . We do not know if the converse holds.) BBS propose the first bidirectional CPA-secure scheme, leaving the construction of a unidirectional scheme as an open problem. In 2005, Ateniese, Fu, Green, and Hohenberger [1, 2] presented a unidirectional CPA-secure scheme. Capturing this useful property comes at a price: the re-encryption algorithm is *single-hop*; that is, a re-encrypted ciphertext cannot be further re-encrypted. In contrast, the BBS scheme is *multi-hop*, namely a ciphertext can be re-encrypted from Alice to Bob to Carol and so on. Constructing a multi-hop unidirectional PRE scheme remains open. Both of these PRE algorithms are only CPA-secure, and CPA security is often not sufficient to guarantee security in general protocol settings.

1.1 Our Contributions

We address the problem of obtaining PRE schemes that are secure in arbitrary protocol settings, or in other words are secure against *chosen ciphertext attacks*. The concept of a CCA secure PRE scheme sounds almost self-contradictory, since on the one hand we want the ciphertexts to be non-malleable, and on the other hand we want to allow the proxy to “translate” the ciphertext from one public key to another. Still, we formulate a meaningful definition of CCA-secure PRE schemes, along with a construction that meets the definition in the standard model and under relatively mild hardness assumptions for bilinear groups.

DEFINITIONS. We formalize definitions of security against CCA attacks for re-encryption schemes. In fact, we provide three definitions: a game-based definition and two simulation-based definitions that guarantee universally composable security. We then demonstrate that the game-based definition lies in between the two simulation-based definitions.

The game-based definition gives the adversary access to the standard encryption and decryption oracles, in addition to a re-encryption oracle (which translates ciphertexts) and a re-encryption key oracle (which returns re-encryption keys). For the definition to make sense, one needs to define the game so as to disallow decryption queries not only on the challenge ciphertext (as usual), but also on any trivial derivative of the challenge ciphertext (e.g., derivatives obtained from a re-encryption).

The simulation-based definitions are formulated within the universally composable (UC) framework [9]. They extend the ideal functionality for CCA-secure encryption [9] to the case of re-encryption, using ideas from the UC definition of replayable CCA (RCCA) security [13]. (Indeed, PRE security is somewhat reminiscent of RCCA security, in that both notions aim at guaranteeing illegitimate “mauling” of ciphertexts, while permitting legitimate “re-encrypting” without modifying the hidden message.) The need to provide two separate UC formulations (one that implies the game-based definition and one that follows from it) stems from the difficulty in capturing the adaptive nature of re-encryption queries. See more details within.

CONSTRUCTIONS. We present a re-encryption scheme secure against chosen-ciphertext attacks. Our scheme is bidirectional and multi-hop. Moreover, our scheme is efficient enough to be used

in practice. We prove it secure under the Decisional Bilinear-Diffie Hellman assumption. For clarity, we present two constructions. First we present and analyze a scheme in the random oracle model. This scheme is simpler and provides intuition into the techniques used. Next we modify the scheme to work in the standard model. The modified scheme and its analysis are somewhat more complicated, but only slightly less efficient.

As in other re-encryption schemes [6, 5, 1, 2, 21, 18], we work in a static adversary model and require that users publicize public keys only if they have the corresponding secret keys. To make this second requirement rigorous, we can either work in the trusted key generation model (where a trusted party generates and distributes all key-pairs), or alternatively in a model where the parties have to demonstrate knowledge of their secret keys when registering their public keys (such as the key registration model of Barak et al. [3]).

We prove that the schemes satisfy the game-based definition. The proof turns out to be non-trivial; in particular, in some cases, one must correctly answer re-encryption queries *without* knowing the corresponding re-encryption key. In fact, the technique we use in the analysis of the scheme is new and may be useful elsewhere.

We leave open many interesting problems in this area, such as designing: (1) unidirectional CCA-secure schemes (without random oracles), (2) any construction that is simultaneously unidirectional and multi-hop, (3) any unidirectional or CCA-secure scheme without bilinear groups, (4) schemes that satisfy the notion of secure obfuscation for CCA-secure re-encryption or other key translation schemes, (5) a CCA-secure PRE scheme that provides unlinkability for re-encrypted ciphertexts, and (6) a definition of security that is equivalent to our game-based definition and also guarantees secure compossibility properties.

PRACTICAL IMPACT. In 2005, the digital rights management (DRM) of Apple’s iTunes was compromised partially due to the fact that an untrusted party (i.e., the client’s resource) could obtain the plaintext during a naive decrypt-and-encrypt operation, albeit with symmetric encryption [26]. This flaw could be prevented by using a secure PRE scheme.

Ateniese et al. [1, 2] built a secure distributed file system, using unidirectional PRE, where a server can re-encrypt and send files to many different clients without ever gaining access to the plaintext.

However, plain CPA-security is clearly not enough for some applications, such as encrypted email forwarding. For instance, an adversary might hope to gain access to a “decryption oracle” by mangling ciphertexts, emailing them to Alice, and then hoping that she responds with, “Did you send the following to me? [Decrypted attachment.]” The present work fills this gap.

THEORETICAL IMPACT. Recently, the notion of re-encryption was considered in the more general framework of program obfuscation. Hohenberger, Rothblum, Shelat, and Vaikuntanathan [21] presented a unidirectional, CPA-secure re-encryption scheme, where the re-encryption key allows the proxy to learn *nothing* more than black-box access to such a program. This positive result for obfuscating an encryption functionality stands in contrast to a series of impossibility results [4] for general obfuscation and negative improbability [20, 16] results for obfuscation of many cryptographic functionalities. Here we relax this obfuscation notion and only require that the proxy learn nothing about the encrypted messages (but the proxy might learn other things not exposed by black-box access.) In particular, our notion of security allows for a deterministic re-encryption procedure, whereas the obfuscation notion of [21] mandates that the re-encryption process be randomized.

1.2 Intuition Behind the Construction

The idea behind our construction begins with the Canetti, Halevi, and Katz [11] paradigm for transforming any selective-identity, CPA-secure identity-based encryption (IBE) scheme into a CCA-secure cryptosystem. Recall the CHK transformation. Let $(\mathcal{G}, \mathcal{S}, \mathcal{V})$ be a strongly-unforgeable one-time signature scheme. Let (G, E, D) be a semantically-secure IBE scheme. Then a CCA-secure cryptosystem can be created as follows. A user’s public key corresponds to the master public key output by G . To encrypt a message m , an encryptor first runs \mathcal{G} to obtain a signature keypair (svk, ssk) . The encryptor then generates the ciphertext $c = E(svk, m)$, using svk as the identity, and signs this ciphertext as $s = \mathcal{S}(ssk, c)$. The output of the encryption algorithm is the tuple (svk, c, s) . To decrypt, a user first checks that $\mathcal{V}(svk, c, s)$ verifies, and if so, proceeds to decrypt c using the master secret key.

Now, there are many similarities between the Boneh and Franklin IBE [8] and the Ateniese et al. PRE [1]. Can we simply apply the CHK paradigm to achieve CCA security? Unfortunately, applying the CHK paradigm to the Boneh-Franklin IBE looks unwieldy at first. These IBE ciphertexts have two parts: let $c = (X, Y)$. If the encryptor signs (X, Y) in the CHK transformation, then the proxy can’t re-encrypt (X, Y) without invalidating the signature. But if the encryptor only signs, say, Y , then the adversary can arbitrarily mutate X , thus changing the decryption value. Our solution is to add an element Z to the ciphertext, such that (Y, Z) will be signed and Z allows anyone to check that the unsigned value X wasn’t mutated in any meaningful way.

The primary technical challenge is to prove this scheme secure via a reduction that can successfully answer the numerous oracle queries allowed to the re-encryption adversary. In Section 3.2, we first show how to navigate these queries in the random oracle model, which simplifies the analysis. Section 3.3 then shows how to replace the random oracles with *specific* concrete hash functions, which will be carefully manipulated to (among other things) sometimes allow the reduction to decrypt or re-encrypt a ciphertext without knowing the correct keys. Our final scheme is only slightly less efficient than its random oracle counterpart.

1.3 Related Work

There is a large related body of work with the very similar name *proxy encryption* (no “re-”) [22, 27, 14]. In proxy encryption, Alice allows Bob to decrypt ciphertexts meant for her with the help of a proxy. In these schemes, Alice’s secret key is shared between Bob and the proxy. An encryption for Alice is first partially decrypted by the proxy, and then fully decrypted by Bob. However, this approach requires that Bob obtain and store an additional secret (for every decryption delegation he accepts!). Proxy encryption schemes are currently realized under a broader class of complexity assumptions than PREs. For example, Dodis and Ivan present both CPA and CCA secure constructions based on RSA, Decisional DH, and bilinear assumptions [14]. PREs are a (strict) subset of proxy encryption schemes [1], where Bob need only have his own decryption key.

In a concurrent and independent work, Green and Ateniese presented the first CPA and CCA-secure *identity-based* PREs in the random oracle model [18]. They employ the above key sharing technique, but alter it as follows: Alice gives one share of her key to the proxy in the clear and the other share to the proxy encrypted under Bob’s public key. To perform re-encryption, the proxy first partially decrypts the ciphertext and then attaches the ciphertext containing the secret share so that Bob can decrypt the rest. In the CCA case, this requires carefully ensuring that the composition of these ciphertexts results in a non-malleable joint-ciphertext. Their CCA IBE scheme

is unidirectional and single-hop in the random oracle model, whereas ours is bidirectional and multi-hop in the standard model. Depending on the application, one may need one or the other. One additional difference between the schemes is that our original and re-encrypted ciphertexts come from the same distribution and thus when Bob receives a re-encrypted ciphertext he cannot tell (based on the ciphertext) whether it is an original or re-encrypted one. In [18], original and re-encrypted ciphertexts come from two separate distributions and moreover Bob *must* know who the ciphertext was originally sent to in order to decrypt a re-encryption. Green and Ateniese [18] also present a game-based unidirectional CCA IBE security definition, which shares many properties with ours. They do not consider simulation-based definitions.

This work should not be confused with the “universal re-encryption” literature [17], which *re-randomizes* ciphertexts instead of changing the underlying public key.

2 Definitions

Throughout this section we concentrate on defining *bidirectional* re-encryption schemes. The case of *unidirectional* schemes, namely where a re-encryption key from pk_1 to pk_2 should *not* provide the ability to re-encrypt from pk_2 to pk_1 can be inferred in a straightforward way. We allow re-encryption proxies between uncorrupted parties to be corrupted adaptively during the course of the computation. In contrast, we restrict the adversary to determine the identities of the corrupted key-holders before the computation starts. Furthermore, do not allow adaptive corruption of proxies between corrupted and uncorrupted parties.¹

We first give the input-output specifications for a proxy re-encryption scheme (PRE). Section 2.1 gives a game-based security definition. Section 2.2 gives a security definition within the Universally Composable security framework. Section 2.3 discusses the relationship between the definitions.

Definition 2.1 (Bidirectional PRE I/O) *A bidirectional, proxy re-encryption scheme is a tuple of algorithms (KeyGen, ReKeyGen, Enc, ReEnc, Dec):*

- $\text{KeyGen}(1^k) \rightarrow (pk, sk)$. On input the security parameter 1^k , the key generation algorithm KeyGen outputs a public key pk and a secret key sk .
- $\text{ReKeyGen}(sk_1, sk_2) \rightarrow rk_{1 \leftrightarrow 2}$. On input two secret keys sk_1 and sk_2 , the re-encryption key generation algorithm ReKeyGen outputs a bidirectional re-encryption key $rk_{1 \leftrightarrow 2}$.
- $\text{Enc}(pk, m) \rightarrow C$. On input a public key pk and a message $m \in \{0, 1\}^*$, the encryption algorithm Enc outputs a ciphertext C .
- $\text{ReEnc}(rk_{1 \leftrightarrow 2}, C_1) \rightarrow C_2$. On input a re-encryption key $rk_{1 \leftrightarrow 2}$ and a ciphertext C_1 , the re-encryption algorithm ReEnc outputs a second ciphertext C_2 or the error symbol \perp .
- $\text{Dec}(sk, C) \rightarrow m$. On input a secret key sk and a ciphertext C , the decryption algorithm Dec outputs a message $m \in \{0, 1\}^*$ or the error symbol \perp .

Remark 2.2 (On generating re-encryption keys:) *Here ReKeyGen is treated as an algorithm that takes for input the two relevant secret keys and generates a re-encryption key. Alternatively, a re-encryption key can be generated via a protocol involving the proxy and the two holders of the secret keys. The requirements are that the proxy learns the re-encryption key as defined here and nothing else. The holders of the secret keys learn nothing from the protocol. These security*

¹This limitation reflects the current state of the art: Unfortunately, we do not have notions of security that adequately capture adaptive corruption of principals, and are at the same time realizable by realistic protocols.

requirements must hold only when at most one of the three parties is corrupted. When any two of the parties are corrupted, we no longer have security requirements.

The input/output specification for a *unidirectional* PRE scheme would change as follows: instead of taking two secret keys as input, the ReKeyGen algorithm would take as input a secret key sk_1 and a public key pk_2 , and output a re-encryption key $rk_{1 \rightarrow 2}$. Here the protocol for providing a proxy with a re-encryption key from pk_1 to pk_2 is simple: The owner of sk_1 locally computes $rk_{1 \rightarrow 2}$ and sends it to the proxy. The owner of sk_2 need not be part of this process.

2.1 Game-Based Definition of Security

We start by formulating the requirements for correctness of decryption.

Definition 2.3 (Bidirectional PRE Correctness) *A bidirectional PRE scheme (KeyGen, ReKeyGen, Enc, ReEnc, Dec) is perfectly correct with respect to domain D if:*

- For all (pk, sk) output by KeyGen and all $m \in D$, it holds that $\text{Dec}(sk, \text{Enc}(pk, m)) = m$;
- For any $n > 1$, any sequence of pairs $(pk_1, sk_1), \dots, (pk_n, sk_n)$ output by KeyGen, any $i < n$, all re-encryption keys $rk_{i \leftrightarrow i+1}$ output by $\text{ReKeyGen}(sk_i, sk_{i+1})$, all messages $m \in D$ and all C_1 output by $\text{Enc}(pk_1, m)$, it holds that $\text{Dec}(sk_n, \text{ReEnc}(rk_{n-1 \leftrightarrow n}, \dots, \text{ReEnc}(rk_{1 \leftrightarrow 2}, C_1) \dots)) = m$. If for any $m \in D$ correctness holds only with probability 1 minus a negligible quantity, we say that the scheme is correct with respect to D .

Next we define the game used for formulating the security requirement. The game defines an interaction between an adversary and a number of oracles, representing the capabilities of the adversary in an interaction with a PRE scheme. It proceeds as follows:

Definition 2.4 (Bidirectional PRE-CCA game)

Let k be the security parameter. Let \mathcal{A} be an oracle TM, representing the adversary. The game consists of an execution of \mathcal{A} with the following oracles, which can be invoked multiple times in any order, subject to the constraints below:

Uncorrupted key generation: Obtain a new key pair as $(pk, sk) \leftarrow \text{KeyGen}(1^k)$. \mathcal{A} is given pk .

Corrupted key generation: Obtain a new key pair as $(pk, sk) \leftarrow \text{KeyGen}(1^k)$. \mathcal{A} is given pk, sk .²

Re-encryption key generation $\mathcal{O}_{\text{rkey}}$: On input (pk, pk') by the adversary, where pk, pk' were generated before by KeyGen, return the re-encryption key $rk_{pk \leftrightarrow pk'} = \text{ReKeyGen}(sk, sk')$ where sk, sk' are the secret keys that correspond to pk, pk' .

We require that either both pk and pk' are corrupted, or alternatively both are uncorrupted. We do not allow for re-encryption key generation queries between a corrupted and an uncorrupted key. (This represents the restriction that the identities of parties whose security is compromised should be fixed in advance.)

Challenge oracle: This oracle can be queried only once. On input (pk^*, m_0, m_1) , where pk^* is called the challenge key, the oracle chooses a bit $b \leftarrow \{0, 1\}$ and returns the challenge ciphertext $C^* = \text{Enc}(pk^*, m_b)$. (As we note later, the challenge key must be uncorrupted for \mathcal{A} to win.)

²Alternatively, the adversary \mathcal{A} could register her public key with an authority by proving knowledge of the corresponding secret key, as in the registered keys model of Barak et al. [3].

Re-encryption $\mathcal{O}_{\text{renc}}$: On input (pk, pk', C) , where pk, pk' were generated before by KeyGen , if pk' is corrupted and (pk, C) is a derivative of (pk^*, C^*) , then return a special symbol \perp which is not in the domains of messages or ciphertexts. Else, return the re-encrypted ciphertext $C' = \text{ReEnc}(\text{ReKeyGen}(sk, sk'), C)$. Derivatives of (pk^*, C^*) are defined inductively, as follows. (See informal discussion immediately below.)

1. (pk^*, C^*) is a derivative of itself.
2. If (pk, C) is a derivative of (pk^*, C^*) and (pk', C') is a derivative of (pk, C) then (pk', C') is a derivative of (pk^*, C^*) .
3. If \mathcal{A} has queried the re-encryption oracle $\mathcal{O}_{\text{renc}}$ on input (pk, pk', C) and obtained response (pk', C') , then (pk', C') is a derivative of (pk, C) .
4. If \mathcal{A} has queried the re-encryption key generation oracle $\mathcal{O}_{\text{rkey}}$ on input (pk, pk') or (pk', pk) , and $\text{Dec}(pk', C') \in \{m_0, m_1\}$, then (pk', C') is a derivative of (pk, C) .

Decryption oracle \mathcal{O}_{dec} : On input (pk, C) , if the pair (pk, C) is a derivative of the challenge ciphertext C^* , or pk was not generated before by KeyGen , then return a special symbol \perp which is not in the domain D of messages. Else, return $\text{Dec}(sk, C)$.

Decision oracle: This oracle can also be queried only once. On input b' : If $b' = b$ and the challenge key pk^* is not corrupted, then output 1; else output 0.

We say that \mathcal{A} wins the PRE-CCA game with advantage ε if the probability, over the random choices of \mathcal{A} and the oracles, that the decision oracle is invoked and outputs 1, is at least $1/2 + \varepsilon$.

We motivate the above definition of derivatives, which is at the heart of the notion of security. Informally, a pair (pk, C) is a derivative of (pk^*, C^*) if decrypting C with the secret key sk that corresponds to pk would give the adversary “illegitimate information” on the hidden bit b . The first three conditions are obvious. The fourth condition represents the fact that if the adversary has the re-encryption key between pk^* and pk (or alternatively a chain of re-encryption keys $(pk^*, pk'''), (pk''', pk''), \dots, (pk', pk)$) then it is possible that C is the result of legitimately re-encrypting C^* to key pk , in which case decrypting C would give the adversary “illegitimate” information on b .

A first attempt to prevent this may be to not allow the adversary to decrypt *any* ciphertext with respect to key pk . However, such a rule would be too restrictive (resulting in an overly weak definition), since C might be generated *not* as a re-encryption of C^* , in which case we want to allow the adversary to obtain the decryption of C . So we would like to allow the adversary to decrypt to key pk ciphertexts that were “not generated as a re-encryption of C^* ”. However, telling whether C was generated as a re-encryption of C^* may be problematic, especially when the re-encryption algorithm is randomized. In particular, requiring that C is not the result of running the re-encryption algorithm on the challenge ciphertext C^* , for *any* randomness, would result in an overly weak definition, since a re-encryption algorithm could artificially output *any* string in $\{0, 1\}^*$ with negligible but positive probability. Restricting to non-negligible outputs does not seem to suffice either.

We circumvent this difficulty by allowing the adversary to decrypt C , while making sure that if C was generated as a re-encryption of C^* then the adversary gains no information from the decryption query. This definitional approach is reminiscent of the definition of *Replayable CCA security* in [13]. In both cases we want to guarantee “CCA security” while allowing re-encryptions that do not change the decrypted value.

This definition too results in a relaxation of plain CCA security, in the same way that RCCA relaxes CCA security. That is, the present definition allows “harmless mauling” of the ciphertext to a different ciphertext that decrypts to the same value (see discussion in [13].) We view this relaxation as an additional feature of the present definition. In particular, it allows for potential re-randomization and unlinkability of the re-encryption process along the lines of [13, 19, 24]. Such re-randomization is ruled out by strict CCA security. The relations with the UC notions of security provide additional justification for the adequacy of the definition.

Definition 2.5 (Bidirectional PRE-CCA security) *A PRE scheme is Bidirectional PRE-CCA secure for domain D of messages if it is correct for D as in Definition 2.3, and any p.p.t. adversary wins the bidirectional PRE-CCA game only with negligible advantage.*

Remark 2.6 (On Private Re-Encryption Keys)

This definition captures the important feature that, from a ciphertext and its re-encrypted value, the adversary does not learn the re-encryption keys. Suppose the adversary could learn a re-encryption key $rk_{c \leftrightarrow d}$ from queries to \mathcal{O}_{renc} , then it could always win the PRE game by re-encrypting challenge (c, C_b) to the pair (d, D_b) using $rk_{c \leftrightarrow d}$, and then querying the decryption oracle \mathcal{O}_{dec} on (d, D_b) .

2.2 UC Definitions of Security

We formulate a definition for proxy re-encryption within the universally composable (UC) security framework, and investigate its relation with the game-based definition. In addition to providing additional confidence in the adequacy of the notion, a definition in the UC framework carries with it strong composability guarantees.

Providing a UC definition amounts to formulating an *ideal functionality* for proxy re-encryption. However, precisely capturing the security guarantees provided by Definition 2.5 as an ideal functionality within the UC framework turns out to be non-trivial. We thus provide *two* such definitions, namely two formulations of an “ideal proxy re-encryption functionality”, and show that one formulation implies Definition 2.5, whereas the second formulation is implied by Definition 2.5. This allows us to “sandwich” the security guarantees provided by Definition 2.5, when formulated in a composable, simulation based framework.

In a nutshell, the difference between the two formulations is that the weak one does not guarantee security to parties that have setup re-encryption keys with corrupted parties (or with other parties that have setup re-encryption keys with corrupted parties), even if the corresponding proxies are uncorrupted. The strong formulation guarantees security even for such parties, as long as the corresponding proxies are uncorrupted.

We first describe the strong formulation, denoted \mathcal{F}_{sPRE} . Here we extend the ideal public-key encryption functionality \mathcal{F}_{PKE} from [9] to handle re-encryption. To do so, we must allow a single instance of \mathcal{F}_{sPRE} to handle multiple pairs of encryption and decryption keys, as well as multiple re-encryption proxies. This is in contrast to the case of \mathcal{F}_{PKE} , where each instance handles only a single pair of encryption and decryption keys. In addition, in order to simplify the presentation, we assume that the size of the domain D of plaintexts is super-polynomial.

Specifically, as in the case of \mathcal{F}_{PKE} , \mathcal{F}_{sPRE} provides a key generation interface for primary encryption and decryption keys, where the keys themselves are formal (“dummy”) values chosen by the adversary, the encryption key (algorithm) is returned to the registering party, and the decryption key (algorithm) is locally recorded. In addition, there is a re-encryption key generation

Functionality \mathcal{F}_{PRE} (with message domain D)

Key Generation: When receiving (KeyGen, sid) from some party P , send (KeyGen, sid, P) to the adversary. When receiving algorithms (e_P, d_P) and a “corrupted” bit from the adversary, register (P, e_P, d_P) and output $(\text{EncryptionAlgorithm}, sid, e_P)$ to P . In addition, if the “corrupted” bit is set then record P as corrupted.

Re-Encryption Key Generation: When receiving $(\text{ReKeyGen}, sid, P, P', X)$ from registered parties P and P' , send $(\text{ReKeyGen}, sid, P, P', X)$ to the adversary. When receiving algorithm $e_{P,P'}$ and a “corrupted” bit from the adversary, record $(P, P', X, e_{P,P'})$ and $(P', P, X, e_{P,P'})$. Output $(\text{Proxy}, sid, P, P')$ to party X (the proxy). In addition, if the “corrupted” bit is set then record X as corrupted.

A registered party P is called **effectively corrupted** if it is corrupted, or it has registered a re-encryption key with an effectively corrupted party where the proxy is corrupted.

Encryption: When receiving (Enc, sid, m, e) from some party E , do: If m is not in the legitimate encryption domain, output an error message. Else, if $e = e_P$ for some registered party P which is not effectively corrupted, then choose a random message $r_m \leftarrow D$, run algorithm $e_P(r_m)$ (making the necessary random choices) and let c be the outcome. Else (i.e., P is either effectively corrupted or not registered), let $c = e(m)$, and let $r_m = \perp$ where $\perp \notin D$. In either case, record (P, m, r_m, c) and output $(\text{Ciphertext}, sid, c)$ to E .

Re-Encryption: When receiving $(\text{ReEnc}, sid, c, e_P, e_{P'})$ from party X , where $(P, P', X, e_{P,P'})$ is recorded, compute $c' = e_{P,P'}(c)$, and return $(\text{Ciphertext}, sid, c')$ to X . In addition, if there is a record (P, m, r_m, c) then add the record (P', m, r_m, c') .

Decryption: When receiving (Dec, sid, c) from a registered party P , do: If there is a recorded entry (P, m, r, c) for some m and r , then set $\mu \leftarrow m$. Else, compute $r = d_P(c)$. If there is a recorded entry (P', m', r, c') for some P', m', c' , then set $\mu \leftarrow m'$. (If there is more than a single such entry then output an error message.) Else set $\mu \leftarrow r$. Return $(\text{Plaintext}, sid, \mu)$ to P' .

Figure 1: The strong version of the ideal functionality for (bidirectional) proxy re-encryption, \mathcal{F}_{sPRE} . The weak version, \mathcal{F}_{wPRE} , is identical except that a party is called effectively corrupted whenever it has registered a re-encryption key with an effectively corrupted party, even when the proxy is not corrupted. The unidirectional case is captured by modifying the re-encryption key generation interface so that it registers proxies with the ability to re-encrypt only unidirectionally.

interface, where two registered parties P and P' can provide a proxy X with the formal capability to re-encrypt ciphertexts from P to P' and vice versa. As part of this operation the adversary provides \mathcal{F}_{sPRE} with a “dummy” re-encryption key $e_{P,P'}$ between P and P' . These interfaces also allow the adversary to determine the identities of corrupted parties and proxies, with effects described below. It is assumed that the adversary specifies parties as corrupted only before the encryption interface is used for the first time. This manifests the restriction of the modeling to non-adaptive party corruption.

We say that a party P is *effectively corrupted* if it is corrupted, or if it has setup a re-encryption key with a party P' that is either corrupted or effectively corrupted, and in addition the re-encryption proxy for P, P' is corrupted. Essentially, if a party P is effectively corrupted then the adversary can “by definition of the problem” freely learn any message encrypted to P . Indeed, \mathcal{F}_{sPRE} will not guarantee secrecy for effectively corrupted parties.

The encryption interface allows arbitrary encryptors to “formally encrypt” messages, namely to obtain “formal ciphertexts” that will decrypt to the correct message, while guaranteeing unconditional secrecy *as long as the target party is not effectively corrupted*. This is done as follows: If the target party P is not effectively corrupted, then the formal ciphertext c for a message m is computed by applying the dummy encryption algorithm provided by the adversary to a random message r in the domain. Then, the tuple (P, m, r, c) is recorded. Recording the random value r is done in order to correctly decrypt re-encrypted ciphertexts (see details below). If the target party is effectively corrupted then c is computed by applying the dummy encryption algorithm to the actual message m to be encrypted. This represents the fact that no secrecy is guaranteed in this case.

In addition to encryption, \mathcal{F}_{sPRE} provides a re-encryption interface. This interface is simple, and allows a re-encryption proxy between P and P' to transform a dummy ciphertext c to a dummy ciphertext c' in a way that guarantees that when P' decrypts c' it will obtain the same value as that obtained by P when decrypting c . The value of c' is computed by applying the recorded algorithm $e_{P,P'}$ to c .

The decryption interface allows a registered party P to decrypt ciphertexts c addressed to it. The decryption values are computed as follows. First, \mathcal{F}_{sPRE} checks whether there is a recorded tuple which contains ciphertext c and target party P . If so, the corresponding plaintext m is returned. This case guarantees correct decryption for messages that were legitimately encrypted and re-encrypted using \mathcal{F}_{sPRE} 's interfaces. Next, \mathcal{F}_{sPRE} “helps” the adversary (namely, the simulator) by proceeding as follows. First, the value $r = d_p(c)$ is computed, where d_p is the dummy decryption algorithm provided by the adversary when P registered. Next, if there is any recorded tuple where the third field (i.e., the random value chosen at encryption time) is r , then the corresponding plaintext m from that tuple is returned. This represents the case where c was generated by applying the re-encryption algorithm not via the interface provided by \mathcal{F}_{sPRE} . (This situation is possible when a re-encryption proxy gets corrupted and its key becomes exposed. In this case we'd like to allow the adversary to create new ciphertexts that decrypt to the same value as existing ones, while still maintaining the secrecy of the decrypted value. This situation is somewhat reminiscent of the case of RCCA security. We note that for this step to be effective, the domain D of plaintexts should be super-polynomial in order to avoid collisions.) Finally, if no matching tuple is found, \mathcal{F}_{sPRE} outputs r as the decryption value. This step is similar to the case of plain \mathcal{F}_{PKE} , and provides assistance to the simulator in the case where the ciphertext c was simply generated by the environment as an encryption of r to party P , not via the legitimate encryption interface.

As stated more precisely below, functionality \mathcal{F}_{sPRE} captures the same security notion as Definition 2.5, except for the following issue: Consider a party P that's not effectively corrupted, but has setup a re-encryption proxy X with a corrupted party P' . (That is, X is not corrupted and the re-encryption key between P and P' is secret.) Now, assume that a ciphertext c , that was generated for P as an encryption of a message m , is being re-encrypted to a ciphertext c' for P' . Recall that, since P is not effectively corrupted, the ciphertext c is a “dummy ciphertext” that's statistically unrelated to m . In contrast, c' must be a value that decrypts to m under a key that's known to

the adversary and the environment. Furthermore, c' must look like a “plausible re-encryption” of c . Thus, a scheme that realizes \mathcal{F}_{sPRE} must allow a “simulator” to “convincingly” transform a ciphertext c that’s unrelated to m into a ciphertext c' linked to m . This is a property that seems to be an artifact of the simulation paradigm; it is not implied by Definition 2.5. In particular, our scheme does not have this property.

The weaker proxy-re-encryption ideal functionality, noted \mathcal{F}_{wPRE} , bypasses this problem by considering any party that has a re-encryption key with an effectively corrupted party to be effectively corrupted, regardless of whether the relevant proxy is corrupted. This provides a somewhat weaker security guarantee, but the gain is that we can now show that this notion is implied by Definition 2.5, thus providing a “lower bound” of the security provided by Definition 2.5 in a simulation-based framework. Figure 1 presents the re-encryption functionalities, \mathcal{F}_{sPRE} and \mathcal{F}_{wPRE} .

Within the UC framework, security is defined via realizing the corresponding ideal functionality. In our case, we make another restriction: We consider only environments that setup all the keys, including the re-encryption keys, before encrypting messages. More precisely, we say that an environment is *non-adaptive corruption* if, as soon as a message is encrypted or re-encrypted to a party, no additional re-encryption keys are generated between this party and other parties (since this may effectively corrupt parties).

UC-secure PRE schemes are then defined as follows:³

Definition 2.7 (UC-secure PRE schemes) *A PRE scheme Σ is non-adaptively strongly (resp., weakly) UC-secure if the protocol π_Σ UC-realizes \mathcal{F}_{sPRE} (resp., \mathcal{F}_{wPRE}) as defined in [9], with respect to non-adaptive environments.*

2.3 Relationships between Security Definitions

To formalize the relationship between realizing \mathcal{F}_{sPRE} , \mathcal{F}_{wPRE} and Definition 2.5, we first describe a natural transformation from a PRE scheme to a protocol geared towards realizing \mathcal{F}_{sPRE} , \mathcal{F}_{wPRE} . We only consider security of the protocol in a non-adaptive setting, where identities of the corrupted parties, including the corrupted proxies, are fixed in advance. Formally, this is captured by defining the protocol so that corruption requests by the adversary are honored only at the onset of the computation, before any encryption request is made.

Let $\Sigma = (\text{KeyGen}, \text{ReKeyGen}, \text{Enc}, \text{ReEnc}, \text{Dec})$ be a bidirectional re-encryption scheme. The protocol π_Σ is defined as follows. It uses an ideal subroutine \mathcal{F}_{RKG} that takes as input the secret keys sk, sk' of two parties and an identity X of a proxy, and outputs to X a re-encryption key $\text{ReKeyGen}(sk, sk')$.

1. On input (KeyGen, sid) , party P computes $(pk, sk) \leftarrow \text{KeyGen}$ and then outputs the tuple $(\text{EncryptionAlgorithm}, sid, pk)$. At this point, if P gets a corruption request from the adversary then it forwards sk to the adversary. Corruption requests at a later point are ignored.
2. On input $(\text{ReKeyGen}, sid, P, P', X)$, and if not corrupted, party P send $(\text{ReKeyGen}, sid, P, P', X, sk)$ to \mathcal{F}_{RKG} . When X receives output $e_{P,P'}$ from \mathcal{F}_{RKG} , it records this value. If at this

³An alternative way to define weak UC-security is to say that a PRE scheme is weakly UC-secure if it UC-realizes the strong functionality, \mathcal{F}_{sPRE} , with respect to restricted environments that do not ask to re-encrypt ciphertexts directed at parties that are not effectively corrupted, to parties that are effectively corrupted. We prefer the current formulation since it makes the security degradation more explicit, and in particular minimizes the restrictions on the environment.

- point X gets a corruption request from the adversary then it forwards $e_{P,P'}$ to the adversary. Corruption requests at a later point are ignored.
3. On input $(\text{Enc}, \text{sid}, m, e)$, party E compute $c = \text{Enc}(e, m)$ and output $(\text{Ciphertext}, \text{sid}, c)$.
 4. On input $(\text{ReEnc}, \text{sid}, c, P, P')$, proxy X computes $c' = \text{ReEnc}(e_{P,P'}, c)$, and outputs $(\text{Ciphertext}, \text{sid}, c')$.
 5. When receiving $(\text{Dec}, \text{sid}, c)$, party P computes $m = \text{Dec}(sk, c)$, where sk is the locally recorded decryption key, and outputs m . (If there is no locally recorded key then output an error message.)

Theorem 2.8 (Strongly UC implies Game-Based) *Let Σ be a PRE scheme that's non-adaptively strongly UC-secure as in Definition 2.7. Then Σ is bidirectional PRE-CCA re-encryption scheme (as in Definition 2.5).*

Proof. We show that if a PRE scheme $\Sigma = (\text{KeyGen}, \text{ReKeyGen}, \text{Enc}, \text{ReEnc}, \text{Dec})$ is strongly UC-secure (i.e., the protocol π_Σ UC-realizes \mathcal{F}_{sPRE} for domain D and non-adaptive corruptions) then it is PRE-CCA secure for domain D . This part of the proof holds for domains D of any size. The proof follows the same general outline as the proof of equivalence between the game-based and UC formulations for RCCA security [13].

Assume that there exists an adversary F that wins the PRE-CCA game against Σ with advantage ε . We construct an environment \mathcal{E} that distinguishes with probability at least ε between an interaction with \mathcal{F}_{sPRE} with *any* adversary \mathcal{S} , and a real-life interaction with the dummy adversary and $\pi = \pi_\Sigma$. Environment \mathcal{E} invokes a copy of F , and proceeds as follows:

1. Invocations by F of the corrupted (resp., uncorrupted) key generation oracles are answered by calling new corrupted (resp., uncorrupted) parties with a key generation request, and forwarding the obtained keys to F .
2. Re-encryption key generation queries are answered by invoking the two relevant parties with requests to provide a *corrupted* proxy X with a re-encryption key between them, and forwarding the obtained key to F . (Recall that F expects to obtain re-encryption keys only between parties which are not effectively corrupted.)
3. Re-encryption queries are answered by invoking the corresponding proxy with a similar re-encryption request and forwarding the obtained ciphertext to F . (Recall that re-encryption requests can be made either between uncorrupted parties, or alternatively from an uncorrupted party to a corrupted party. Furthermore, \mathcal{F}_{sPRE} provides a meaningful security guarantee even in the latter case: Consider a ciphertext addressed at an uncorrupted party P . Then secrecy of the plaintext is guaranteed until the point where this ciphertext is re-encrypted to a corrupted party.)
4. The challenge query (pk^*, m_0, m_1) is answered by choosing a random bit b , invoking some party to encrypt message m_b with public key pk^* , and returning the response to F as the challenge ciphertext C^* .
5. Decryption queries (pk, C) are answered as \mathcal{O}_{dec} does. That is, \mathcal{E} checks, by making the necessary decryption requests to parties, if C is a derivative of C^* . If so, then \mathcal{E} responds with \perp . Else, \mathcal{E} asks the holder of key pk to decrypt C and forwards the response to F .

6. When F invokes the decision oracle with a guess bit b' , if $b' = b$ then \mathcal{E} outputs “real”; else it outputs “ideal”.

Analyzing \mathcal{E} , we first claim that in an interaction of \mathcal{E} with π , the view of the simulated F is identical to its view of a PRE-CCA interaction with scheme Σ ; thus, in this case F will predict the bit b with non-negligible advantage ε . Next, we claim that, in an interaction of \mathcal{E} with \mathcal{F}_{sPRE} and *any* adversary, the view of F is *statistically independent* of the bit b ; thus F cannot predict b with positive advantage. \square

Theorem 2.9 (Game-Based implies Weakly UC) *Let Σ be a bidirectional PRE-CCA secure re-encryption scheme (as in Definition 2.5) over message space D , where $|D|$ is super-polynomial in the security parameter. Then it holds that Σ is non-adaptively **weakly UC**-secure as in Definition 2.7.*

Proof. We show that if a PRE scheme $\Sigma = (\text{KeyGen}, \text{ReKeyGen}, \text{Enc}, \text{ReEnc}, \text{Dec})$ is PRE-CCA secure then it is weakly UC-secure.

Let $\pi = \pi_\Sigma$ be the protocol constructed from Σ as defined above. We show that π securely realizes \mathcal{F}_{wPRE} . That is, we construct an ideal-process adversary \mathcal{S} such that no environment \mathcal{E} can tell with non-negligible probability whether it interacts with \mathcal{S} and the ideal protocol for \mathcal{F}_{wPRE} , or with π and an adversary \mathcal{A} . (In fact, as shown in [9], it suffices to consider the *dummy adversary* $\tilde{\mathcal{A}}$ which fully follows the instructions of the environment.)

Recall that the only possible activities of \mathcal{S} are, first, to give \mathcal{F}_{wPRE} encryption and decryption algorithms (e_P, d_P) for each registering party P , and second, to give \mathcal{F}_{wPRE} a re-encryption key $e_{P,P'}$ for each pair of parties P and P' that request it. These activities are carried out as follows:

1. To provide algorithms (e_P, d_P) for a registering party P , first run $(pk, sk) \leftarrow \text{KeyGen}$. Algorithm $d_P(\cdot)$ is the decryption algorithm $\text{Dec}(sk, \cdot)$. Algorithm $e_P(m)$ proceeds as follows: If $m = \perp$ then choose a random message $\alpha \leftarrow D$ and run $\text{Enc}(pk, \alpha)$. else, return $\text{Enc}(pk, m)$.
2. To provide a re-encryption algorithm for parties P, P' , where neither party is effectively corrupted. Recall the corresponding secret keys $sk_P, sk_{P'}$, compute $rk_{P \leftrightarrow P'} = \text{ReKeyGen}(sk_P, sk_{P'})$ and return algorithm $\text{ReEnc}(rk_{P \leftrightarrow P'}, \cdot)$.

Analyzing \mathcal{S} , assume for contradiction that there is an environment \mathcal{E} that distinguishes between the real and ideal interactions. We use \mathcal{E} to construct an adversary F that breaks the PRE-CCA security of scheme Σ . More precisely, assume that for some value of the security parameter k we have $\text{EXEC}_{\mathcal{F}_{wPRE}, \mathcal{S}, \mathcal{E}}(k) - \text{EXEC}_{\pi, \tilde{\mathcal{A}}, \mathcal{E}}(k) > \varepsilon$. We show that F guesses the bit b correctly in the PRE-CCA game with probability at least $1/2 + \varepsilon/4p$, where p is the total number of messages encrypted throughout the run of the system. (Without loss of generality, we assume that in each execution of the protocol \mathcal{E} asks to encrypt exactly p messages.)

Adversary F proceeds as follows. F first randomly chooses a number $h \leftarrow \{1, \dots, p\}$. Next, F runs \mathcal{E} on the following simulated interaction with a system running π (and the dummy adversary $\tilde{\mathcal{A}}$). Let m_i denote the i th message that \mathcal{E} asks to encrypt in an execution.⁴

1. When \mathcal{E} activates some *uncorrupted* party P with input $(\text{KeyGen}, \text{sid})$, then F calls the uncorrupted key generation oracle, obtains a key pk and forwards pk to \mathcal{E} .

⁴Without loss of generality we assume that \mathcal{E} only asks to encrypt messages with registered public keys e_P where P is registered and uncorrupted. Indeed, when \mathcal{E} asks to encrypt a message m with another public key e , it receives a value $c = \text{Enc}(e, m)$ that it can compute (or, sample) by itself.

2. When \mathcal{E} activates some *corrupted* party P with input $(\text{KeyGen}, \text{sid})$, then F calls the corrupted key generation oracle, obtains a key pair (sk, pk) and forwards (sk, pk) to \mathcal{E} .
3. When \mathcal{E} activates *uncorrupted* parties P and P' with input $(\text{ReKeyGen}, \text{sid}, P, P', X)$, invoke the re-encryption key generation oracle \mathcal{O}_{rkey} on input $pk_P, pk_{P'}$ and record the returned value $rk_{P \leftrightarrow P'}$ and send it to X . If X is corrupted then also return the value $rk_{P \leftrightarrow P'}$ to \mathcal{E} .
4. Encryption queries are handled as follows:
 - (a) For the first $h - 1$ times that \mathcal{E} asks to encrypt some message, m_i , with key e_i , F returns to \mathcal{E} the ciphertext $C_i = \text{Enc}(e_i, m_i)$.
 - (b) At the h th time that \mathcal{E} asks to encrypt a message, m_h , with key e_h , F queries its challenge oracle with the pair of messages (e_h, m_h, r_h) where r_h is chosen at random from the domain D , obtains the test ciphertext C_h , and forwards C_h to \mathcal{E} . Record the tuple (e_h, m_h, r_h, C_h) .
 - (c) For the remaining $p - h$ times that \mathcal{E} asks to encrypt some message, m_i , with key e_i , F lets the encrypting party return $C_i = \text{Enc}(e_i, r_i)$ where r_i is chosen at random from D . Record the tuple (e_i, m_i, r_i, C_i) .
5. When \mathcal{E} activates party X with input $(\text{ReEnc}, \text{sid}, c, e_P, e_{P'})$, where none of P, P' or X are corrupted, compute $C' = \text{ReEnc}(rk_{P \leftrightarrow P'}, C)$ and return C' to \mathcal{E} . Record the tuple $(e_{P'}, m_i, r_i, C')$.
6. When \mathcal{E} activates party P with input $(\text{Dec}, \text{sid}, C)$, do:
 - (a) If P is corrupted, return $m = \text{Dec}(sk_P, C)$ to \mathcal{E} , where sk_P is the decryption key generated for P .
 - (b) Else, invoke the decryption oracle \mathcal{O}_{dec} with input (pk_P, C) , and let r^* denote the response. (If the response is \perp then let $r^* = r_h$). Next, if there is a recorded tuple (e_P, m, r^*, C) , return m to \mathcal{E} ; else, return r^* to \mathcal{E} . (The rationale here is as follows: As long as the ideal encryption procedure does not choose the same plaintext r twice, the above procedure perfectly mimics the behavior of \mathcal{F}_{wPRE} .)
7. When \mathcal{E} halts, F outputs whatever \mathcal{E} outputs and halts.

Analyzing the success probability of F is done via a hybrid argument. Let the random variable H_i denote the output of \mathcal{E} from an interaction that is identical to an interaction with \mathcal{S} and \mathcal{F}_{wPRE} , with the exception that the first i ciphertexts are computed as an encryption of the real plaintexts, rather than encryptions of random values.

Let B denote the event that two activations of the encryption interface of \mathcal{F}_{wPRE} choose the same random value r . It can be seen that, conditioned on the event that B never occurs, H_0 is distributed identically to the output of \mathcal{E} in the ideal interaction. Also, H_p is statistically close to the output of \mathcal{E} in the interaction with Σ (this follows from the correctness of Σ). Furthermore, in a run of F , if the value C_h that F obtains from its challenge oracle is an encryption of m_h then the output of the simulated \mathcal{E} has the distribution of H_{h-1} . If C_h is an encryption of r_h then the output of the simulated \mathcal{E} has the distribution of H_h . Since \mathcal{E} distinguishes between H_h and H_{h-1} with advantage $\varepsilon/2p$, we have that F predicts the hidden bit b with advantage $\varepsilon/4p$.

Note that we made crucial use of the distinction between \mathcal{F}_{wPRE} and \mathcal{F}_{sPRE} , more specifically of the fact that the environment can never ask to transform a ciphertext directed to an uncorrupted party into a ciphertext directed to a corrupted party. \square

3 Constructions

Following some preliminaries on bilinear forms, we first present a simple construction in the random oracle model, and then show how replace the random oracles with concrete hash functions.

3.1 Bilinear Forms

We write $\mathbb{G} = \langle g \rangle$ to denote that g generates the group \mathbb{G} . Let \mathbf{BSetup} be an algorithm that, on input the security parameter 1^k , outputs the parameters for a bilinear map as $(q, g, h, \mathbb{G}, \mathbb{G}_T, \mathbf{e})$, where \mathbb{G}, \mathbb{G}_T are of prime order $q \in \Theta(2^k)$ and $\langle g \rangle = \langle h \rangle = \mathbb{G}$. The efficient mapping⁵ $\mathbf{e} : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ is both: (*Bilinear*) for all $g \in \mathbb{G}$ and $a, b \in \mathbb{Z}_q$, $\mathbf{e}(g^a, g^b) = \mathbf{e}(g, g)^{ab}$; and (*Non-degenerate*) if g generates \mathbb{G} , then $\mathbf{e}(g, g) \neq 1$.

The security of our schemes depend only on this mild assumption:

Decisional Bilinear Diffie-Hellman (DBDH) [8]: Let $\mathbf{BSetup}(1^k) \rightarrow (q, g, \mathbb{G}, \mathbb{G}_T, \mathbf{e})$, where $\langle g \rangle = \mathbb{G}$. For all p.p.t. adversaries \mathcal{A} , the following probability is strictly less than $1/2 + 1/\text{poly}(k)$:

$$\Pr[a, b, c, d \leftarrow \mathbb{Z}_q; x_0 \leftarrow \mathbf{e}(g, g)^{abc}; x_1 \leftarrow \mathbf{e}(g, g)^d; z \leftarrow \{0, 1\}; z' \leftarrow \mathcal{A}(g, g^a, g^b, g^c, x_z) : z = z'].$$

It will simplify the reading of our proofs to use the following *equivalent* formulation of DBDH, known as **Modified DBDH** [25]. The modified DBDH assumption is identical to the DBDH assumption, except that $x_0 \leftarrow \mathbf{e}(g, g)^{ab/c}$ (instead of $x_0 \leftarrow \mathbf{e}(g, g)^{abc}$).

Lemma 3.1 *If mDBDH is solvable in $(\mathbb{G}, \mathbb{G}_T)$ with probability ε , then DBDH is solvable in $(\mathbb{G}, \mathbb{G}_T)$ with probability ε ; and vice versa.*

Proof. (DBDH \implies mDBDH.) On DBDH input (g, g^a, g^b, g^c, Q) , query the mDBDH solver on input $(g^c, g^a, g^b, g, Q) = (y, y^A, y^B, y^C, Q)$ and output its response. Observe that when $Q = \mathbf{e}(y, y)^{AB/C}$ for the mDBDH solver, by substitution, we have $\mathbf{e}(g^c, g^c)^{(a/c)(b/c)/(1/c)} = \mathbf{e}(g, g)^{abc}$ for the DBDH solver.

(mDBDH \implies DBDH.) Omitted. \square

3.2 PRE Π_{RO} in the Random Oracle Model

Notation and Configuration. Let 1^k be the security parameter and $(q, g, h, \mathbb{G}, \mathbb{G}_T, \mathbf{e})$ be the bilinear map parameters output by $\mathbf{BSetup}(1^k)$. Let $\text{Sig} = (\mathcal{G}, \mathcal{S}, \mathcal{V})$ be a strongly unforgeable one-time signature scheme, where $\ell = \ell_{sig}(k)$ denotes the length of the verification keys output by $\mathcal{G}(1^k)$. Moreover, we assume that the verification key space produced by \mathcal{G} has super-logarithmic minimum

⁵As we've written our bilinear map above, we've restricted ourselves to certain implementations, e.g., supersingular curves. This is done for clarity of exposition. We could implement our construction using a more general bilinear mapping of the form $\mathbf{e} : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$, where efficient isomorphisms between \mathbb{G}_1 and \mathbb{G}_2 may not exist. A benefit of using more general curves is that we could reduce the bits per ciphertext transmitted (see [15] for more details).

entropy; that is, any given key has a negligible chance of being sampled. Let $H : \{0, 1\}^{\leq \ell} \rightarrow \mathbb{G}$ and $F : \{0, 1\}^{\leq \ell} \rightarrow \mathbb{G}$ be two independent hash functions, which we will treat as random oracles.

Define the algorithm *Check* on input a ciphertext tuple (A, B, C, D, E, S) and a key pk as follows:

1. Run $\mathcal{V}(A, (C, D, E), S)$ to verify signature S on message (C, D, E) with respect to key A .
2. Check that $\mathbf{e}(B, F(A)) = \mathbf{e}(pk, D)$ and that $\mathbf{e}(B, h) = \mathbf{e}(pk, E)$.
3. If any of these checks fail, output 0; else output 1.

Scheme $\Pi_{RO} = (\text{KeyGen}, \text{ReKeyGen}, \text{Enc}, \text{ReEnc}, \text{Dec})$ is described as follows:

Key Generation (KeyGen): On input 1^k , select random $x \in \mathbb{Z}_q$. Set $pk = g^x$ and $sk = x$.

Re-Encryption Key Generation (ReKeyGen): On input $sk_X = x$ and $sk_Y = y$, output the bidirectional re-encryption key $rk_{X \leftrightarrow Y} = x/y \pmod q$.

Encryption (Enc): On input pk and a message $m \in \mathbb{G}_T$:

1. Select a one-time signature keypair as $\mathcal{G}(1^k) \rightarrow (svk, ssk)$. Set $A = svk$.
2. Select a random $r \in \mathbb{Z}_q$ and compute

$$\begin{aligned} B &= pk^r & , & & C &= \mathbf{e}(g, H(svk))^r \cdot m, \\ D &= F(sv k)^r & , & & E &= h^r. \end{aligned}$$

3. Run the signing algorithm $\mathcal{S}(ssk, (C, D, E))$, where the message to sign is the tuple (C, D, E) , and denote the signature S .
4. Output the ciphertext (A, B, C, D, E, S) .

Re-Encryption (ReEnc): On input a re-encryption key

$rk_{X \leftrightarrow Y} = x/y$ and a ciphertext $K = (A, B, C, D, E, S)$ under key pk_Y , re-encrypt the ciphertext to be under key pk_X as:

1. Compute $B' = B^{rk_{X \leftrightarrow Y}} = g^{(yr)(x/y)} = g^{xr}$.
2. If $\text{Check}(K, pk_Y) = 1$, output the new ciphertext (A, B', C, D, E, S) ; otherwise, output \perp .

Decryption (Dec): On input a secret key sk and any ciphertext $K = (A, B, C, D, E, S)$, if $\text{Check}(K, g^{sk}) = 1$, then output the message $C/\mathbf{e}(B, H(A))^{1/sk}$; otherwise, output \perp .

Remark 3.2 *Our KeyGen and ReKeyGen algorithms are the same as those originally proposed by BBS [5]. Parties X and Y can compute $rk_{X \leftrightarrow Y}$ efficiently as: X selects a random $r \in \mathbb{Z}_q$ and sends rx to Y and r to the proxy. Y sends rx/y to the proxy, who computes x/y . (Operations are modulo q .) Recall that in a bidirectional scheme, no security is guaranteed if the proxy colludes with either party.*

Theorem 3.3 *If the DBDH assumption holds in $(\mathbb{G}, \mathbb{G}_T)$, then scheme Π_{RO} is Bidirectional PRE-CCA secure for domain \mathbb{G}_T of messages in the random oracle model.*

Proof sketch. Recall that DBDH and mDBDH are equivalent. (It will be less tedious to use mDBDH here.) We begin by observing that if the mDBDH assumption holds in $(\mathbb{G}, \mathbb{G}_T)$, then there exist strongly unforgeable one-time signature schemes. Let $\text{Sig} = (\mathcal{G}, \mathcal{S}, \mathcal{V})$ be such a scheme, where \mathcal{G} has super-logarithmic minimum entropy. Let \mathcal{A} be any p.p.t. adversary. Then, we show how to

construct a p.p.t. adversary \mathcal{B} , with black-box access to \mathcal{A} , that succeeds in breaking the mDBDH assumption with probability:

$$\Pr[\mathcal{B} \text{ breaks mDBDH}] \geq 1/2 + \Pr[\mathcal{A} \text{ breaks } \Pi]/2 - \Pr[\mathcal{A} \text{ breaks Sig}] - q_O \cdot \delta,$$

where \mathcal{A} makes q_O oracle queries and δ is the maximum probability that any given verification key is output by \mathcal{G} (which by assumption is negligible). Also by assumption, $\Pr[\mathcal{A} \text{ breaks Sig}]$ is negligible for any p.p.t. adversary \mathcal{A} . Let us now describe how \mathcal{B} operates.

On mDBDH input (g, g^a, g^b, g^c, Q) , where \mathcal{B} 's goal is to decide if $Q = \mathbf{e}(g, g)^{ab/c}$ or not, \mathcal{B} sets up the global parameters for \mathcal{A} as follows: the description of the groups $\langle g \rangle = \mathbb{G}, \mathbb{G}_T$, their prime order q , and the mapping \mathbf{e} , which are implicit in the mDBDH input, will also be used in the re-encryption game. Set $h = g^{cw}$, where $w \in \mathbb{Z}_q$ is chosen randomly. The system parameters are $(q, g, h, \mathbb{G}, \mathbb{G}_T, \mathbf{e}, H, F)$, where H and F are random oracles. The security parameter is $k \geq |q|$.

\mathcal{B} runs $\mathcal{G}(1^k) \rightarrow (svk^*, ssk^*)$, and records these values. Oracle queries from \mathcal{A} are handled as:

- *Key Generation*: \mathcal{B} chooses a random $x_i \in \mathbb{Z}_q$. If user i is *uncorrupted*, then \mathcal{B} outputs $pk_i = (g^c)^{x_i} = g^{cx_i}$. If user i is *corrupted*, then \mathcal{B} outputs $(pk_i, sk_i) = (g^{x_i}, x_i)$.
- *Hash H* : On input m to hash function H , check to see if pair (m, Y) is recorded in database D_H . If not select random $Y \in \mathbb{G}$, record (m, Y) in D_H and do:

$$H(m) = \begin{cases} g^a & \text{if } m = svk^* \\ Y & \text{otherwise, where } (m, Y) \in D_H. \end{cases} \quad (1)$$

- *Hash F* : On input m to hash function F , check to see if pair (m, Z, s) is recorded in database D_F . If not select random $s \in \mathbb{Z}_q$, record (m, g^s, s) in D_F and do:

$$F(m) = \begin{cases} g^c & \text{if } m = svk^* \\ Z & \text{otherwise, where } (m, Z, s) \in D_F. \end{cases} \quad (2)$$

- *Decryption*: On input (i, K) to \mathcal{O}_{dec} , if $Check(K, pk_i) = 0$, then the ciphertext is not well-formed, so \mathcal{B} halts and returns \perp . Otherwise, \mathcal{B} proceeds as follows.
If user i is corrupted, then $sk_i = x_i$ and \mathcal{B} returns $Dec(sk_i, K)$. If user i is uncorrupted, then \mathcal{B} parses K as (A, B, C, D, E, S) . Since the ciphertext is well-formed, we know that $B = pk_i^r$ and $D = F(A)^r$ for the same value of $r \in \mathbb{Z}_q$. Finally, \mathcal{B} checks that $A \neq svk^*$ and aborts if this is false. If all checks pass, then \mathcal{B} finds $(A, F(A), s) \in D_F$ and decrypts the ciphertext by computing $C/\mathbf{e}(D^{1/s}, H(A)) = C/\mathbf{e}(F(A)^{r/s}, H(A)) = C/\mathbf{e}(g^r, H(A))$.
- *Re-Encryption Key*: On input (i, j) to \mathcal{O}_{rkey} , if one of i and j is uncorrupted and the other is corrupted, then this call is illegal. Otherwise, \mathcal{B} outputs the re-encryption key x_j/x_i .
- *Re-Encryption*: On input (i, j, K) to \mathcal{O}_{renc} , if the value $Check(K, pk_i) = 0$, then the ciphertext is not well-formed, so \mathcal{B} halts and returns \perp . Otherwise, \mathcal{B} parses $K = (A, B, C, D, E, S)$.
 - If users i and j are *both* corrupted or if they are *both* uncorrupted, \mathcal{B} computes the re-encryption key x_j/x_i and executes $ReEnc(x_j/x_i, K)$.
 - If user i is corrupted and user j is uncorrupted, then \mathcal{B} computes $E^{x_j/w} = h^{rx_j/w} = g^{cwr x_j/w} = g^{cr x_j} = pk_j^r = B'$ (where $B = pk_i^r$) and outputs (A, B', C, D, E, S) .
 - If user i is uncorrupted and user j is corrupted, if $A = svk^*$, then \mathcal{B} outputs \perp . Otherwise, \mathcal{B} finds $(A, F(A), s) \in D_F$ and computes $D^{x_j/s} = F(A)^{rx_j/s} = g^{rx_j} = pk_j^r = B'$ (where $B = pk_i^r$) and outputs (A, B', C, D, E, S) .

- *Challenge:* At some point, \mathcal{A} will output a challenge tuple (i, m_0, m_1) , where i is the index of an uncorrupted user. \mathcal{B} responds choosing a random $d \in \{0, 1\}$ and setting:

$$\begin{aligned} A &= \text{svk}^* & , & & B &= (g^b)^{x_i} = pk_i^{b/c}, \\ C &= Q \cdot m_d & , & & D &= g^b = F(A)^{b/c}, \\ E &= (g^b)^w = h^{b/c} & , & & S &= \mathcal{S}_{\text{ssk}^*}(C, D, E). \end{aligned}$$

- *Decision:* Finally, \mathcal{A} will output a guess $d' \in \{0, 1\}$. If $d = d'$, then \mathcal{B} outputs 1 (i.e., mDBDH instance), otherwise \mathcal{B} outputs 0 (i.e., not an mDBDH instance).

The setup, keys and hash responses of \mathcal{B} are perfectly distributed according to a real instance of re-encryption scheme Π_{RO} . The decryption and re-encryption queries are also perfect, except that \mathcal{B} cannot always answer them when $A = \text{svk}^*$. First, consider that *before* the challenge is given, \mathcal{A} has a $q_O \cdot \delta$ chance of querying either oracle on a ciphertext with $A = \text{svk}^*$. *After* the challenge is given, \mathcal{A} 's chance of querying these oracles on a well-formed ciphertext where $A = \text{svk}^*$ and yet the ciphertext is not a derivative of the challenge is $\Pr[\mathcal{A} \text{ breaks Sig}]$. Consider that a *well-formed* ciphertext (A, B, C, D, E, S) decrypts uniquely *regardless* of the corresponding public key. That is, $C = \mathbf{e}(g, H(A))^r \cdot m$ and $D = F(A)^r$ uniquely fixes m . If we recall the definition of a derivative from Definition 2.4, then we see by inspection that if $A = \text{svk}^*$ then for the ciphertext not to be a derivative of the challenge (C, D, E) must not be identical to the challenge ciphertext. If the ciphertext is well-formed, then S is a valid forgery against Sig.

When \mathcal{B} receives an mDBDH instance as input, its challenge ciphertext is also a perfectly distributed, proper encryption of message m_d . Thus, in this case, \mathcal{A} guesses $d' = d$ with the same advantage as it would in a real execution of Π_{RO} . To see this, recall that $H(\text{svk}^*) = g^a$, and in this case, $Q = \mathbf{e}(g, g)^{ab/c}$. Thus, $C/\mathbf{e}(B, H(A))^{1/\text{sk}_i} = \mathbf{e}(g, g)^{ab/c} \cdot m_d / \mathbf{e}(g^{bx_i}, g^a)^{1/cx_i} = m_d$.

When \mathcal{B} does not receive an mDBDH instance, then the challenge ciphertext contains no information about m_d , since ciphertext component C is uniformly distributed in \mathbb{G}_T independent of d , and \mathcal{A} succeeds in predicting $d' = d$ with exactly $1/2$ probability.

Finally, from Lemma 3.1, if there exists a p.p.t. adversary that breaks mDBDH with probability ε , then there exists a p.p.t. adversary that breaks DBDH with probability ε . \square

The following is immediate from Theorems 2.9 and 3.3.

Corollary 3.4 *If the DBDH assumption holds in $(\mathbb{G}, \mathbb{G}_T)$, then bidirectional re-encryption scheme Π_{RO} is non-adaptively, weakly UC-secure with respect to Definition 2.7 in the random oracle model.*

3.3 PRE Π without Random Oracles

Now, we will remove the random oracles from the construction in Section 3.2. To do so, we will change the way hash functions H and F operate. Our setup is similar to that of Π_{RO} , except that:

Function H : We replace random oracle H with a universal one-way hash function family (following Canetti, Halevi, and Katz [10] and Boneh and Boyen [7].) Let \mathcal{H} be a pairwise independent family of hash functions $H : \{0, 1\}^{\leq \ell} \rightarrow \mathbb{G}$, where given an element $x \in \{0, 1\}^{\leq \ell}$ and $y \in \mathbb{G}$, there is an efficient algorithm for sampling $H \in \mathcal{H}$ such that $H(x) = y$.

Function F : While H can be replaced by *any* universal one-way hash satisfying the above constraints, F will be replaced by the following specific one [10, 7]. Let g_2, g_3 be random

generators of \mathbb{G} . Then we define the function $F : \mathbb{Z}_q \rightarrow \mathbb{G}$ as $F(y) \stackrel{\text{def}}{=} g_2^{\tilde{y}} \cdot g_3$, where \tilde{y} is a fixed one-to-one representation of y in \mathbb{Z}_q . (An additional hash function can be included in the public key to implement this one-to-one mapping from y to \tilde{y} .) For simplicity, we will write y instead of \tilde{y} .

To change from scheme Π_{RO} to scheme Π only the encryption algorithm changes as:

Encryption (Enc): On input pk and message $m \in \mathbb{G}_T$, do:

1. Select a one-time signature keypair as $\mathcal{G}(1^k) \rightarrow (svk, ssk)$. Set $A = svk$.
2. Select a random $r \in \mathbb{Z}_q$ and compute

$$\begin{aligned} B &= pk^r & , & & C &= \mathbf{e}(g, H(svk))^r \cdot m, \\ D &= F(svk)^r = (g_2^{svk} \cdot g_3)^r & , & & E &= h^r. \end{aligned}$$

3. Run the signing algorithm $\mathcal{S}(ssk, (C, D, E))$, where the message is (C, D, E) , and denote the signature S .
4. Output the ciphertext (A, B, C, D, E, S) .

Efficiency comparison of Π and Π_{RO} . Scheme Π remains surprisingly efficient compared to Π_{RO} . It requires two additional elements from \mathbb{G} in the global parameters. The cost to compute a ciphertext also increases by two multi-base exponentiations in \mathbb{G} to cover the hashes.

We are now ready to present our main result. Removing oracle H is rather straight-forward, because H was only programmed on one point in the proof of Theorem 3.3. Canetti, Halevi, and Katz [10] present one method of designing a universal one-way hash function that satisfies a series of polynomial constraints, i.e., pairs (x, y) such that $H(x) = y$. Removing oracle F , however, is more involved, because F was programmed for *exponentially* many points. In particular, the proof of Theorem 3.3 required that: (1) F satisfied one (x, y) constraint and (2) that for all other inputs $z \neq x$, the discrete log of $F(z)$ base g was known to the party \mathcal{B} playing the security game with the adversary. In scheme Π , we specially designed a function F such that: (1) F satisfies one constraint (x, y) and (2) such that *together with the other parts of the ciphertext*, \mathcal{B} can compute the necessary information related to $F(z)$ base g for all inputs $z \neq x$.

Theorem 3.5 *If the DBDH assumption holds in $(\mathbb{G}, \mathbb{G}_T)$, then scheme Π is Bidirectional PRE-CCA secure for domain \mathbb{G}_T of messages in the standard model.*

Proof sketch. We now describe how to modify the proof of Theorem 3.3 to our new scheme without random oracles by specifying a different algorithm \mathcal{B} , with black-box access to \mathcal{A} , that succeeds in breaking the mDBDH assumption with probability:

$$\Pr[\mathcal{B} \text{ breaks mDBDH}] \geq 1/2 + \Pr[\mathcal{A} \text{ breaks } \Pi]/2 - \Pr[\mathcal{A} \text{ breaks Sig}] - q_O \cdot \delta,$$

where \mathcal{A} makes q_O oracle queries and δ is the maximum probability that any given verification key is output by \mathcal{G} (which by assumption is negligible.)

On mDBDH input (g, g^a, g^b, g^c, Q) , where \mathcal{B} 's goal is to decide if $Q = \mathbf{e}(g, g)^{ab/c}$ or not, \mathcal{B} sets up the global parameters for \mathcal{A} as follows: \mathcal{B} runs $\mathcal{G}(1^k) \rightarrow (svk^*, ssk^*)$. Next, \mathcal{B} sets the generators $h = g^{cw}$, $g_2 = g^{\alpha_1}$, and $g_3 = g^{-\alpha_1 svk^*} \cdot g^{c\alpha_2}$ for a randomly chosen $w, \alpha_1, \alpha_2 \in \mathbb{Z}_q$.

Finally, \mathcal{B} samples a pairwise independent hash function $H \in \mathcal{H}$ such that $H(\text{svk}^*) = g^a$. The system parameters are $(g, g, \mathbb{G}, \mathbb{G}_T, \mathbf{e}, h, g_2, g_3, H)$. \mathcal{B} runs $\mathcal{G}(1^k) \rightarrow (\text{svk}^*, \text{ssk}^*)$, and remembers these values. Oracle queries from \mathcal{A} to $\mathcal{O}_{\text{rkey}}$ are handled the same as in the proof of Theorem 3.3. The other oracles behave as follows:

- *Key Generation:* \mathcal{B} chooses a random $x_i \in \mathbb{Z}_q$. If user i is *uncorrupted*, then \mathcal{B} outputs $pk_i = (g^c)^{x_i} = g^{cx_i}$. Else, \mathcal{B} sets $sk_i = x_i$, $pk_i = g^{x_i}$, and outputs (pk_i, sk_i) .
- *Decryption:* On input (i, K) to \mathcal{O}_{dec} , if $\text{Check}(K, pk_i) = 0$, then the ciphertext is not well-formed, so \mathcal{B} halts and returns \perp . Next, \mathcal{B} checks that $A \neq \text{svk}^*$ and aborts if this is false. Otherwise, \mathcal{B} proceeds as follows. Notice that in a well-formed ciphertext, $B = pk_i^r$ and $D = F(A)^r$ for the same value of $r \in \mathbb{Z}_q$. \mathcal{B} decrypts the ciphertext by computing

$$t = \frac{D}{B^{\frac{\alpha_2}{x_i}}} \quad , \quad \lambda = \frac{1}{\alpha_1(A - \text{svk}^*)}.$$

Then, \mathcal{B} outputs the message $C/\mathbf{e}(t^\lambda, H(A))$.

Note that when $A \neq \text{svk}^*$, then \mathcal{B} can solve for $t^\lambda = g^r$:

$$t = \frac{F(A)^r}{(pk_i^r)^{\frac{\alpha_2}{x_i}}} = \frac{g_2^{rA} g_3^r}{pk_i^{\frac{r\alpha_2}{x_i}}} = \frac{(g^{\alpha_1})^{rA} (g^{-\alpha_1 \text{svk}^* + c\alpha_2})^r}{(g^{cx_i})^{\frac{r\alpha_2}{x_i}}} = \frac{g^{r\alpha_1(A - \text{svk}^*) + r c\alpha_2}}{g^{rc\alpha_2}} = g^{r\alpha_1(A - \text{svk}^*)}.$$

- *Re-Encryption:* On input (i, j, K) to $\mathcal{O}_{\text{renc}}$, if the value $\text{Check}(K, pk_i) = 0$, then the ciphertext is not well-formed, so \mathcal{B} halts and returns \perp . Otherwise, \mathcal{B} parses $K = (A, B, C, D, E, S)$.
 - If both users i and j are uncorrupted or both are corrupted, \mathcal{B} computes the re-encryption key x_j/x_i and executes the algorithm $\text{ReEnc}(x_j/x_i, K)$.
 - If user i is corrupted and user j is uncorrupted, then \mathcal{B} computes $E^{x_j/w} = h^{rx_j/w} = g^{cwx_j/w} = g^{cx_j} = pk_j^r = B'$ (where $B = pk_i^r$) and outputs (A, B', C, D, E, S) .
 - If user i is uncorrupted and user j is corrupted, if $A = \text{svk}^*$, then \mathcal{B} outputs \perp . Otherwise, \mathcal{B} solves for g^r as it does in decryption (where $B = pk_i^r$), computes $g^{rx_j} = pk_j^r = B'$ and outputs (A, B', C, D, E, S) .
- *Challenge:* At some point, \mathcal{A} will output a challenge (i, m_0, m_1) , where i is the index of an honest user. \mathcal{B} responds choosing a random $d \in \{0, 1\}$ and setting:

$$\begin{aligned} A &= \text{svk}^* \quad , \quad B = (g^b)^{x_i} = pk_i^{b/c}, \\ C &= Q \cdot m_d \quad , \quad D = (g^b)^{\alpha_2} = (g_2^{\text{svk}^*} \cdot g_3)^{b/c}, \\ E &= (g^b)^w = h^{b/c} \quad , \quad S = \mathcal{S}_{\text{ssk}^*}(C, D, E). \end{aligned}$$

- *Decision:* Finally, \mathcal{A} will output a guess $d' \in \{0, 1\}$. If $d = d'$, then \mathcal{B} outputs 1 (i.e., mDBDH instance), otherwise \mathcal{B} outputs 0 (i.e., not an mDBDH instance).

This ends our description of \mathcal{B} . The analysis follows the previous proof. \square

The following is immediate from Theorems 2.9 and 3.5.

Corollary 3.6 *If the DBDH assumption holds in $(\mathbb{G}, \mathbb{G}_T)$, then bidirectional re-encryption scheme Π is non-adaptively, weakly UC-secure with respect to Definition 2.7.*

4 Acknowledgments

We thank Matthew Green, Manoj Prabhakaran and Shengbao Wang for helpful comments and discussions.

References

- [1] Giuseppe Ateniese, Kevin Fu, Matthew Green, and Susan Hohenberger. Improved Proxy Re-encryption Schemes with Applications to Secure Distributed Storage. In *NDSS*, pages 29–43, 2005.
- [2] Giuseppe Ateniese, Kevin Fu, Matthew Green, and Susan Hohenberger. Improved Proxy Re-encryption Schemes with Applications to Secure Distributed Storage. *ACM TISSEC*, 9(1):1–30, Feb 2006.
- [3] Boaz Barak, Ran Canetti, Jesper Buus Nielsen, and Rafael Pass. Universally composable protocols with relaxed set-up assumptions. In *FOCS '04*, pages 186–195, 2004.
- [4] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In *CRYPTO*, vol. 2139, pages 1–18, 2001.
- [5] Matt Blaze, Gerrit Bleumer, and Martin Strauss. Divertible protocols and atomic proxy cryptography. In *EUROCRYPT*, vol. 1403, pages 127–144, 1998.
- [6] Matt Blaze and Martin Strauss. Atomic proxy cryptography. Technical report, AT&T Research, 1997.
- [7] Dan Boneh and Xavier Boyen. Efficient selective-ID secure identity-based encryption without random oracles. In *EUROCRYPT '04*, vol. 3027 of LNCS, pages 223–238, 2004.
- [8] Dan Boneh and Matthew K. Franklin. Identity-Based Encryption from the Weil Pairing. In *CRYPTO '01*, vol. 2139 of LNCS, pages 213–229, 2001.
- [9] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *FOCS*, 2001. See Cryptology ePrint Archive: Report 2000/067.
- [10] Ran Canetti, Shai Halevi, and Jonathan Katz. A forward-secure public-key encryption scheme. In *EUROCRYPT*, vol 2656 of LNCS, pp. 255–271, 2003.
- [11] Ran Canetti, Shai Halevi, and Jonathan Katz. Chosen-ciphertext security from identity-based encryption. In *EUROCRYPT*, vol. 3027 of LNCS, pages 207–222, 2004.
- [12] Ran Canetti and Susan Hohenberger. Chosen-ciphertext secure proxy re-encryption. Cryptology ePrint Report 2007/171, 2007.
- [13] Ran Canetti, Hugo Krawczyk, and Jesper B. Nielsen. Relaxing chosen-ciphertext security. In *CRYPTO '03*, vol. 2729 of LNCS, pages 565–582, 2003.
- [14] Yevgeniy Dodis and Anca-Andreea Ivan. Proxy cryptography revisited. In *NDSS '03*, 2003.

- [15] Steven D. Galbraith, Kenneth G. Paterson, and Nigel P. Smart. Pairings for cryptographers, 2006. Cryptology ePrint Archive: Report 2006/165.
- [16] Shafi Goldwasser and Yael Tauman Kalai. On the impossibility of obfuscation with auxiliary input. In *FOCS '05*, pages 553–562, 2005.
- [17] Philippe Golle, Markus Jakobsson, Ari Juels, and Paul F. Syverson. Universal re-encryption for mixnets. In *CT-RSA*, vol 2964 of LNCS, pages 163–178, 2004.
- [18] Matthew Green and Giuseppe Ateniese. Identity-based proxy re-encryption. In *ACNS '07*, vol. 4521 of LNCS, pages 288–306, 2007.
- [19] Jens Groth. Re-randomizable and replayable adaptive chosen ciphertext attack secure cryptosystems. In *TCC '04*, pages 152–170, 2004.
- [20] Satoshi Hada. Zero-knowledge and code obfuscation. In *ASIACRYPT*, vol 1976 of LNCS, pp. 443–457, 2000.
- [21] Susan Hohenberger, Guy N. Rothblum, abhi shelat, and Vinod Vaikuntanathan. Securely obfuscating re-encryption. In *TCC*, vol. 4392, pages 233–252, 2007.
- [22] Markus Jakobsson. On quorum controlled asymmetric proxy re-encryption. In *PKC '99*, pages 112–121, 1999.
- [23] Masahiro Mambo and Eiji Okamoto. Proxy Cryptosystems: Delegation of the Power to Decrypt Ciphertexts. *IEICE Trans. Fund. Elect. Communications and CS*, E80-A/1:54–63, 1997.
- [24] Manoj Prabhakaran and Mike Rosulek. Rerandomizable RCCA encryption. In *CRYPTO 2007*.
- [25] Amit Sahai and Brent Waters. Fuzzy identity-based encryption. In *EUROCRYPT*, pages 457–473, 2005.
- [26] Tony Smith. DVD Jon: buy DRM-less Tracks from Apple iTunes, March 18, 2005. Available at http://www.theregister.co.uk/2005/03/18/itunes_pymusique.
- [27] Lidong Zhou, Michael A. Marsh, Fred B. Schneider, and Anna Redz. Distributed blinding for El Gamal re-encryption. TR 1924, Cornell CS Dept., 2004.