Proxy Re-Encryption: Analysis of Constructions and its Application to Secure Access Delegation

David Nuñez, Isaac Agudo, Javier Lopez

Network, Information and Computer Security Laboratory (NICS Lab)
Universidad de Málaga, Spain
Email: {dnunez, isaac, jlm}@lcc.uma.es

Abstract

This paper analyzes the secure access delegation problem, which occurs naturally in the cloud, and postulate that Proxy Re-Encryption is a feasible cryptographic solution, both from the functional and efficiency perspectives. Proxy re-encryption is a special type of public-key encryption that permits a proxy to transform ciphertexts from one public key to another, without the proxy being able to learn any information about the original message. Thus, it serves as a means for delegating decryption rights, opening up many possible applications that require of delegated access to encrypted data. In particular, sharing information in the cloud is a prime example. In this paper, we review the main proxy re-encryption schemes so far, and provide a detailed analysis of their characteristics. Additionally, we also study the efficiency of selected schemes, both theoretically and empirically, based on our own implementation. Finally, we discuss some applications of proxy re-encryption, with a focus on secure access delegation in the cloud.

Keywords: Proxy Re-Encryption, Cloud Computing, Access Delegation, Cryptography

1. Introduction

The materialization of the cloud computing paradigm has raised great expectations regarding performance, simplification of business processes, and, foremost, cost reduction. At the same time, these expectations come with new security and privacy risks. Threat scenarios radically change when moving from resources fully controlled by the data owner to resources administrated by third party entities like public clouds. Nowadays, the great majority of cloud systems base their security on preventing potential attackers from accessing internal servers and databases, where users' data is stored. To this end, there is a great variety of measures, with access control systems and network defense techniques being the most prominent. However, the premise of this approach is that the attackers should not be able to break a predetermined security perimeter, where the protected assets (e.g., users' data) reside. These types of measures, although crucial, are often not enough. In addition to external attackers, which may include not only "hackers" but also nation-scale adversaries, accidental data disclosures and insider attacks are also a menacing possibility.

Countermeasures to these threats include the establishment of internal security policies and governance rules, and the reinforcement of access control strategies, but these simply reduce the situation to a trust problem. That is, in the end, there are no actual mechanisms that prevent cloud providers from breaking these measures, either by accident or intentionally, and, in most cases, there is almost no risk of being discovered accessing users' information without their consent. An interesting conflict appears in this scenario – users want to go to the cloud for its benefits, but at the same time, they are unwilling to provide their data to entities that they do not necessarily trust. The adoption of cloud services has been slowed by this dichotomy from the beginning. The introduction of more advanced security mechanisms that enable users to benefit from cloud services and still ensure the confidentiality of their information could help to reduce the trust assumptions in the cloud, and hence, to break the aforementioned dichotomy.

Therefore, it is necessary to depart from the traditional premise that shapes current cloud security and to assume that the measures defined above can be bypassed. A more realistic premise is to assume that the attackers have potential access to users' data [1]. Under this assumption, the only plausible solution is the use of cryptography, so outsourced data is stored in encrypted form. Thus, when traditional security measures fail, attackers will only obtain encrypted data. In a way, the deployed encryption mechanisms become the ultimate safeguard of data confidentiality. A critical principle of this solution is to design the system in such a way that even the provider itself does not have access to the corresponding decryption key; not doing this would again imply a strong trust assumption on the provider. However, a naive combination of this principle with traditional encryption primitives, both symmetric and asymmetric, can hinder the proper processing and sharing of outsourced information and negatively impact the functionality of the system. Therefore, this requirement implies the use of cryptographic primitives that transcend traditional ones, so data confidentiality can be guaranteed, but functionality still remain unaffected.

In this paper we analyze the problem of secure access delegation in the cloud, which is one of the most basic functionalities in this environment, and justify why it cannot be solved by traditional encryption techniques without resorting to complex key management procedures. Different types of cryptosystems have been proposed as solutions, with *Proxy Re-Encryption* as the most prominent candidate. Proxy Re-Encryption (PRE) is a type of public-key encryption that allows a proxy entity to transform ciphertexts from one public key to another, without learning anything about the underlying data. Therefore, from a functional point of view, it can be seen as a means of sharing data securely. The core postulate of this paper is that proxy re-encryption is a prime candidate to construct cryptographically-enforced access control systems where the protected data is stored externally, since it enables dynamic delegation to encrypted information. In a PRE-based solution, private data can reside in the cloud in encrypted form and be shared to authorized users by means of re-encryption, while still remaining confidential with regard to unauthorized parties and the cloud provider itself. In addition, PRE allows the data owner to delegate the access after the data is encrypted, which is important since in a typical access delegation scenario it may not always be possible to identify beforehand the access conditions. The use of encryption to protect data at rest can decrease the risks associated to data disclosures in this kind of scenario, since outsourced information can only be effectively shared if access has been delegated by its owner.

1.1. Contributions

In this paper we analyze the research landscape on proxy re-encryption, and in particular, we make the following contributions:

- We present a profound examination of the proxy re-encryption cryptosystem by reviewing its basic concepts (such as definitions, security models, and properties) and analyzing the main PRE schemes so far, in the light of the attained properties and security notions.
- We provide a comparative analysis of the performance of selected PRE schemes, both from the theoretical and experimental points of view.
- We review the state of research on applications of proxy re-encryption, in particular for the case of access delegation in the cloud. A central standpoint of this paper is that proxy re-encryption constitutes a feasible solution to this problem, and we support this claim by a thorough analysis that includes literature review and study of incentives and economic viability.
- We identify several research directions that cover challenging areas with respect to the fundamentals, construction and application of proxy re-encryption schemes.

1.2. Research Methodology

In order to perform a thorough review on PRE schemes and applications, we followed a methodology to identify and filter publications based on bibliometric criteria. A comprehensive bibliography on PRE schemes and applications, carefully maintained by Shao [2], served as a first raw source of publications. On top of that, we manually added several relevant publications originated from our own study of the literature or from queries for relevant keywords to search engines. The result of this phase is two lists of publications, one focused on schemes (83 papers) and the other on applications (69 papers). Next, it was necessary to filter the list of PRE schemes, given the workload associated to their analysis. Although, in general, most of the papers were preliminarily studied, some of them were filtered out. We used the number of cites for each paper, as measured by Google Scholar, as a heuristic metric of the relevance of the paper. For instance, non-recent publications (e.g., before 2009) which have no cites yet, were marked as not relevant. However, manual verification of the discarded publications was required in order to discard false negatives. Note that we focused exclusively in standard PRE schemes, ruling out other variants (e.g., conditional, certificateless, broadcast, etc.) that imply strong changes to the syntax, security notions and properties, which makes comparisons less meaningful. The result of this phase is a collection of 58 publications (13 schemes and 45 applications).

1.3. Organization

The rest of this paper is organized as follows: Section 2 describes the secure access delegation problem and discuss the suitability of PRE. Section 3 introduces the basic definitions, properties and security models of PRE. In Section 4, we describe the main PRE schemes and analyze them according to their properties; we also perform a theoretical and experimental analysis of the efficiency of selected schemes. Section 5 discusses some possible applications of PRE, with a focus on secure access delegation in the cloud. Finally, Section 6 presents our conclusions and foreseeable research directions for PRE, regarding both constructions and applications.

2. The Secure Access Delegation Scenario

The need for weakening the traditional security assumptions that govern the current security architectures of cloud systems makes the encryption of data prior to outsourcing an essential requirement. At the same time, it is also necessary that the implemented encryption techniques allow to delegate access for sharing purposes, which is one of the most basic functionalities. We refer generically to this setting as the *secure access delegation scenario*. There are, in fact, more advanced functionalities, such as searching, or even computing, over encrypted data; yet, the access delegation functionality is very challenging as the distribution of access rights becomes difficult once the information has been encrypted and outsourced.

In any access delegation scenario (regardless of whether there is outsourcing or not), there are three main separate roles: data producers, data owner, and data consumers. The most generic usage relation in this setting is that multiple data producers generate data which is owned by a data owner, who in

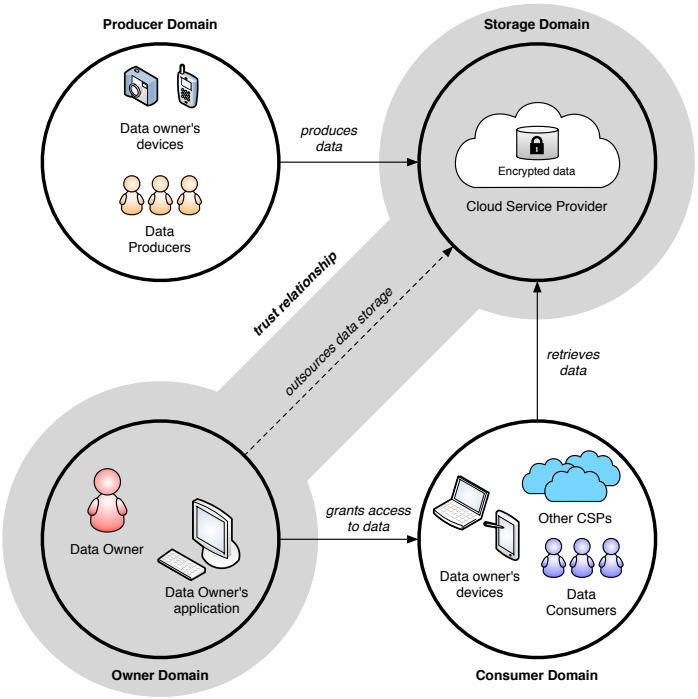

Figure 1: Domains in the Secure Access Delegation Scenario

turn can share it with multiple data consumers. An important aspect of this usage relation is that the data owner and data producers can be separate entities, not necessarily within the same security domain. This latter characteristic has great implications when it comes to designing secure, yet scalable, solutions, since this rules out conventional techniques, such as the use of public-key encryption alone. In addition, under our motivating scenario, we also introduce a fourth role that provides a data storage service, where information is encrypted. Figure 1 depicts these roles and domains, as well as the relationships and interactions between them.

The *Data Producer Domain* comprises those entities that generate data. Since generation does not imply ownership, the distinction between data producer and data owner becomes necessary. However, data producers can participate in the protection of the data from the beginning, by encrypting it from the source, and send it directly to the storage provider, in this way decoupling this interaction from the data owner.

The *Data Owner Domain* is centered on the subject that owns the data whose access is to be delegated. The main function of the data owner is to authorize consumers to access his data. Note that the data owner can also (and most times does) act as a data producer, but this does not rule out scenarios where separate entities participate in data production (e.g., data owner's devices). We assume that the data owner interacts with other actors through a user agent, usually a browser or a specific application, running on a trusted computer. This domain is assumed to be completely trusted by the data owner.

The *Secure Storage Provider Domain* is controlled by specialized entities that steward data owners' information and provide a secure access service, without being able to learn anything. Given that cloud computing is the most prominent instantiation of the considered scenario, we will also refer to this actor as the *cloud provider*. This domain is assumed to be semi-trusted, since we assume that the provider will provide the service correctly, but, at the same time, it may have incentives for trying to read the data. This trust assumption is explained below in more detail.

The *Data Consumer Domain* comprises the entities that are legitimate recipients of the information shared by the data owner, which include not only people, but also third-party services and data owner's devices. Consumers access this information through the storage service provided by the cloud.

A fifth domain, omitted in Figure 1, comprises all the external actors that may have an interest in the protected information, such as hackers and nation-scale adversaries. However, none of these actors are either in charge of managing data or granted with any permission to read it. This type of adversary would have to deal first with traditional security measures (e.g., physical security, firewalls, access control systems, etc.), and in a worst-case scenario, they should see nothing more than encrypted data. If the encryption scheme in use achieves an adequate security notion (e.g., indistinguishability under chosen-plaintext/ciphertext attacks), then we can assume that the data is secure.

Any information that is to be protected must be encrypted from the source (i.e., data producer domain) and decrypted by the legitimate recipients (i.e., data consumer domain). Therefore, from a visibility point of view, the goal is that the storage provider domain and any external actor should only see encrypted data, in any case.

2.1. Threat model and trust assumptions

As argued, in this scenario the most powerful threats may come from the cloud provider domain, since we assume that it can bypass its internal security safeguards. Thus, we consider it as the main adversary. It is also important to remark that, in this scenario, data consumers represent a lesser danger, since they are not as empowered as storage providers: if they are not previously authorized by the data owner, then they can be considered as belonging to the external domain; conversely, if they have been granted access to the protected data, then they are legitimized to read it, so once the information is released, it is impossible for any access control system to prevent it from being disclosed. Additionally, we assume that cloud providers are semi-trusted, in the sense that their behavior is presumed correct with respect to protocol fulfillment, but they may have some incentive to read users' data without their consent. This kind of behavior is usually called *honest-but-curious*.

2.2. Proxy re-encryption as a solution

A naive solution for the secure access delegation problem would be to use conventional encryption techniques (e.g., AES, RSA) and to share the decryption key with the parties designated by the data owner. Symmetric encryption cannot be used alone, since it implies that the same key is shared between producers, owner and consumers or, at least, that for each piece

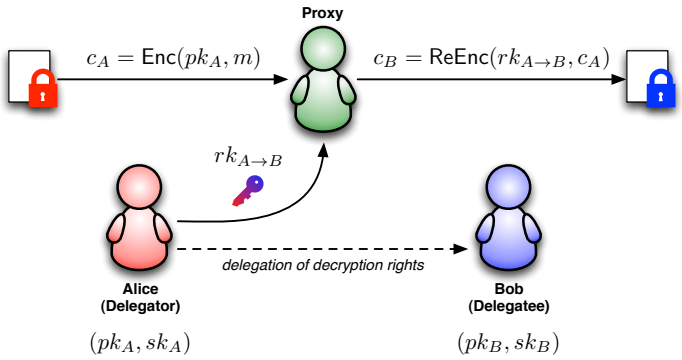

Figure 2: Main actors and interactions in a PRE environment

Table 1: Correspondence between PRE and the secure data sharing scenario

Proxy Re-Encryption	Secure Data Sharing
Delegator	Data Owner
Delegatee	Data Consumer
Proxy	Secure Storage Provider
Anyone	Data Producer
Ciphertexts	Outsourced information
Re-Encryption	Enforcement of access delegation

of encrypted data, producers and owner agree on a key, which is extremely inefficient. With respect to public-key cryptography, the problem in our scenario is that the producers do not necessarily know in advance who are the intended consumers of the encrypted information. Therefore, the only possibility is that they encrypt the data under some common public key controlled by the data owner (e.g., the data owner's public key); this implies that the data owner has to decrypt the data and subsequently encrypt it with a key known by the intended consumers; this *decrypt-and-encrypt* solution requires the data owner to be available online to re-encrypt the data when needed, which is not always possible, apart from being extremely inefficient. The problem gets increasingly complex when one considers multiple pieces of data, and diverse producers and consumers.

It can be seen that this functionality, although simple, cannot be solved by traditional encryption techniques without resorting to complex key management procedures. Proxy Re-Encryption (PRE), however, can be seen as a means for the *delegation of decryption rights*, representing then a natural candidate to construct cryptographically-enforced access control mechanisms compatible with our motivating scenario. For this reason, in the PRE literature, the parties involved are labeled in terms of a relationship of delegation, namely:

- Delegator: This actor is the one that *delegates* his decryption rights using proxy re-encryption. In order to do this he creates a re-encryption key, which he sends to the proxy. We usually refer to the delegator as "Alice".
- Delegatee: The delegatee is granted a delegated right to decrypt ciphertexts that, although were not intended for him in the first place, where re-encrypted for him with permission from the original recipient (i.e., the delegator). This actor usually takes the name "Bob".
- Proxy: It handles the re-encryption process that transforms ciphertexts under the delegator's public key into ciphertexts that the delegatee can decrypt using his private key. The proxy uses the re-encryption key during this process, and does not learn any additional information.

Figure 2 depicts the main actors in a PRE environment and their interactions. Since PRE is a special type of PKE, users also have a pair of public and private keys, as shown in the figure. Hence, anyone that knows a public key is capable of producing ciphertexts intended for the corresponding recipient; conversely, these ciphertexts can only be decrypted using the corresponding decryption key. The distinctive aspect is that ciphertexts can be re-encrypted in order to be decrypted by a different private key than the one originally intended.

Note that there is a direct correspondence between the actors involved in a PRE setting and those associated to the secure access delegation scenario. In a PRE-based solution, private data is initially encrypted by a data producer (which can be any entity that knows the proper public key) and outsourced to a semi-trusted proxy (i.e., storage provider in the cloud). By creating the corresponding re-encryption keys and giving them to the proxy, the data owner is effectively authorizing data consumers to access his data. The proxy enforces these access delegations through the re-encryption process using the corresponding re-encryption keys, while the information that is protected still remains confidential with respect to unauthorized parties and the proxy itself. Table 1 summarizes these parallelisms.

3. Basic definitions and concepts

In this section we provide the basic definitions and concepts that will serve as the basis of our analysis. This includes syntax definition, security models and relevant properties.

The basic idea of a proxy re-encryption scheme is embodied by the ability of a proxy to transform ciphertexts under the public key of Alice into ciphertexts decryptable by Bob; to do so, the proxy must be in possession of a re-encryption key that enables this process. In addition, the proxy cannot learn any information about the encrypted messages, under any of the keys. Most of proxy re-encryption schemes comply with the diagram shown in Figure 3, which depicts the flow of messages, ciphertexts and keys in a PRE environment.

Basically, a PRE scheme has two types of functions: those that generate key material (KeyGen and ReKeyGen), and those that deal with ciphertexts and messages (Enc, ReEnc, and Dec). Some of these functions are like PKE functions: KeyGen produces pairs of public and secret keys, Enc generates a ciphertext that encrypts a message according to a certain public key, while Dec deciphers the ciphertext using the corresponding secret key. On top of these functions, a PRE scheme also defines

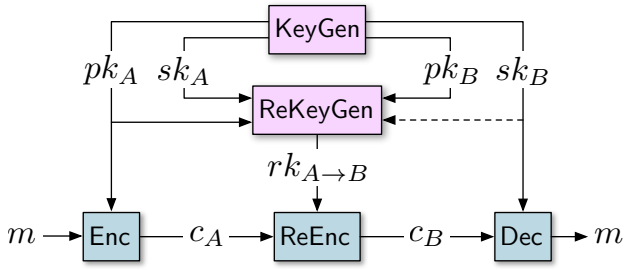

Figure 3: General diagram of a proxy re-encryption scheme

functions to support the re-encryption functionality: ReKeyGen produces a re-encryption key between Alice and Bob, and ReEnc uses this key to transform ciphertexts originally intended for Alice into ciphertexts decryptable by Bob using his secret key.

3.1. Syntax of PRE schemes

The following is a general definition of the syntax of a proxy re-encryption scheme, based on the ones from Canetti and Hohenberger [3] and Ateniese *et al.* [4]:

Definition 3.1. *A proxy re-encryption scheme is a tuple of algorithms (*KeyGen*,* ReKeyGen*,* Enc*,* ReEnc*,* Dec*):*

- KeyGen$(n) \to (pk_A, sk_A)$. *On input security parameter n, the key generation algorithm* KeyGen *outputs a pair of public and secret keys (pk_A, sk_A) for user A.*
- ReKeyGen$(pk_A, sk_A, pk_B, sk_B) \to rk_{A\to B}$. *On input the pair of public and secret keys (pk_A, sk_A) for user A and the pair of public and secret [1] keys (pk_B, sk_B) for user B, the re-encryption key generation algorithm* ReKeyGen *outputs a re-encryption key $rk_{A\to B}$.*
- Enc$(pk_A, m) \to c_A$. *On input the public key pk_A and a message $m \in \mathcal{M}$, the encryption algorithm* Enc *outputs a ciphertext $c_A \in C$.*
- ReEnc$(rk_{A\to B}, c_A) \to c_B$. *On input a re-encryption key $rk_{A\to B}$ and a ciphertext $c_A \in C$, the re-encryption algorithm* ReEnc *outputs a second ciphertext $c_B \in C$ or the error symbol $\perp$ indicating c_A is invalid.*
- Dec$(sk_A, c_A) \to m$. *On input the secret key sk_A and a ciphertext $c_A \in C$, the decryption algorithm* Dec *outputs a message $m \in \mathcal{M}$ or the error symbol $\perp$ indicating c_A is invalid.*

The plaintext and ciphertext spaces are denoted by $\mathcal{M}$ and C, respectively.

[1] If the secret key of the delegatee sk_B is not needed, then the scheme is *not interactive*, since re-encryption keys for B can be produced without her interaction; otherwise, the scheme is *interactive*. See Section 3.3.5 for more details about this.

Note that, although this definition is wide enough, there are others more general, such as the one by Ateniese *et al.* [4], where instead single encryption and decryption algorithms, there are sets of algorithms $\overrightarrow{\text{Enc}}$ and $\overrightarrow{\text{Dec}}$, defined over different ciphertext spaces. Figure 4 shows the relations among plaintext and ciphertext spaces for different kinds of PRE schemes, where (4a) represents PRE schemes with a single ciphertext space, while (4b) shows the case of two ciphertext spaces. Examples of PRE schemes with a single ciphertext space are [5, 3, 6, 7, 8, 9], while [4, 10, 11, 12] are schemes with two ciphertext spaces. Some schemes, such as [13], exhibit an expansive nature on re-encryption, and technically, are defined over an infinite number of ciphertext spaces, since each re-encryption induces a different space. For the sake of simplicity, we opt for the generic (and simpler) syntax with a single ciphertext space.

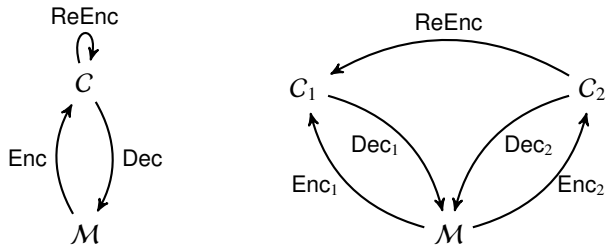

(a) Single ciphertext space (b) Two ciphertext spaces

Figure 4: Transformations between plaintext and ciphertext spaces (extracted from [14])

3.2. Security Models of Proxy Re-Encryption

Being proxy re-encryption an extension of public-key encryption, it is natural that the security models for PRE extend those of PKE. However, the ability to re-encrypt ciphertexts presents an interesting challenge when facing the definitions of security for PRE. On the one hand, PRE constructions have to guarantee the security objectives of the scheme, such as confidentiality and validity of ciphertexts. On the other hand, they have to allow the re-encryption of ciphertexts. Intuitively, both goals seem to conflict with each other.

3.2.1. Security Notions for PRE

Similarly to PKE, the most usual security notions in PRE are *indistinguishability against chosen-plaintext attacks* (IND-CPA) and *indistinguishability against chosen-ciphertext attacks* (IND-CCA) [15]. Both notions capture the inability of an adversary to distinguish ciphertexts for known messages, and differ from each other by when the oracles are available for the adversary. These security notions are formally defined as a two-phase security game: during the first phase, the adversary can use the available oracles, constrained by some conditions; next, before the second phase starts, the adversary freely chooses two messages and receives the challenge ciphertext, which is an encryption of one of them at random; next, he can use the available oracles, again, constrained by some conditions; and finally, he has to guess which of the messages was encrypted.

The restrictions applicable during the game are what actually differentiates the security notions. If the oracles are completely restricted throughout the game, the security notion is IND-CPA; if they are forbidden during the second phase, then one obtains *indistinguishability against non-adaptive chosen-ciphertext attacks* (IND-CCA1), which is seldom targeted in PRE; otherwise, when they are permitted during the whole game, we are alluding to the *indistinguishability against adaptive chosen-ciphertext attacks* (IND-CCA2) notion. The security models provided by Ateniese *et al.* in [4] (which targets CPA-security for unidirectional single-use PRE schemes) and Canetti and Hohenberger in [3] (which targets CCA2-security for bidirectional multi-use PRE schemes) have been used as reference to define and analyze security in most PRE schemes to date, and hence, these are two of the most representative.

In the context of PRE, however, it is possible to define even more fine-grained notions. Since there are two main oracles that consume ciphertexts (i.e., the decryption and re-encryption oracles), and consequently, that can be used to construct chosen-ciphertext attacks, Nuñez *et al.* proposed in [14] a family of attack models (and subsequent security notions), parametrized by the availability of the decryption and re-encryption oracles. These attack models are of the form $\text{CCA}_{i,j}$, where the indices i and j represent the last phase of the game when the decryption and re-encryption oracles, respectively, are available. As an illustration, $\text{CCA}_{0,0}$ represents a notion where both oracles are not available to the adversary (and, therefore, it is equivalent to CPA), while $\text{CCA}_{2,1}$ is a notion where the decryption oracle is present during the whole game, but the re-encryption oracle only before the challenge. In addition, [14] provides a unified framework of security definitions, combining the models of Ateniese *et al.* and Canetti and Hohenberger, which makes it a relevant PRE security model, regardless if the analyzed scheme is unidirectional/bidirectional or single-use/multi-use. We introduce this framework due to its generality and its usefulness to reason about PRE security notions. It also has led to interesting and previously unknown results, such as that it is impossible for a PRE scheme to achieve CCA2-security if it leaks re-encryption keys through queries to the re-encryption oracle [14]. Although it would be possible to analyze each PRE scheme under this general framework, in the end most PRE schemes target and achieve the most representative notions, namely, CPA, CCA1 and CCA2; therefore, in our analysis we will use these more traditional notions.

The definition of PRE security notions is complemented by the definition of certain necessary restrictions to queries from the adversary. Just as in PKE the adversary cannot query the decryption oracle with the challenge ciphertext, in PRE the adversary should not be able to trivially win the game through queries to the decryption, re-encryption and re-encryption key generation oracles; however, such restrictions should allow the adversary to still query the decryption and re-encryption oracles with any unrelated ciphertext. The restrictions to the oracles in PRE are usually defined using the concept of *derivatives of the challenge ciphertext*, which captures the idea of ciphertexts that are connected to the challenge by means of oracle queries. See [3] for more details about the concept of derivatives and [14] for a detailed discussion about attack models and oracle restrictions.

An interesting and orthogonal concept to the aforementioned security notions is *Replayable CCA* (RCCA) originally defined for PKE [16]. This is a weaker form of IND-CCA2 where the adversary is able to make innocuous modifications to ciphertexts, as long as the original message is not altered. This notion naturally fits in the context of PRE, since the goal of PRE is to transform ciphertexts from one user to another, without changing the message.

Finally, it is worth mentioning that there are variations to the traditional PRE security game (in which the challenge is the encryption of a message provided by the adversary, as in PKE). In [17], Canard *et al.* define three different games, which differ from each other by the nature of the challenge ciphertext. The first one is the usual game, where the challenge is an encryption of a message that can be re-encrypted; in the second game, the challenge is an encryption of a message that cannot be re-encrypted; and in the third game, it is a re-encrypted ciphertext (which cannot be re-encrypted, since these games are focused on single-use schemes). However, these security definitions are not very common and not applicable to all types of PRE schemes, so we will not consider them in our analysis.

3.2.2. Assumptions

The security models for PRE are also shaped by some additional assumptions. Perhaps the more important of these assumptions in the PRE context is the *corruption model*. In a *static corruption model*, the adversary must decide in advance whether to *corrupt* a user or not before asking for the generation of the users keypair; with the term corruption we are referring to the adversary knowing the secret keys of the user. In contrast, in an *adaptive corruption model*, the adversary is free to corrupt users in any moment.

Another interesting assumption is related to how the adversary obtains keys. In the *knowledge of secret key model*, the challenger generates the key material of all users, while in the *chosen key model* the adversary can adaptively choose public keys for malicious users [12]. See [18] for more insights about these models.

As in PKE, the vast majority of proxy re-encryption schemes are examples of provable security. In some cases, the proofs of the security are given in the *random oracle model*, where hash functions are assumed to behave as random oracles, an idealization where the hash function is deterministic but its output is uniformly distributed at random in its image domain. When this assumption is not present, we say that we are in the *standard model*.

The security of provable-secure schemes is defined in terms of reductions to hard problems. In other words, the schemes are proven secure, assuming that certain problem is hard. There is a multitude of hardness assumptions, some of them more prominent than others. The most usual are the Computational and Decisional Diffie-Hellman (CDH and DDH) problems in the case of generic groups, the Bilinear Decisional Diffie-Hellman (DBDH) problem in the case of groups with bilinear pairings,

and the Learning With Errors (LWE) problem in the case of lattice-based schemes.

3.3. Properties

In this section we review the main properties associated to proxy re-encryption schemes. Some of these properties are related to the ability of an adversary (including the proxy and the delegatees) to derive key material.

3.3.1. Directionality

This property is associated to direction of the delegation, and is embodied by the re-encryption key. The delegation can be either *unidirectional* or *bidirectional*. Unidirectional delegation means that decryption rights are delegated only from delegator to delegatee, and not in the other direction. Otherwise, the scheme is bidirectional, and it is possible to compute $rk_{B\to A}$ from $rk_{A\to B}$. The fact that a PRE scheme is bidirectional is not necessarily negative, but depending on the situation, it may not be a desirable characteristic. Since PRE can be seen as a mechanism for delegation of decryption rights, bidirectional schemes would be appropriate for scenarios where the trust relationship between delegators and delegatees is symmetric. However, this situation is not common for most applications.

For bidirectional schemes, there is a recurring pattern that appears in some of them, related to how the re-encryption keys are constructed. Informally, these re-encryption keys contain a combination of the delegatee's secret key and the inverse of the delegator's secret key, so when they are used during re-encryption, the delegator's secret key in the ciphertext can be substituted by the delegatee's. More formally, this pattern appears in schemes where public keys are a function of the secret keys, and the secret key space and re-encryption key space are the same. Let us assume that the underlying structure of the secret key space has a multiplicative group structure. Then, in this pattern, re-encryption keys are of the form $rk_{A\to B} = sk_A^{-1} \cdot sk_B$, where sk_A^{-1} is the multiplicative inverse of sk_A. Therefore, it is clear that when this pattern emerges, the converse re-encryption key can be easily computed as $rk_{B\to A} = rk_{A\to B}^{-1}$, so the scheme is bidirectional. A second result is that it is possible that a collusion between the proxy and one of the participants extracts the secret key of the other (e.g., $sk_A = rk_{A\to B}^{-1} \cdot sk_B$). We refer to this pattern as the "*BBS pattern*", as the BBS98 scheme [5] was the first to show it, although [3, 19, 6, 9] are other examples. Usually, the underlying group is multiplicative, although, for example, the scheme in [6] uses an additive group.

3.3.2. Number of Uses

We say a PRE scheme is *single-use* if ciphertexts are re-encryptable just once. This characteristic is usually associated to schemes with multiple ciphertext spaces, since this way the re-encryption can be constructed as a one-way transformations between ciphertext spaces, e.g., by means of pairings. On the contrary, if ciphertexts are re-encryptable multiple times, the PRE scheme is said to be *multi-use*. This characteristic is usually associated to schemes with a single ciphertext space. We further classify multi-use schemes in three types according to the way they achieve this property:

True multi-use schemes. The main characteristic of this type of schemes is that the re-encryption function does not alter the form of the ciphertexts. That is, re-encrypted ciphertexts are identical in shape to original ciphertexts, except maybe for the random elements. For this reason, re-encryption preserves the size of ciphertexts and the running time of the decryption does not depend on the number of re-encryptions of the ciphertext. Examples of this type of multi-use schemes are the BBS98 [5] and CH07 [3] schemes. Section 4.1.1 shows how the BBS98 achieves the true multi-use property.

Expansive multi-use schemes. Some multi-use schemes are based on an iterative method of key encapsulation. The idea is that a random secret, generated for each re-encryption key, acts as trapdoor for a trapdoor function applied to the random secret of the previous re-encryption; the current random secret is encrypted with the encryption function of the PRE scheme, so it can be further re-encrypted. Thus, the ciphertext contains a sort of chain of random secrets for each re-encryption, scrambled using trapdoor functions. The result of this procedure is that the ciphertext size grows linearly with the number of re-encryptions, and, as a consequence, the cost of decryption depends on the number of previous decryptions.

Limited multi-use schemes. Whereas the previous kinds of multiuse schemes support an indefinite number of re-encryptions, some recent proxy re-encryption schemes, in particular those based on lattices [6, 7, 8, 9], present a limited version of the multi-use property, since the re-encryption function introduces noise to the ciphertext. For this reason, and depending on the parameters used, the accumulated noise makes the decryption procedure to fail after a certain number of re-encryptions. This may even happen after just one re-encryption. Thus, schemes of this type are in an intermediate area between single-use and multi-use. For instance, the scheme from Nuñez *et al.* [9] is of this type, as the number of possible re-encryptions varies with the parameters used; in particular, the average number of re-encryptions that is supported varies from 5 to 50. Another interesting example is the scheme from Kirshanova [8], which is allegedly single-use, although that would ultimately depend on the choice of parameters.

3.3.3. Collusion-safeness

This property conveys the safeness of the delegator's secret key against collusion attacks made by the delegatee and the proxy; that is, delegator's secret key sk_A cannot be derived from the re-encryption key $rk_{A\to B}$ and the delegatee's secret key sk_B. However, this property may not be sufficient for some purposes. In most cases, such a collusion does reveal a *weak secret* associated to the delegator's secret key sk_A. In some schemes, such as [4] and [12], this weak secret can be used to create new re-encryption keys or to decrypt re-encryptable ciphertexts; to a certain extent, the weak secret key enables to achieve the functionality of the re-encryption key generation and the decryption functions. Ateniese *et al.* further define the concept of collusion-safeness and formalize a special security notion to this respect, called *master secret security* [4], focused in unidirectional single-use schemes.

3.3.4. Transitivity

A PRE scheme is said to be transitive if the proxy alone is able to re-delegate decryption rights. That is, it can combine re-encryption keys to produce new ones (e.g., from $rk_{A\to B}$ and $rk_{B\to C}$ one can obtain $rk_{A\to C}$). As in the bidirectional case, transitiveness is not negative per se; it depends heavily on the scenario where the scheme is applied. However, transitive delegation is troublesome in general, because it is difficult for the original delegator to foresee the potential delegations which could occur. Schemes that present the BBS pattern are also transitive. This is easy to check: if $rk_{A\to B} = sk_A^{-1} \cdot sk_B$ and $rk_{B\to C} = sk_B^{-1} \cdot sk_C$, then $rk_{A\to C} = rk_{A\to B} \cdot rk_{B\to C}$.

3.3.5. Interactivity

Recall that the general syntax of PRE presented in Section 3.1 included the secret key of the delegatee sk_B in the re-encryption key generation function. If this key is not needed, then the scheme is *not interactive*, since re-encryption keys for delegatees can be produced without their participation; that is, the delegator is able to create a re-encryption key using his own secret key and the delegatee's public key. On the contrary, an interactive scheme implies the participation of the delegatee; in most cases interactivity is an undesired property, as it introduces a communication overhead, and more worryingly, may be vulnerable to collusion-attacks for extracting the secret keys involved. However, interactive schemes can be used to implement *delegation acceptance* by the delegatee [4]; that is, re-encryption keys cannot be generated without the consent and participation of the delegatee.

The reason why the syntax we proposed for the re-encryption key algorithm requires the secret key of the delegatee is to accommodate interactive PRE schemes. We should note, however, that even with these schemes, it is possible to devise a procedure to compute the re-encryption key without any party learning the secret key of the other, eliminating this concern. Therefore, the provided definition is merely syntactical, and does not imply necessarily that the delegator ends up knowing the secret of delegatee; instead, it means that the re-encryption key algorithm somehow uses the secret key of delegatee.

Finally, it is worth mentioning, that the BBS pattern introduced before produces interactive schemes, since the re-encryption key generation necessarily uses the secret keys of both users. Canetti and Hohenberger describe in [3] a simple three-party protocol to compute the re-encryption key that involves the proxy and both users. The final outcome of the protocol is that the proxy learns the re-encryption key but no secret keys, while the users do not reveal their secret keys to each other.

3.3.6. Other properties

Temporary. Some schemes take the temporal dimension into consideration, so the delegation of decryption rights is only valid for a specific period of time. This property was first introduced by Ateniese *et al.* in [4].

Conditional. It is also possible to define keywords that restrict the re-encryption functionality, in a conditional vein. Therefore, conditional PRE represents a fine-grained generalization of traditional PRE. In conditional PRE, re-encryption keys are associated to a certain keyword, so the proxy is only capable of re-encrypting ciphertexts that are tagged with that keyword. This notion was introduced by Weng *et al.* in [20].

Non-transferability. This property, first considered by Ateniese *et al.* in [4], captures the idea of the inability of a collusion of proxy and delegatees to re-delegate decryption rights (i.e., producing new re-encryption keys).

Proxy invisibility. A PRE scheme is said to be proxy-invisible if a delegatee is unable to distinguish a ciphertext computed under her public key from a re-encrypted ciphertext, originally encrypted under another public key [4]. That is, the proxy is "invisible" in the sense that the delegatee cannot discern whether the proxy has transformed the ciphertexts.

Perfect Key-Switching. A stronger property than proxy invisibility is perfect key-switching. Informally, a PRE scheme satisfies this property when the re-encryption cleanly switches one public key for another, without altering the original randomness. Examples of this type of scheme are the BBS98 [5] and CH07 [3] schemes. This property is used in [21] to adapt the Fujisaki-Okamoto transformation to PRE, achieving a weak variant of CCA2.

4. Analysis of Proxy Re-Encryption Schemes

In this section we review and analyze the main proxy re-encryption schemes, which result from the bibliometric process presented in the introduction. A total of 13 publications were selected, and since some of them proposed several schemes, the total number of analyzed schemes is 19. We only considered those proposed schemes which were accompanied by a proof of security. The goal of this analysis is to study the characteristics of each of these schemes, taking in consideration the concepts presented in the previous section. It is out of the scope of this paper to explain in detail the constructions from the cryptographic point of view, due to space limitations. The review makes a comparative analysis possible, which is described subsequently. Finally, we also present a performance analysis of a selection of these schemes, both theoretically and empirically.

4.1. Review of Proxy Re-Encryption Schemes

An early notion, reminiscent of proxy re-encryption, was presented in 1997 by Mambo and Okamoto [22], although their proposal implied that the original recipient must be available for re-encrypting ciphertexts when needed, which is not always feasible. Blaze, Bleumer and Strauss proposed in 1998 the first proxy re-encryption scheme [5], which complies with the established notion of proxy re-encryption. Since then numerous schemes have been proposed. In this Section we review a selection of these schemes. In order to properly identify schemes, each scheme was labeled with the author's initials and year of publication, and if necessary, an additional alphabetic index to distinguish schemes within the same publication (e.g., AFGH06a). Figure 5 shows a timeline depicting the influence

and branches among the proxy re-encryption schemes we selected for our analysis.

4.1.1. BBS98 scheme

This scheme was the first to show an actual construction that complied with the notion of proxy re-encryption. Given the simplicity and elegance of the construction, we describe it here for illustration purposes. The BBS98 scheme is based on ElGamal, and it is constructed upon a group $\mathbb{G}$ of prime order q, with generator g. As in traditional ElGamal, secret keys are of the form $sk = a \in \mathbb{Z}_q$, while public keys are $pk = g^a \in \mathbb{G}$. Ciphertexts are tuples of the form $(pk^r, g^r \cdot m) = ((g^a)^r, g^r \cdot m)$, for a random $r \in \mathbb{Z}_q$. Re-encryption keys are computed as $rk_{A\to B} = \frac{sk_B}{sk_A} \bmod q = \frac{b}{a} \bmod q \in \mathbb{Z}_q$, using the secret keys of both users, which implies that the scheme is interactive. In order to re-encrypt a ciphertext, the proxy simply raises the first component of the ciphertext to the re-encryption key, obtaining $((g^{ar})^{\frac{b}{a}}, m \cdot g^r) = (g^{br}, m \cdot g^r)$. It can be seen that the re-encryption process simply switches one key for another, cleanly, making this scheme an example of the perfect key-switching property. Therefore, the decryption process is exactly the same for all the ciphertexts. In order to retrieve the original message, the recipient (e.g., Bob) uses his secret key to compute $m = c_{B,2} \cdot (c_{B,1})^{\frac{-1}{sk_B}}$. We can also see that the scheme is true multi-use, since re-encrypted ciphertexts are indistinguishable from original ones from the point of view of the re-encryption process. Additionally, it is the first representative of the bidirectional pattern presented in Section 3.3.1 (i.e., the BBS pattern), and as such, has been highly influential to other bidirectional schemes (e.g., the re-encryption key generation procedures of both [3] and [19] are identical to this one, while [6] and [9] have analogous procedures but in a lattice-based setting). Since the scheme follows this bidirectional patter, it is interactive, transitive and not resistant to collusions. Finally, we note that this scheme is proven CPA-secure in the standard model.

4.1.2. AFGH06 schemes

Ateniese, Fu, Green and Hohenberger proposed in [4] the first unidirectional PRE schemes (see Figure 5), based on bilinear pairings. These schemes were also the first to present the idea of multiple ciphertext spaces, as shown in Figure 4b. Original encryptions are referred as "second-level ciphertexts", while re-encrypted ciphertexts are "first-level ciphertexts". Hence, the re-encryption function is a transformation between the second-level ciphertext space and the first-level one. Their first scheme, AFGH06a, is unidirectional, single-use, not interactive, not transitive, and proxy invisible; it is also collusion-safe, although colluding adversaries may compute the weak secret g^{a_1} of the delegator that allows to decrypt second-level ciphertexts and create re-encryption keys. The second scheme, AFGH06b, is a temporary variation of the previous one, where the validity of re-encryption keys is bounded to a specific time period. It is similar to AFGH06a, but it introduces a trusted third party that broadcasts a random value associated to each time period. Thus, this scheme permits the revocation of all previous delegations just by making a change in a global parameter (i.e., the current time period). However, in order to generate re-encryption keys, the delegator needs that the delegatee compute and publish a *delegation acceptance value*, which makes this scheme interactive. Moreover, the re-encryption key generation must occur in the same time period (or before) than the encryption process, which limits its flexibility. Both schemes, AFGH06a and AFGH06b, are proven CPA-secure in the standard model.

The work of Ateniese *et al.* has been highly influential since it provides the first formalizations of PRE syntax and security notions (although focused on unidirectional single-use PRE), as well as an initial description of PRE properties. The authors also propose several applications of PRE, and in particular, an access control server for a secure file system. Additionally, the authors provide the first implementation [23] and performance measurements of a proxy re-encryption scheme.

4.1.3. GA07 schemes

Green and Ateniese proposed in [13] the first identity-based proxy re-encryption schemes (IB-PRE). Being identity-based, these schemes use the identities of the delegator and delegatees as their public keys. The authors present two IB-PRE schemes, which we will refer as GA07a and GA07b. Their first scheme, GA07a, is unidirectional and offers multi-use capabilities, at the expense of being only CPA-secure in the random oracle model. The scheme is also non-interactive, and non-transitive. However, this scheme is not collusion-safe, as the proxy and the delegatee can collude and pool their keys to obtain the secret key of the delegator. This scheme is based in the Boneh-Franklin IBE scheme [24], and, in principal, can reuse an existing deployment, as it uses the same type of parameters and keys. The multi-use property is achieved using an expansive construction, so ciphertexts grow on each re-encryption. This scheme was the first to show this type of multi-use construction, influencing others such as the schemes from Chu and Tzeng [10] (see Figure 5).

Their second scheme, GA07b is single-use, and is allegedly CCA-secure in the random oracle model, although Koo *et al.* present in [25] an attack that uses the re-encryption oracle during the second phase of the security game, so we will consider it as CCA1-secure. As in the previous case, GA07b is not collusion-safe. This scheme is based on the Gentry-Silverberg HIBE scheme [26], and uses the technique proposed by Canetti *et al.* [3] to achieve CCA-security. It is also important to note that, in the same vein as most IBE-based schemes, these schemes require a Key Generation Center that issues a private key for each identity and that must maintain a master secret key. Both schemes are also pairing-based.

4.1.4. CH07 scheme

Canetti and Hohenberger present in [3] the first CCA-secure bidirectional scheme in the standard model. This construction is a variation of an initial construction that is CCA-secure in the random oracle model; we will only consider the version in the standard model, since the modifications are minimal. The CH07 scheme presents the BBS pattern, so it is also interactive, transitive, and not resistant to collusions. CH07 introduces

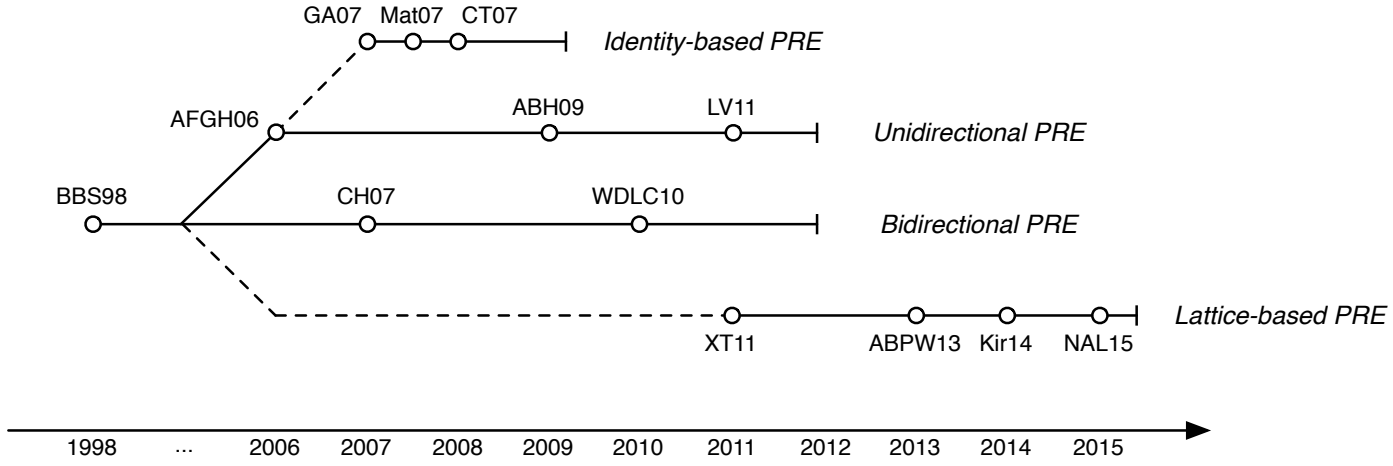

Figure 5: Timeline of selected PRE schemes

CCA-security by integrating a one-time signature into the ciphertexts, following the CHK paradigm [16]. Informally, the solution proposed by the authors is to sign a portion of the ciphertext, which remains unaffected by the re-encryption; otherwise, the signature is invalidated. The remaining part of the ciphertext is what actually changes during re-encryption. In order to validate this part, the signed portion includes some extra information that permits to check that the re-encrypted part has only changed the underlying public key. Another main contribution of the authors is several definitions of CCA-security for bidirectional proxy re-encryption schemes, both game- and simulation-based. In particular, the game-based definition introduced the notion of derivatives of the challenge ciphertext (see Section 3.2.1), which is essential for defining a meaningful CCA security notion for proxy re-encryption. Additionally, the authors show the relationships between their game-based and simulation-based security definitions.

4.1.5. *CT07 schemes*

Chu and Tzeng presented in [10] two IB-PRE schemes, built upon Waters IBE construction [27]. In fact, their security proofs are reductions to the security of the Waters IBE scheme, which in turn is secure under the DBDH assumption. Similarly to the GA07a scheme, the proposed schemes follow the expansive multi-use construction. The first scheme, CT07a, is CPA-secure in the standard model, unidirectional and not interactive, but not collusion-resistant, as shown by [12]. The second scheme, CT07b, is reported CCA-secure, but Shao and Cao show in [28] that anyone is able to re-randomize ciphertexts, which makes the scheme secure in a weaker notion, namely RCCA; this scheme is also not resistant to collusions.

4.1.6. *Mat07 scheme*

Matsuo presented in [29] two IB-PRE schemes. The first one is a rather peculiar proposal, called *"hybrid proxy re-encryption"*, where ciphertexts encrypted using a public key encryption scheme can be re-encrypted to ciphertexts under an identity-based encryption scheme; due to this unusual nature, we will not analyze it here. The second one, Mat07, is an identity-based proxy re-encryption scheme. A significant characteristic of this scheme is that it introduces a new entity called "Re-Encryption Key Generator" that is in charge of producing re-encryption keys and that receives a copy of the master secret key; in principle, the Private Key Generator could also take this role. Thus, in this setting the original delegator is deprived of the re-encryption key generation capabilities, which are now taken by the Re-Encryption Key Generator. The scheme is unidirectional, single-use, collusion-resistant and is proven CPA-secure on the standard model. It is also interactive, as it requires the secret key of the delegatee, and transitive, since it is trivial to compute the re-encryption key $rk_{ID \to ID''}$ from $rk_{ID \to ID'}$ and $rk_{ID' \to ID''}$.

4.1.7. *ABH09 scheme*

Another interesting proposal is presented in [30] by Ateniese *et al.* that defines the notion of *key privacy* for the first time in the context of PRE, which prevents the proxy to derive the identities of both sender and receiver from a re-encryption key. Their scheme, ABH09, is constructed with bilinear pairings and is proven CPA-secure in the standard model. The scheme is unidirectional, single-hop, resistant to collusions, not interactive and not transitive. The key privacy property is proven by means of a specific security game defined by the authors, which serves as a reference to other key-private PRE schemes (such as the one from Aono *et al.* [7]). Additionally, they also present two necessary conditions for a PRE scheme to be key private. The first condition is that the re-encryption function should not be deterministic. The second condition is that a restricted PKE-version of the scheme should also be key private, as defined in [31] for PKE.

4.1.8. *WDLC10 scheme*

Weng *et al.* proposed in [19] two bidirectional schemes without pairings, both CCA-secure in the random oracle model under the hardness of CDH problem. Unlike most of previous proposals, these schemes are not based in bilinear pairings, which makes them, in principle, more efficient. These schemes are designed to achieve CCA-security by integrating

Schnorr signatures [32] with a PRE-version of hashed ElGamal encryption scheme; we note, however, that their security proof is not valid with respect to some re-encryption oracle queries, as shown by Nuñez *et al.* in [21], so we will consider that they achieve a weak form of CCA-security (marked as CCA* in Table 2). The first scheme, WDLC10a, is single-use and follows the BBS pattern for bidirectional schemes. Therefore, it is interactive, transitive and not resistant to collusions. The second scheme, WDLC10b, is very similar to the previous one, but produces re-encryption keys in a different way in order to be non-transitive.

4.1.9. LV11 schemes

Libert and Vergnaud proposed in [12] several unidirectional schemes with RCCA security in the standard model. The first scheme, LV11a, is similar to AFGH06a, but takes ideas from CH07, such as the use of one-time signatures. It is unidirectional, single-use, collusion-resistant, not interactive and not transitive. Since it is defined in the RCCA security model (Section 3.2.1), it is possible to publicly re-randomize ciphertexts. Seo *et al.* [33] detected an error in one of their security proofs, in which the adversary was able to distinguish the simulation from a real attack, and propose a way to amend it. The second scheme, LV11b, is a temporal version of the previous one; as opposed to the temporal scheme from Ateniese, AFGH06b, this scheme is not interactive. Additionally, it is the first PRE scheme that considers the chosen-key model. In addition to these schemes, Libert and Vergnaud also propose several variations, without providing extensive descriptions and proofs. In particular, they show how the temporal scheme can be extended to allow temporal windows, instead of just single periods of time; they also introduce a conditional version of their scheme. Another remarkable contribution of the work of Libert and Vergnaud is the use of the chosen-key model for the first time (see Section 3.2.2)

4.1.10. XT10 scheme

Xagawa and Tanaka presented in [6] the first PRE scheme based on lattices (see Figure 5). In particular, this scheme is based on the Learning With Errors (LWE) problem. This scheme is reminiscent to the BBS98 scheme, but adapted to a lattice-based setting. It presents the same BBS pattern, although represented additively, so $rk_{A\to B} = sk_B - sk_A$. Therefore, this scheme is bidirectional, interactive, transitive and not resistant to collusions. The scheme is also multi-use, more specifically, of the limited multi-use type.

4.1.11. ABPW13 scheme

Aono *et al.* proposed in [7] a lattice-based encryption scheme, ABPW13, which is proven CPA-secure in the standard model. This scheme is based upon a lattice cryptosystem from Lindner and Peikert [34], whose hardness relies on the LWE problem. ABPW13 is unidirectional, interactive, not transitive and limited multi-use. The scheme is not resistant to collusions from the proxy and the delegatee, as shown recently by [35]. The authors also state that ABPW13 is key-private; however, Nishimaki and Xagawa recently noted in [36] that, although the scheme ensures the anonymity of the delegator, the delegatee is exposed trivially, since its public key is contained in the re-encryption key, achieving then a partial form of key-privacy. On top of this scheme, the authors construct a "CCA-secure" version in the random oracle model, using the generic Fujisaki-Okamoto conversion. However, as shown in [21], this version is flawed because it does not perform the necessary validation step during decryption of re-encrypted ciphertexts. If one forces this check, decryption will always fail in the case of re-encryption, breaking the correctness of the scheme.

4.1.12. Kir14 scheme

Kirshanova presented in [8] a lattice-based PRE scheme, which was reported to be CCA1-secure. The scheme is based on the CCA1-secure PKE scheme from Micciancio and Peikert [37], and its security is associated to the hardness of the LWE problem. Basically, Kirshanova extended the original PKE scheme to support re-encryptions, using the technique presented in [37] for trapdoor delegation. The scheme Kir14 is unidirectional, not interactive and resistant to collusions. The authors state that the scheme is single-use, although technically we consider it of the limited multiuse type, as the choice of parameters is what determines the number of possible re-encryptions. The authors prove on what conditions the re-encryption process preserves the correctness, but this proof only considers one hop; still, it would be possible that some sets of parameters permit multiple re-encryptions.

In 2015, Nuñez *et al.* described an attack to this scheme, and show that it does not satisfy CCA1 security[14]. The attack makes use of a property first introduced by Canetti and Hohenberger in [3], called *privacy of re-encryption keys*, which captures the notion of an adversary being unable to extract re-encryption keys from the knowledge of original and re-encrypted ciphertexts. Nuñez *et al.* generalized this property and show its impact in the security notions that consider the re-encryption oracle; in particular, they show how the Kir14 scheme leaks re-encryption keys through the re-encryption oracle, which implies it cannot not achieve a strict CCA1 security notion. Recently, Fan and Liu [38] show that there is another error in their security proof that does not allow to answer correctly to decryption and re-encryption key generation queries simultaneously, which together with the attack from Nuñez *et al.* , implies that the scheme Kir14 is at most CPA-secure.

4.1.13. NAL15 schemes

Nuñez *et al.* describe in [9] schemes based on the NTRU cryptosystem [39]. The first scheme, NAL15a (originally named "NTRUReEncrypt"), is bidirectional, multihop, and interactive, but not collusion-resistant, as it follows the usual BBS pattern. As the original NTRU scheme, the underlying structure of this scheme is a polynomial ring, so the operations performed are simply additions and multiplications of polynomials, which can be computed very efficiently, and can even be parallelized. Therefore, the key strength of this scheme is its performance. The authors present experimental results that show that this scheme outperforms others by an order of magnitude, in a similar way than NTRU with respect to other PKE schemes. How-

ever, as the original NTRU scheme, it lacks a proof of security. The parameters for NTRU are then computed with regard to its resistance to known attacks. To overcome this problem, the authors also propose a provably-secure variant that is proven CPA-secure under the hardness of the Ring-LWE problem, a variation of the LWE problem under a polynomial ring; this scheme is an extension of a provable-secure version of NTRU proposed by Stehlé and Steinfeld [40]. This second PRE scheme, NAL15b, is identical to the previous one with respect to the properties.

4.2. Summary and Comparison

The schemes analyzed in the previous section are also depicted in Figure 5, temporally organized and separated by the major branches. It can be seen that the BBS98 scheme has had a huge influence on the vast majority of PRE schemes to date. In fact, most bidirectional schemes follow the same structure. Unidirectional schemes are mostly based on pairings, following the same ideas introduced by the AFGH06 schemes, which in turn, can be seen as a transposition of the BBS98 scheme to a pairing setting, where the re-encryption keys are "protected" as exponents. Identity-based PRE is highlighted as one of the major branches of PRE, first started with the GA07 schemes, which also introduced the expansive multi-use construction. An interesting trend appeared in 2011 with the first lattice-based PRE scheme; since then, more lattice-based schemes have appeared. However, to date, lattice-based schemes only achieve weak security notions.

We now summarize and compare the analyzed schemes in a single table. Table 2 shows this comparison, according to the following criteria:

- Directionality: A single arrow ($\rightarrow$) is used to represent a unidirectional scheme, whereas a double arrow ($\leftrightarrow$) denotes a bidirectional one.
- Use: Single-use schemes are represented by the letter 'S', true multi-use schemes by the initials 'TM', expansive multi-use schemes by 'EM', and limited multi-use schemes by 'LM'.
- Number of ciphertext spaces (#spaces): This column specifies the number of ciphertext spaces on which the scheme is defined.
- Security: The achieved notion of security (e.g., CPA, CCA, RCCA) is presented, as well as whether is based or not in the random oracle model (RO or SM).
- Based on: This column describes the underlying structure on which the scheme is constructed. The possible values are "Group" for generic groups, "Pairing" for groups with bilinear pairings, and "Lattice" for lattice-based schemes. Additionally, the hardness assumption is also presented.
- Collusion resistance: Schemes that are resistant to collusions are marked with ✓, and with × otherwise.
- Non-transitive: Schemes that are non-transitive are marked with ✓, and with × otherwise.
- Non-interactive: Schemes that are non-interactive are marked with ✓, and with × otherwise.
- The last column shows additional relevant characteristics of the analyzed schemes

4.3. Performance

This section is devoted to analyzing the performance of several PRE schemes. This analysis will be made from two points of view: theoretical and empirical. The latter kind of study is is seldom tackled in the literature.

As described in the general syntax of PRE, there are several functions in a PRE scheme. However, from a functional point of view, we are mainly interested in studying the computational costs for three of them: encryption of messages intended to be re-encrypted, re-encryption of ciphertexts, and decryption of re-encrypted ciphertexts. This decision is justified by the actual use of proxy re-encryption in applications, where ciphertexts are meant to be re-encrypted. Therefore, functions such as non-delegatable encryption (denoted usually by "first-level encryption") are out of the scope of this analysis, since could be achieve by means of traditional encryption schemes.

4.3.1. Theoretical analysis

In this section we analyze and compare schemes from the theoretical standpoint. In order to present an analysis that supports meaningful comparisons, we will focus on group- and pairing-based PRE schemes, since all of them share similar underlying structures, namely, cyclic groups, whose operations' computational cost can be used as measuring units. Therefore, we omit from the main analysis the schemes that are lattice-based, since their underlying structure varies greatly between each other (e.g., polynomials vs matrices), the theoretical analysis of costs is very intricate (e.g., vector-matrix multiplications of various dimensions), and the factors that influence computational costs cannot be standardized, which invalidates any meaningful comparison between different schemes. However, a separate study of lattice-based schemes is provided at the end of this section.

Usually, computational costs are analyzed in terms of the main operations performed, which in this case are the exponentiation and the pairing. These costs are denoted by t_e and t_p, respectively. It is important to note that, in the case of pairing groups, the computational costs of operations are different depending on the group where are performed. For this reason, when necessary we will make the distinction between operations in $\mathbb{G}$ (denoted by t_e) and in $\mathbb{G}_T$ (denoted by t_{e_T}). Some of the schemes make use of a one-time signatures (OTS) to reach CCA-security. In this case, the signature of a message using OTS, including key pair generation, is denoted by t_S, while the verification of a signature is denoted by t_V. Finally, the decryption of expansive multi-use schemes varies with the number N of re-encryptions.

Table 2: Comparison of PRE schemes

Scheme	Dir.	Use	#spaces	Security	Based on	Coll. res.	Non-trans.	Non-int.	Other characteristics
BBS98	↔	TM	1	CPA (SM)	Group (DDH)	×	×	×	Perfect key-switching
AFGH06a	→	S	2	CPA (SM)	Pairing (eDBDH)	✓	✓	✓	Proxy invisible
AFGH06b	→	S	2	CPA (SM)	Pairing (DBDH)	✓	×	×	Temporary
CH07	↔	TM	1	CCA (SM)	Pairing (DBDH)	×	×	×	Perfect key-switching
GA07a	→	EM	N	CPA (RO)	Pairing (DBDH)	×	✓	✓	Identity-based
GA07b	→	S	2	CCA1 (RO)	Pairing (DBDH)	×	✓	✓	Identity-based
CT07a	→	EM	N	CPA (SM)	Pairing (DBDH)	×	✓	✓	Identity-based
CT07b	→	EM	N	RCCA (SM)	Pairing (DBDH)	×	✓	✓	Identity-based
Mat07	→	S	1	CPA (SM)	Pairing (DBDH)	✓	×	×	Identity-based
ABH09	→	S	2	CPA (SM)	Pairing (eDBDH)	✓	✓	✓	Key-private
WDLC10a	↔	S	2	CCA* (RO)	Group (CDH)	×	×	×	
WDLC10b	↔	S	2	CCA* (RO)	Group (CDH)	×	✓	×	
LV11a	→	S	2	RCCA (SM)	Pairing (3w-DBDHI)	✓	✓	✓	
LV11b	→	S	2	RCCA (SM)	Pairing (1w-DBDHI)	✓	✓	✓	Temporary, Chosen-Key model
XT10	↔	LM	1	CPA (SM)	Lattice (LWE)	×	×	×	
ABPW13	→	LM	1	CPA (SM)	Lattice (LWE)	×	✓	×	Partial Key-Private
Kir14	→	LM	1	CPA (SM)	Lattice (LWE)	✓	×	✓	
NAL15a	↔	LM	1	-	Lattice (NTRU)	×	×	×	
NAL15b	↔	LM	1	CPA (SM)	Lattice (Ring-LWE)	×	×	×	

The results of this analysis are presented in Table 3. It is interesting to see the great difference in performance between schemes. For instance, the BBS98 schemes only needs a few exponentiations, as opposed to more complex schemes, such as LV11a and CH07, which require up to five pairing operations for the decryption. This is due to the additional costs incurred by the achievement of CCA-security. Notable schemes are WDLC10a and WDLC10b, since they provide a weak form of CCA-security but only require a few exponentiations. Note also how in the case of expansive multi-use schemes, the cost of decryption depends on the number of re-encryptions.

With regard to space costs, these are mainly driven by the size of group elements. As in the case of computational costs, some schemes make use of one-time signatures; the size of a signature is denoted by $|\sigma|$, while the size of a verification key is $|svk|$. The cost of ciphertexts in some schemes also depends on the size of the original message, $|m|$. Finally, in some cases, elements of $\mathbb{Z}_q$ are also used. Table 3 shows the results of this analysis. It can be seen how expansive schemes increase the size of ciphertext linearly on each re-encryption. Similarly to the analysis of computational costs, CPA-schemes usually have lower size of ciphertexts, since CCA-schemes need to include additional elements for the validation of the ciphertexts (e.g., signatures).

Finally, although the computational cost analysis of lattice-based schemes is omitted for the reasons described before, it is still possible to study their main operations from the point of view of asymptotic computational complexity. Table 4 summarizes the results of this analysis. However, we remark that this study does not allow for direct comparisons between schemes, since it is of asymptotic nature (i.e., assuming n approaches infinity). It is also worth mentioning that lattice-based schemes usually sample random values from special Gaussian distributions, each of them using different sets of parameters, and the cost of this sampling process may not be negligible, although we have omitted it since it can be done off-line, in advance.

4.3.2. Empirical analysis

We complement the theoretical analysis with an experimental evaluation, based on previous work [9]. A selection of these schemes was implemented, with representation of all the kinds of PRE schemes; this time, lattice-based schemes were also included. The selected schemes are BBS98, AFGH06a, WDLC10a, LV11a, ABPW13, and NAL15a. Our execution environment was an Intel Core 2 Duo processor @ 2.66 GHz with 4 GB of RAM.

Group-based schemes were implemented in Java using elliptic curve cryptography over a prime field. The NIST P-256 curve was used, which provides 128 bits of security [41].

Pairing-based schemes were implemented using the jPBC library [42], a pairing-based cryptography library for Java. As for the cryptographic details, all the studied PRE schemes use symmetric pairings (also called Type 1 pairings), which we instantiated with a supersingular curve over fields of large prime characteristic (with embedding degree 2). In order to achieve 128 bits of security (against the discrete logarithm problem in $\mathbb{G}$ and $\mathbb{G}_T$) we took a 256-bit group order and 3072 bits for the field size [43][44]. For efficiency reasons, we have made extensive use of exponentiation and pairing preprocessing of frequently-used elements.

With regard to lattice-based schemes, we implemented NAL15a in Java [45] and used the *ees1171ep1* parameter set from [46]

Table 3: Computational and space costs of selected PRE schemes

Scheme	Encryption	Re-encryption	Decryption	Ciphertext size	Re-encrypted ciphertext size
BBS98	$2t_e$	t_e	t_e	$2\|\mathbb{G}\|$	$2\|\mathbb{G}\|$
AFGH06a	$t_e + t_{e_T}$	t_p	t_{e_T}	$\|\mathbb{G}\| + \|\mathbb{G}_T\|$	$2\|\mathbb{G}_T\|$
AFGH06b	$t_p + t_e + t_{e_T}$	t_p	t_{e_T}	$\|\mathbb{G}\| + \|\mathbb{G}_T\|$	$2\|\mathbb{G}_T\|$
CH07	$t_p + 4t_e + t_{e_T} + t_S$	$4t_p + 2t_e + t_V$	$5t_p + 2t_e + t_{e_T} + t_V$	$\|svk\| + 3\|\mathbb{G}\| + \|\mathbb{G}_T\| + \|\sigma\|$	$\|svk\| + 3\|\mathbb{G}\| + \|\mathbb{G}_T\| + \|\sigma\|$
GA07a	$t_p + 2t_e$	t_p	Nt_p	$\|\mathbb{G}\| + \|\mathbb{G}_T\|$	$(N+1)(\|\mathbb{G}\| + \|\mathbb{G}_T\|)$
GA07b	$t_p + 3t_e$	$4t_p + 2t_e$	$2t_p + t_e$	$2\|\mathbb{G}\| + \|\mathbb{G}_T\| + \|m\|$	$\|\mathbb{G}\| + \|\mathbb{G}_T\| + 2\|m\| + \|id\|$
CT07a	$2t_e + t_{e_T}$	$2t_p$	$(N+2)t_p$	$2\|\mathbb{G}\| + \|\mathbb{G}_T\|$	$(N+2)\|\mathbb{G}\| + (N+1)\|\mathbb{G}_T\|$
CT07b	$3t_e + t_{e_T} + t_S$	$2t_e + t_V$	$(7N+3)t_p + 2t_e + (N+1)t_V$	$3\|\mathbb{G}\| + \|\mathbb{G}_T\| + \|svk\| + \|\sigma\|$	$(6N+3)\|\mathbb{G}\| + (N+1)(\|\mathbb{G}_T\| + \|svk\| + \|\sigma\|)$
Mat07	$t_p + 3t_e + t_{e_T}$	$t_p + t_{e_T}$	$2t_p$	$2\|\mathbb{G}\| + \|\mathbb{G}_T\|$	$2\|\mathbb{G}\| + \|\mathbb{G}_T\|$
ABH09	$2t_e + t_{e_T}$	$4t_p + 2t_{e_T}$	t_{e_T}	$2\|\mathbb{G}\| + \|\mathbb{G}_T\|$	$2\|\mathbb{G}_T\|$
WDLC10a	$3t_e$	$3t_e$	$2t_e$	$2\|\mathbb{G}\| + \|m\| + \|\mathbb{Z}_q\|$	$2\|\mathbb{G}\| + \|m\|$
WDLC10b	$3t_e$	$3t_e$	$2t_e$	$2\|\mathbb{G}\| + \|m\| + \|\mathbb{Z}_q\|$	$\|\mathbb{G}\| + \|m\|$
LV11a	$3t_e + t_{e_T} + t_S$	$2t_p + 4t_e + t_V$	$5t_p + t_e + t_{e_T} + t_V$	$\|svk\| + 2\|\mathbb{G}\| + \|\mathbb{G}_T\| + \|\sigma\|$	$\|svk\| + 4\|\mathbb{G}\| + \|\mathbb{G}_T\| + \|\sigma\|$
LV11b	$5t_e + t_{e_T} + t_S$	$4t_p + 8t_e + t_V$	$9t_p + t_e + t_{e_T} + t_V$	$\|\mathbb{Z}_q\| + \|svk\| + 3\|\mathbb{G}\| + \|\mathbb{G}_T\| + \|\sigma\|$	$\|\mathbb{Z}_q\| + \|svk\| + 7\|\mathbb{G}\| + \|\mathbb{G}_T\| + \|\sigma\|$

Table 4: Asymptotic complexity of lattice-based PRE schemes

Scheme	Encryption	Re-encryption	Decryption
XT10	$O(n^2 \log n)$	$O(n^2)$	$O(n^2)$
ABPW13	$O(n^3)$	$O(n^3 \log^2 n)$	$O(n^2)$
Kir14	$O(n^3 \log n)$	$O(n^2 \log^2 n)$	$O(n^3 \log^2 n)$
NAL15a	$O(n \log n)$	$O(n \log n)$	$O(n \log n)$
NAL15b	$O(n \log n)$	$O(n \log n)$	$O(n \log n)$

Table 5: Experimental performance of several PRE schemes (in ms.)

Scheme	Encryption	Re-encryption	Decryption
BBS98	11.07	11.48	11.21
AFGH06a	22.76	83.52	13.76
WDLC10a	22.52	22.29	11.89
LV11a	155.27	386.93	443.87
ABPW13	1.17	20.50	0.47
NAL15a	0.43	1.15	1.22

(designed for 256 bits of security). The ABPW13 was implemented using the SageMath software, using the set of parameters proposed by the authors that correspond to 143 bits of security and 128-bit plaintexts.

Table 5 shows the cost in ms. of the main operations of the selected PRE schemes. These figures were measured as the mean CPU time of 10.000 executions for each operation. Roughly, the theoretical costs presented before are translated to an empirical setting. It is worth mentioning the high performance that lattice-based schemes exhibit. Both schemes, ABPW13 and NAL15a, present similar figures regarding encryption and decryption, although ABPW13 is one order of magnitude slower when it comes to re-encryption. Note, however, that the results of these experiments are highly dependent on implementation issues, such as the choice of programming language, parameters (type of curve, size of fields, type of pairing, etc.), the underlying libraries and the use of preprocessing. For this reason, any comparative analysis based on these figures has to consider these aspects.

5. Applications of Proxy Re-Encryption

The last part of this paper is devoted to the analysis of the applications of proxy re-encryption. As described in the introduction, we have performed a review of almost 70 papers regarding applications, of which 45 were finally analyzed in detail. We also followed a bibliometric approach for drawing conclusions on this part. In particular, we classified each of the reviewed application according to certain criteria: objective, scenario and functionality. The first criterion is related to the security objectives that are intended to tackle with the application of PRE; possible objectives are confidentiality, privacy, authentication and accountability. The second criterion was clearly dominated by the cloud scenario, although other scenarios are also considered, such as wireless networks. The third criterion is associated to the intended functionality that is constructed with PRE, being access control and key management the most prominent.

Once the publications were classified, we counted the number of occurrences for each category and obtained the following results, summarized in Table 6. From the viewpoint of the security objective, a vast majority of the papers tackled confidentiality (80%), while the rest addressed privacy (10%), authentication (8%) and accountability (2%). Concerning the scenario, the Cloud is the most used (53%), although there are examples of wireless networks (8%) and other scenarios (33%). As for the functionality, two thirds of the papers deal with access con-

Table 6: Bibliometric Analysis of PRE Applications

Criteria	Classification	Coverage
Objective	**Confidentiality**	**80%**
	Privacy	10%
	Authentication	8%
	Accountability	2%
Scenarios	**Cloud**	**53%**
	Wireless Network	8%
	Others	33%
Functionality	**Access Control**	**67%**
	Key Management	24%
	Communication	4%

trol (67%), followed by key management (24%) and communications (4%). An immediate conclusion is that, not surprisingly, most of devised PRE applications address the combination of the *confidentiality* objective, the *cloud computing* scenario, and the *access control* functionality. In other words, the typical PRE use case is the secure access delegation scenario, described in Section 2. The next section discusses the application of PRE to this scenario in more detail.

5.1. Secure Access Delegation in the Cloud using PRE

From an abstract point of view, sharing information securely in the cloud can be seen as a problem of access control, where the protected resource is encrypted data. Proxy re-encryption is then used to construct a cryptographically-enforced access control. Ateniese *et al.* proposed in [4] a generic way to do this, as illustrated in Figure 6. The data to be protected is encrypted by the data producer with a fresh symmetric encryption key, the *data key*, which is in turn encrypted under the data owner's public key using the PRE scheme, thus creating an *encrypted lockbox*. This lockbox, which contains the encrypted data and data key, can now be stored in an untrusted repository, such as the cloud. Note also that the production of encrypted data can be performed by any entity that knows the public key of the data owner, supporting this way multiple data sources. For granting access to his data, the data owner generates re-encryption keys for authorized users and hands them to the proxy entity; these can be seen as access delegation tokens created by the data in order to delegate access to certain consumen. When the data consumer requests access to the encrypted data, the proxy re-encrypts the encrypted key. Finally, the consumer decrypts the lockbox with his private key.

This template for cryptographically-enforced access control stems naturally from the main functionality of PRE, which is delegation of access rights to encrypted data. Note that it is not necessary for the data owner to be online. The only requirement is that the cloud is available (which is intrinsic to the cloud paradigm). The data owner can be off-line, while consumers interact directly with the cloud provider in order to access to authorized information. Note that this assumes a honest-but-curious trust model (as discussed in Section 2.1), where the cloud provider may have an incentive for accessing users' data without their permission, but at the same time, it is assumed to behave honestly with respect its functionality.

5.1.1. Review of the literature

In this section we review the state of research on solutions for access delegation in the cloud that use PRE. Most of them share a similar essence, although vary greatly with regards to specific constructions and designs.

In [47], Yu *et al.* propose a system for access delegation in the cloud, using a combination of Key-Policy Attribute-Based Encryption (KP-ABE) and PRE. Data is encrypted using KP-ABE and stored in the cloud, so only users in possession of an appropriate collection of attribute secret keys can decrypt the data. Besides managing encrypted data, the cloud also manages attribute secret keys of the users, except for one special secret key which is required for all decryptions, so the cloud cannot decipher anything. The reason why the cloud manages these secret keys is to handle revocation of users. When a data owner revokes certain users, then new keys must be provided for the remaining users, and encrypted data must be re-encrypted. Both issues are handled by the cloud using PRE, so the data owner simply generates re-encryption keys for transforming not only ciphertexts, but also attribute public and secret keys. This added functionality is possible by carefully integrating the BBS98 scheme with a KP-ABE scheme. For efficiency reasons, the re-encryptions are performed in a "lazy" way, that is, only when an access request from a user is made. [48] propose a modification to Yu *et al.* 's design in order to avoid collusions between the provider and revoked users, but their proposal consists basically in replacing the cloud provider with a trusted third party, which implies relying on stronger trust assumptions. [49, 50, 51] describe similar approaches but for integration with Ciphertext-Policy Attribute-Based Encryption (CP-ABE), where the access structure is associated to the encrypted data rather than to the user attribute key.

Jia *et al.* describe in [52] a classical instantiation of the PRE-based access control template. Their proposal uses the GA07a scheme, where re-encryption keys represent authorization tokens between users. Revocation is handled by the data owner by asking the provider to renew re-encryption keys.

Lin *et al.* propose in [53] a combination of threshold encryption with PRE. This proposal fits the general template of PRE application for access delegation, with the exception being that the proxy entity is distributed among several servers in order to support a decentralized architecture. On an access request, a randomly chosen subset of these servers re-encrypt the data, and since each of these servers store a share of the data owner's secret key, they also perform partial decryptions of the encrypted data. The delegatee user combines the partially decryptions in order to obtain the requested data.

Liu *et al.* propose in [54] a time-constrained access control scheme, combining once again PRE and ABE. In this case, ABE is used for describing time-based access control policies, whereas PRE is used for updating the time attributes. Their proposed system follows the typical template for PRE-based access control in the cloud.

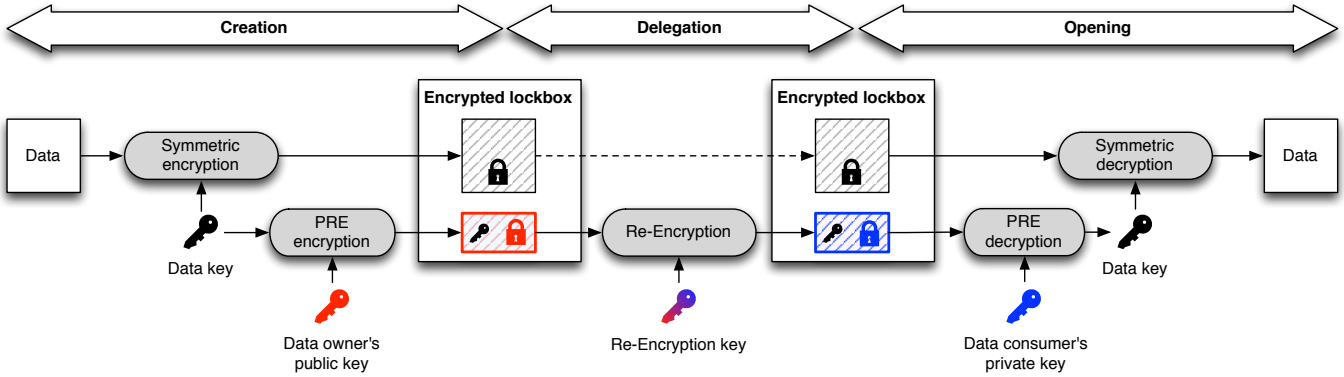

Figure 6: Lifecycle of a PRE-based lockbox

The proposal by Yang *et al.* [55] is different from the previous ones sincce it is a general model that can be instantiated with different PRE schemes. Previous proposals required a specific PRE scheme (most of the times, a variation of the classic BBS98 scheme), carefully tailored for its integration with other primitives and protocols, such as ABE schemes. This general model modifies the encrypted lockbox scheme presented before, including ABE in the following manner. The data key k is split in two different keys k_1 and k_2 by means of a XOR operation, so $k = k_1 \oplus k_2$. Next, PRE is used for encrypting k_1 and ABE for k_2. If a consumer requests access to the encrypted data, then he must possess the necessary ABE secret keys and the cloud must have the corresponding re-encryption key for performing the re-encryption. In that case, then the consumer is able to retrieve k_1 and k_2, and hence, to decrypt the data with the data key k. In this model, fine-grained access control is provided by ABE, whereas PRE makes revocation possible. This model is an example of use of re-encryption keys as authorization tokens, where the presence of re-encryption keys in the cloud provider means that the consumer is authorized to access the data.

Xiong *et al.* present CloudSeal in [56], a cloud-based access delegation system that integrates PRE with a secret sharing scheme. This proposal is slightly similar to the one from Yu *et al.* , but instead of ABE, it uses a secret sharing scheme. The main drawback is that it assumes the existence of secure channels between the data owner and data consumers, which nullifies the necessity of the proposal.

Han *et al.* propose in [57] an identity-based PRE scheme suitable for intra- and inter-domain data sharing. The main feature of this scheme it that it bounds the re-encryption keys not only to the consumerss identity but also to a specific ciphertext (i.e., a shared file). This design choice implies that the data owner must create a different re-encryption key for each pair of data consumer and shared file, but also limits the chances that the cloud provider re-encrypts arbitrary data. Lin *et al.* also propose a similar idea [58], but with a hierarchical PRE instead of identity-based PRE.

Other similar examples are found for more specific applications. For example, the PRE-based access control template is used in [59, 60, 61] for creating an Identity-as-a-Service model where the cloud identity provider cannot access the identity information. In [62], PRE is used for delegating access to encrypted search indexes, in the context of privacy-aware searches. Proxy re-encryption is integrated to the MapReduce paradigm for privacy-preserving Big Data processing in [63, 64]. Other examples are found for service aggregation [65], de-duplication in secure cloud storage [66], and privacy-preserving location-based services [67].

5.1.2. *Findings*

There are common patterns that appear among some of these proposals. For instance, the most prominent is that access delegation is realized by considering re-encryption keys as authorization tokens. Thus, $rk_{A\to B}$ can be seen as an authorization from data owner A to user B. In this case, the enforcement of access rights is naturally realized by means of re-encryption.

When it comes to revocation of access rights, PRE can be used in two ways. The first way is applicable when re-encryption keys are used as authorization tokens between users; in this case, the owner simply instructs the cloud provider to delete the re-encryption keys [4, 55]. Therefore, it is necessary that the cloud provider is trusted enough to ensure that the keys are actually deleted. The second way, used by Yu *et al.* [47], is to re-encrypt the data so it cannot be decrypted by the revoked user. In this case, it is also necessary to trust that the cloud provider is performing the re-encryption and that it is not maintaining a copy of the unre-encrypted data.

Another interesting finding concerns other cryptographic primitives that usually accompany PRE. The main is Attribute-Based Encryption (ABE), which is used in combination with PRE in at least 9 of the 45 papers (that is, a fifth part). Another cryptographic technique commonly used in conjunction with PRE is Threshold Encryption.

5.1.3. *Incentives*

We have seen that PRE realizes access delegation to encrypted information, and, therefore, has a natural application to sharing data securely in the cloud. We believe that this represents a great advance with respect to the state of current cloud services, where users' data is fully controlled by the cloud provider, or, at its best, data is encrypted with keys controlled by the

provider. As discussed in the introduction of this paper, data owners are obliged to trust that the provider will make proper use of his data and will guarantee its protection. Therefore, a PRE-based solution could enable users to outsource their data to cloud providers without necessarily establishing a strong trust relation with them, relying instead in the protection granted by the underlying cryptographic mechanism. At this point, it is worth studying what kind of incentives may motivate cloud providers to implement this kind of solutions, that is, handle data in encrypted form. Note that, in the first place, cloud providers would lose control over the user's data, which is currently a valuable asset. Moreover, providers may incur in more expenses as a result of implementing these additional security mechanisms, However, as pointed out in [59], there are two main incentives that could encourage cloud providers to adopt this solution.

The first incentive is compliance and minimization of liability. An intrinsic characteristic of the cloud is that it enables ubiquitous access to data, which can, potentially, infringe privacy and data protection regulations. Moreover, the cloud can also be used to host and process information that is of problematic nature, such as, illegal or defamatory material [68]. Therefore, cloud providers are affected by specific laws and regulations regarding privacy, data protection and copyright, among others. Although cloud providers currently try to reduce their liability through specific clauses in SLAs, legal responsibility for the data in the cloud also lies on the side of the provider. In contrast, in a PRE-based solution where the outsourced data is encrypted prior to reaching the cloud, liability of the cloud provider is drastically minimized, since it does not hold the decryption keys to read users' data. As a consequence, users should be the ones designated as liable and subject to the enforcement of key disclosure laws.

Data confidentiality as an added value is a second type of incentive. Cloud providers could offer secure data processing and storage as an added value, establishing a competitive advantage over the rest. This characteristic could be an important driver in users' decision process, as there is an increasing interest in this kind of services. In our opinion, there is room in the current cloud service landscape for a business model based on the respect for users' privacy and data confidentiality.

5.1.4. Economic analysis

One of the postulates of this paper is that the application of PRE as a cryptographic mechanism to support secure access delegation in the cloud is not only functional, but practical. A thorough analysis of the practicality of this kind of solution should also determine whether it is economically viable. There is a vast amount of proposals for improving the security, reliability and functionality of cloud services, but, at the same time, there are almost no critical analyses about the economic impact that arises from implementing them. In particular, most cryptographic proposals only provide theoretical analysis of security and computation complexity, but do not study whether their proposals are economically feasible, even when their solutions often imply an intensive use of computation or communication resources.

Table 7: Comparison of re-encryption costs in the cloud

Scheme	Time (ms)	Cost (cents)	#re-enc/cent
BBS98	11.48	4.58E-05	21 844
AFGH06a	83.52	3.33E-04	3 003
WDLC10a	22.29	8.89E-05	1 1250
LV11a	386.9	1.54E-03	648
ABPW13	20.5	8.17E-05	12 233
NAL15a	1.15	4.59E-06	218 063

In [69], Chen and Sion analyze the economic impact of the outsourcing of computation and storage to the cloud. The key contribution of this work is the quantification of the costs associated to this outsourcing, which are driven by several factors, such as hardware, energy and personnel. The authors break down the expenses derived from computation, storage and network services, using the *picocent* (10^{-12} USD cents) as unit of cost. Focusing on computation only, in 2010 the cost was estimated in approximately 2 picocents per CPU cycle. A simple analysis based on the prices of Amazon EC2 shows that these costs have decreased roughly a 25% from 2010 to 2016 (a yearly 4%) Although some factors that influence these costs remain stable (e.g., infrastructures, energy and personnel), the core resource, which is hardware, increases with time its efficiency per unit of cost. Computing resources are the best example, since the cost per CPU cycle has decreased exponentially through history, following Moore's law. For our analysis, we will consider an adjusted computation cost of 1.5 picocents per CPU cycle. Based on this, we estimate the cost per re-encryption operation for several of the PRE schemes analyzed in this survey. A similar economic assessment about the use of proxy re-encryption in a cloud setting is presented in [59]. Recall that the re-encryption is the basic operation that the cloud provider performs in a PRE-based access control solution for the cloud. Table 7 shows these estimations. These figures vary greatly depending on the selected scheme and its properties, but demonstrate that the cloud provider could perform in the order of thousands of re-encryptions per cent.

For illustration purposes, let us assume a scenario similar to that suggested in [59], where a cloud provider implements a PRE-based access control solution that performs a million re-encryptions per day, one for each access request it receives. From the figures of the last table, it can be seen that the total cost of these operations over the course of a year range from $5,631 for the LV11a scheme, to just $17 for the NAL15a scheme. In our opinion, these expenses are reasonable for a cloud provider considering the costs that it could incur in the case of a disclosure or security breach, without taking into account intangible costs such as loss of customers and reputation damage.

5.2. Other Applications of PRE

Although protecting outsourced information in the cloud is a natural application, there are other interesting uses for proxy

re-encryption, where *group key management* is the most remarkable. In this case, PRE can be used for different purposes such as distributing keys, revoking access, performing key escrows, etc. This has multiple applications, such as DRM protection [70, 71, 72, 73], and security in multicast communications [74, 75, 76, 77]. Among alternative uses of PRE we find privacy-preserving solutions for RFID [78, 79, 80], authentication in VANETs [81, 82], location privacy [83], privacy in online social networks [84], anonymity in P2P communication [85], and access control in other scenarios [86, 87].

6. Conclusions and Research Directions

In this paper we study the secure access delegation problem, which occurs naturally in the cloud setting, and postulate that proxy re-encryption is a feasible cryptographic solution to this problem, both from the functional and efficiency perspectives. Proxy re-encryption permits to delegate access to encrypted data, which is of special interest in scenarios where outsourced data must be protected (e.g., the cloud).

We review and analyze the current state of research on PRE, for both constructions and applications. The most prominent proxy re-encryption schemes are studied in the light of relevant properties and security models. We additionally provide a comparative analysis of the performance of selected schemes, both from the theoretical and experimental points of view. With regards to the applications, we perform an extensive analysis of the available literature following a bibliometric approach. As a result, we conclude that secure access delegation in the cloud is currently the use case with the most potential for PRE, given the functionalities and performance levels that it provides, and the regulatory and economic incentives. Moreover, it is more realistic than other solutions, such as those based on homomorphic encryption.

In addition to the analysis of PRE schemes and applications, we identify three main research directions, each of them located at a different abstraction level. The first line is focused on the fundamentals of security, and in particular, on defining unified security notions and properly understanding their interconnections. The second line shows a less-explored – although increasingly targeted – area of research, which is the construction of proxy re-encryption schemes with lattice-based techniques. Finally, the third identified line is the application of PRE to real-world solutions, and therefore, it is oriented to the practitioner. We describe next these lines in more detail.

6.1. Security Definitions for Proxy Re-Encryption

There are several challenges that are related with the very definition of what is considered secure in PRE, from the perspective of provable security. In order to reason about the security of a cryptosystem, it becomes essential to have a proper definition and understanding of the notions that model its security. However, in the PRE literature, security notions are often ad-hoc, presenting subtle differences and constraints, which results in the lack of a unified framework of security definitions. A reason for this is the many possible combinations of PRE properties (unidirectional vs. bidirectional, single-use vs. multi-use, interactive vs. non-interactive, etc), which affect the relations between security notions and the necessary restrictions. It is therefore a challenge to achieve definitional unity for PRE security. To this respect, we also note that there are few attempts focused on analyzing PRE security notions (e.g., [3, 18, 14]), which reinforces our view that this a worthwhile area of research. The work of Nuñez *et al.* in [14] is an initial effort to this respect, identifying some general implications and separations between different PRE security notions; however, more of these relations are to be found. Another open research line is the analysis of security of PRE schemes using a simulation-based approach, which was first proposed by Canetti and Hohenbeger in [3] but has received little attention. Other topics of interest are the modelization of adversaries and their restrictions, such as in [18], as well as investigation the relation between PRE properties and security notions [14].

6.2. Lattice-based Proxy Re-Encryption Schemes

The vast majority of PRE schemes to date are based on traditional number-theoretic foundations, usually groups (either pairing-based or not) where the Discrete Logarithm is hard. This can be problematic if efficient cryptanalytic methods against this problem are found or quantum computers are available. Lattice-based cryptography is a promising alternative due to its potential for post-quantum security, and in some cases, its performance. With respect to the latter, it is worth mentioning that the experimental results presented in Section 4.3.2 show a very interesting trend. Lattice-based PRE schemes seem to be more efficient in comparison to previous constructions. However, at the same time, the highest security notion attempted so far by lattice-based PRE was CCA1 by the Kir14 scheme, although as discussed in Section 4.1.12, it only achieves CPA-security. This clearly represents a current open challenge. Note also that it is difficult to compare the performance of schemes that have very different underlying structures. An obvious reason is that the presumed levels of security depend, essentially, on the effectiveness of known attacks; lattice-based cryptography is a fairly young area, so potential improvements on attack strategies, or even new kinds of attacks, are to be expected.

6.3. Applications to the real world

All the applications we reviewed in Section 5 are mere scientific proposals. It is also worth finding out whether there are industrial initiatives towards the application of PRE in real commercial solutions. Although we are not presently aware of widely-used real-world applications of PRE, there are some indications that show that the IT industry is becoming increasingly interested in it, particularly with regard to its application to the secure access delegation scenario. Nishimaki and Xagawa recently noted in [36] that Toshiba has released, on the Japanese market, a cloud storage service that use proxy re-encryption to ensure data confidentiality, called "Digital Kashikinko" (which means "digital safety box"). To the best of our knowledge, this is the first commercial application of PRE. Although the specific details are not known (as for most commercial applications), the description in [88] shows a typical PRE-based

architecture, like the one discussed in this paper. Other emerging examples are the encrypted database ZeroDB [89], which uses PRE to share access keys between users, and the EU H2020 research project CREDENTIAL [90], which targets applications for the eHealth, eBusiness, and eGovernment domains using PRE as a basic cryptographic primitive.

An interesting way of probing the industry's interest in the topic is to look for patents involving PRE. As conjectured in this paper, the application of PRE to the cloud seems to be recurring, given its natural potential in this scenario. We have found patents from Toshiba [91], Huawei [92, 93], Nokia [94], and Gemalto [95], among others, which present approaches that are, essentially, instances of the secure access delegation scenario. With regard to other use cases, Apple has patented a solution for DRM protection based in PRE [96]. The authors of one of the first PRE schemes [4] also presented a patent [97] that essentially covers their original proposal. We foresee a growing number of patents in these areas, which would ultimately lead to new products and services based on this fascinating cryptosystem.

Acknowledgements

This work was partly supported by the Junta de Andalucía through the project FISICCO (P11-TIC-07223) and by the Spanish Ministry of Economy and Competitiveness through the PERSIST project (TIN2013-41739-R).

References

[1] Raluca Ada Popa. *Building practical systems that compute on encrypted data*. PhD thesis, Massachusetts Institute of Technology, 2014.

[2] Jun Shao. Bibliography on proxy re-cryptography. `http://ndc.zjgsu.edu.cn/~jshao/prcbib.htm`, 2015.

[3] R. Canetti and S. Hohenberger. Chosen-ciphertext secure proxy re-encryption. In *Proceedings of the 14th ACM conference on Computer and communications security*, pages 185–194. ACM, 2007.

[4] G. Ateniese, K. Fu, M. Green, and S. Hohenberger. Improved proxy re-encryption schemes with applications to secure distributed storage. *ACM Transactions on Information and System Security*, 9(1):1–30, 2006.

[5] M. Blaze, G. Bleumer, and M. Strauss. Divertible protocols and atomic proxy cryptography. *Advances in Cryptology—EUROCRYPT'98*, pages 127–144, 1998.

[6] Keita Xagawa and Keisuke Tanaka. Proxy re-encryption based on learning with errors. In *Proceedings of the 2010 Symposium on Cryptography and Information Security (SCIS 2010)*, 2010.

[7] Yoshinori Aono, Xavier Boyen, Le Trieu Phong, and Lihua Wang. Key-private proxy re-encryption under LWE. In *Progress in Cryptology–INDOCRYPT 2013*, pages 1–18. Springer, 2013.

[8] Elena Kirshanova. Proxy re-encryption from lattices. In *Public-Key Cryptography–PKC 2014*, pages 77–94. Springer, 2014.

[9] David Nuñez, Isaac Agudo, and Javier Lopez. NTRUReEncrypt: An efficient proxy re-encryption scheme based on NTRU. In *Proceedings of the 10th ACM Symposium on Information, Computer and Communications Security*, ASIA CCS '15, pages 179–189, New York, NY, USA, 2015. ACM.

[10] C.K. Chu and W.G. Tzeng. Identity-based proxy re-encryption without random oracles. *Information Security*, pages 189–202, 2007.

[11] R.H. Deng, J. Weng, S. Liu, and K. Chen. Chosen-ciphertext secure proxy re-encryption without pairings. In *Proceedings of the 7th International Conference on Cryptology and Network Security*, pages 1–17. Springer-Verlag, 2008.

[12] B. Libert and D. Vergnaud. Unidirectional chosen-ciphertext secure proxy re-encryption. *Information Theory, IEEE Transactions on*, 57(3):1786–1802, 2011.

[13] M. Green and G. Ateniese. Identity-based proxy re-encryption. In *Applied Cryptography and Network Security*, pages 288–306. Springer, 2007.

[14] David Nuñez, Isaac Agudo, and Javier Lopez. A parametric family of attack models for proxy re-encryption. In *Proceedings of the 28th IEEE Computer Security Foundations Symposium*, CSF'15, pages 290–301. IEEE Computer Society, 2015.

[15] Mihir Bellare, Anand Desai, David Pointcheval, and Phillip Rogaway. Relations among notions of security for public-key encryption schemes. In *Advances in Cryptology—CRYPTO'98*, pages 26–45. Springer, 1998.

[16] Ran Canetti, Shai Halevi, and Jonathan Katz. Chosen-ciphertext security from identity-based encryption. In *Advances in Cryptology-Eurocrypt 2004*, pages 207–222. Springer, 2004.

[17] Sébastien Canard, Julien Devigne, and Fabien Laguillaumie. Improving the security of an efficient unidirectional proxy re-encryption scheme. *Jounal of Internet Services and Information Security*, pages pp140–160, 2011.

[18] Jiang Zhang, Zhenfeng Zhang, and Yu Chen. PRE: Stronger security notions and efficient construction with non-interactive opening. *Theoretical Computer Science*, 542:1–16, 2014.

[19] Jian Weng, Robert H Deng, Shengli Liu, and Kefei Chen. Chosen-ciphertext secure bidirectional proxy re-encryption schemes without pairings. *Information Sciences*, 180(24):5077–5089, 2010.

[20] Jian Weng, Robert H Deng, Xuhua Ding, Cheng-Kang Chu, and Junzuo Lai. Conditional proxy re-encryption secure against chosen-ciphertext attack. In *Proceedings of the 4th International Symposium on Information, Computer, and Communications Security*, pages 322–332. ACM, 2009.

[21] David Nuñez, Isaac Agudo, and Javier Lopez. On the application of generic CCA-secure transformations to proxy re-encryption. *Security and Communication Networks*, 2016.

[22] Masahiro Mambo and Eiji Okamoto. Proxy cryptosystems: Delegation of the power to decrypt ciphertexts. *IEICE transactions on fundamentals of electronics, Communications and computer sciences*, 80(1):54–63, 1997.

[23] The JHU-MIT Proxy Re-cryptography Library. `http://spar.isi.jhu.edu/~mgreen/prl/`.

[24] Dan Boneh and Matt Franklin. Identity-based encryption from the weil pairing. In *Advances in Cryptology—CRYPTO 2001*, pages 213–229. Springer, 2001.

[25] Woo Kwon Koo, Jung Yeon Hwang, and Dong Hoon Lee. Security vulnerability in a non-interactive ID-based proxy re-encryption scheme. *Information Processing Letters*, 109(23):1260–1262, 2009.

[26] Craig Gentry and Alice Silverberg. Hierarchical id-based cryptography. In *Advances in cryptology—ASIACRYPT 2002*, pages 548–566. Springer, 2002.

[27] Brent Waters. Efficient identity-based encryption without random oracles. In *Advances in Cryptology–EUROCRYPT 2005*, pages 114–127. Springer, 2005.

[28] J. Shao and Z. Cao. CCA-secure proxy re-encryption without pairings. In *Proceedings of the 12th International Conference on Practice and Theory in Public Key Cryptography: PKC'09*, pages 357–376. Springer-Verlag, 2009.

[29] Toshihiko Matsuo. Proxy re-encryption systems for identity-based encryption. In *Pairing-Based Cryptography–Pairing 2007*, pages 247–267. Springer, 2007.

[30] G. Ateniese, K. Benson, and S. Hohenberger. Key-private proxy re-encryption. *Topics in Cryptology–CT-RSA 2009*, pages 279–294, 2009.

[31] Mihir Bellare, Alexandra Boldyreva, Anand Desai, and David Pointcheval. Key-privacy in public-key encryption. In *Advances in Cryptology—ASIACRYPT 2001*, pages 566–582. Springer, 2001.

[32] Claus-Peter Schnorr. Efficient signature generation by smart cards. *Journal of cryptology*, 4(3):161–174, 1991.

[33] J. W. Seo, D. H. Yum, and P. J. Lee. Comments on "unidirectional chosen-ciphertext secure proxy re-encryption". *Information Theory, IEEE Transactions on*, PP(99):1, 2012.

[34] Richard Lindner and Chris Peikert. Better key sizes (and attacks) for LWE-based encryption. In *Topics in Cryptology–CT-RSA 2011*, pages 319–339. Springer, 2011.

[35] Kunwar Singh, C Pandu Rangan, and AK Banerjee. Cryptanalysis of unidirectional proxy re-encryption scheme. In *Information and Commu-*

nication Technology, pages 564–575. Springer, 2014.

[36] Ryo Nishimaki and Keita Xagawa. Key-private proxy re-encryption from lattices, revisited. *IEICE TRANSACTIONS on Fundamentals of Electronics, Communications and Computer Sciences*, 98(1):100–116, 2015.

[37] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *Advances in Cryptology–EUROCRYPT 2012*, pages 700–718. Springer, 2012.

[38] Xiong Fan and Feng-Hao Liu. Various proxy re-encryption schemes from lattices. Cryptology ePrint Archive, Report 2016/278, 2016. `http://eprint.iacr.org/2016/278`.

[39] Jeffrey Hoffstein, Jill Pipher, and Joseph H Silverman. NTRU: A ring-based public key cryptosystem. In *Algorithmic number theory*, pages 267–288. Springer, 1998.

[40] Damien Stehlé and Ron Steinfeld. Making NTRU as secure as worst-case problems over ideal lattices. In *Advances in Cryptology–EUROCRYPT 2011*, pages 27–47. Springer, 2011.

[41] E. Barker, L. Chen, A. Roginsky, and M. Smid. Recommendation for Pair-Wise Key Establishment Schemes Using Discrete Logarithm Cryptography. NIST special publication 800-56A (Revision 2), NIST, May 2013.

[42] Angelo De Caro and Vincenzo Iovino. jPBC: Java pairing based cryptography. In *Proceedings of the 16th IEEE Symposium on Computers and Communications, ISCC 2011*, pages 850–855. IEEE, 2011.

[43] Steven D Galbraith, Kenneth G Paterson, and Nigel P Smart. Pairings for cryptographers. *Discrete Applied Mathematics*, 156(16):3113–3121, 2008.

[44] Elaine Barker, William Barker, William Burr, William Polk, and Miles Smid. Recommendation for key management — part 1: General. Technical report, 2005.

[45] Java implementation of NTRUReEncrypt. `https://github.com/cygnusv/ntrureencrypt`.

[46] W Whyte, N Howgrave-Graham, J Hoffstein, J Pipher, JH Silverman, and P Hirschhorn. IEEE P1363.1: Draft standard for public-key cryptographic techniques based on hard problems over lattices. Technical report, IEEE, 2008.

[47] Shucheng Yu, Cong Wang, Kui Ren, and Wenjing Lou. Achieving secure, scalable, and fine-grained data access control in cloud computing. In *INFOCOM, 2010 Proceedings IEEE*, pages 1–9. Ieee, 2010.

[48] Namje Park. Secure data access control scheme using type-based re-encryption in cloud environment. In *Semantic Methods for Knowledge Management and Communication*, pages 319–327. Springer, 2011.

[49] Guojun Wang, Qin Liu, Jie Wu, and Minyi Guo. Hierarchical attribute-based encryption and scalable user revocation for sharing data in cloud servers. *computers & security*, 30(5):320–331, 2011.

[50] Junbeom Hur. Improving security and efficiency in attribute-based data sharing. *Knowledge and Data Engineering, IEEE Transactions on*, 25(10):2271–2282, 2013.

[51] Piotr K Tysowski and M Anwarul Hasan. Hybrid attribute-and re-encryption-based key management for secure and scalable mobile applications in clouds. *Cloud Computing, IEEE Transactions on*, 1(2):172–186, 2013.

[52] Weiwei Jia, Haojin Zhu, Zhenfu Cao, Lifei Wei, and Xiaodong Lin. Sdsm: a secure data service mechanism in mobile cloud computing. In *Computer Communications Workshops (INFOCOM WKSHPS), 2011 IEEE Conference on*, pages 1060–1065. IEEE, 2011.

[53] Hsiao-Ying Lin and W-G Tzeng. A secure erasure code-based cloud storage system with secure data forwarding. *Parallel and Distributed Systems, IEEE Transactions on*, 23(6):995–1003, 2012.

[54] Qin Liu, Guojun Wang, and Jie Wu. Time-based proxy re-encryption scheme for secure data sharing in a cloud environment. *Information Sciences*, 258:355–370, 2014.

[55] Yanjiang Yang and Youcheng Zhang. A generic scheme for secure data sharing in cloud. In *Parallel Processing Workshops (ICPPW), 2011 40th International Conference on*, pages 145–153. IEEE, 2011.

[56] Huijun Xiong, Xinwen Zhang, Danfeng Yao, Xiaoxin Wu, and Yonggang Wen. Towards end-to-end secure content storage and delivery with public cloud. In *Proceedings of the second ACM conference on Data and Application Security and Privacy*, pages 257–266. ACM, 2012.

[57] Jinguang Han, Willy Susilo, and Yi Mu. Identity-based data storage in cloud computing. *Future Generation Computer Systems*, 29(3):673–681, 2013.

[58] Hsiao-Ying Lin, John Kubiatowicz, and Wen-Guey Tzeng. A secure fine-grained access control mechanism for networked storage systems. In *Software Security and Reliability (SERE), 2012 IEEE Sixth International Conference on*, pages 225–234. IEEE, 2012.

[59] David Nuñez and Isaac Agudo. BlindIdM: A privacy-preserving approach for identity management as a service. *International Journal of Information Security*, 13(2):199–215, 2014.

[60] Bernd Zwattendorfer, Daniel Slamanig, Klaus Stranacher, and Felix Hörandner. A federated cloud identity broker-model for enhanced privacy via proxy re-encryption. In *Communications and Multimedia Security*, pages 92–103. Springer, 2014.

[61] David Nuñez, Isaac Agudo, and Javier Lopez. Integrating openid with proxy re-encryption to enhance privacy in cloud-based identity services. In *Proceedings of the 2012 IEEE 4th International Conference on Cloud Computing Technology and Science (CloudCom)*, pages 241 – 248, Taipei, Taiwan, Dec 2012 2012. IEEE Computer Society, IEEE Computer Society.

[62] Zeeshan Pervez, Ammar Ahmad Awan, Asad Masood Khattak, Sungyoung Lee, and Eui-Nam Huh. Privacy-aware searching with oblivious term matching for cloud storage. *The Journal of Supercomputing*, 63(2):538–560, 2013.

[63] Lei Xu, Weidong Shi, and Taeweon Suh. Pfc: Privacy preserving fpga cloud-a case study of mapreduce. In *Cloud Computing (CLOUD), 2014 IEEE 7th International Conference on*, pages 280–287. IEEE, 2014.

[64] David Nuñez, Isaac Agudo, and Javier Lopez. Delegated access for hadoop clusters in the cloud. In *Proceedings of the 2014 IEEE 6th International Conference on Cloud Computing Technology and Science (CloudCom)*, pages 374–379, Dec 2014.

[65] Peishun Wang, Yi Mu, Willy Susilo, and Jun Yan. Privacy preserving protocol for service aggregation in cloud computing. *Software: Practice and Experience*, 42(4):467–483, 2012.

[66] Chuanyi Liu, Xiaojian Liu, and Lei Wan. Policy-based de-duplication in secure cloud storage. In *Trustworthy Computing and Services*, pages 250–262. Springer, 2013.

[67] Jun Shao, Rongxing Lu, and Xiaodong Lin. Fine: A fine-grained privacy-preserving location-based service framework for mobile devices. In *INFOCOM, 2014 Proceedings IEEE*, pages 244–252. IEEE, 2014.

[68] Rolf H Weber and Dominic Nicolaj Staiger. Cloud computing: a cluster of complex liability issues. *Web Journal of Current Legal Issues*, 20(1), 2014.

[69] Y. Chen and R. Sion. On securing untrusted clouds with cryptography. In *Proc. 9th annual ACM workshop on Privacy in the electronic society*, pages 109–114. ACM, 2010.

[70] Gelareh Taban, Alvaro A Cárdenas, and Virgil D Gligor. Towards a secure and interoperable DRM architecture. In *Proceedings of the ACM workshop on Digital rights management*, pages 69–78. ACM, 2006.

[71] Sangho Lee, Heejin Park, and Jong Kim. A secure and mutual-profitable drm interoperability scheme. In *Computers and Communications (ISCC), 2010 IEEE Symposium on*, pages 75–80. IEEE, 2010.

[72] Qin Qiu, Zhi Tang, and Yinyan Yu. A decentralized authorization scheme for drm in p2p file-sharing systems. In *Consumer Communications and Networking Conference (CCNC), 2011 IEEE*, pages 136–140. IEEE, 2011.

[73] Nakul Joshi and Ronald Petrlic. Towards practical privacy-preserving digital rights management for cloud computing. In *Consumer Communications and Networking Conference (CCNC), 2013 IEEE*, pages 265–270. IEEE, 2013.

[74] Yun-Peng Chiu, Chin-Laung Lei, and Chun-Ying Huang. Secure multicast using proxy encryption. In *Information and Communications Security*, pages 280–290. Springer, 2005.

[75] Ritesh Mukherjee and J William Atwood. Scalable solutions for secure group communications. *Computer Networks*, 51(12):3525–3548, 2007.

[76] Chun-Ying Huang, Yun-Peng Chiu, Kuan-Ta Chen, and Chin-Laung Lei. Secure multicast in dynamic environments. *Computer Networks*, 51(10):2805–2817, 2007.

[77] Yiliang Han, Xiaolin Gui, Xuguang Wu, and Xiaoyuan Yang. Proxy encryption based secure multicast in wireless mesh networks. *Journal of network and computer applications*, 34(2):469–477, 2011.

[78] Thomas S Heydt-Benjamin, Hee-Jin Chae, Benessa Defend, and Kevin Fu. Privacy for public transportation. In *Privacy Enhancing Technologies*, pages 1–19. Springer, 2006.

[79] Florian Kerschbaum and Alessandro Sorniotti. Rfid-based supply chain partner authentication and key agreement. In *Proceedings of the second ACM conference on Wireless network security*, pages 41–50. ACM, 2009.

[80] Qiang Yan, Yingjiu Li, Robert H Deng, et al. Anti-tracking in rfid discovery service for dynamic supply chain systems. *International Journal of RFID Security and Cryptography (IJRFIDSC)*, 1(1/2):25–35, 2012.

[81] Jun Liu, Xiaoyan Hong, Qunwei Zheng, and Lei Tang. Privacy-preserving quick authentication in fast roaming networks. In *Local Computer Networks, Proceedings 2006 31st IEEE Conference on*, pages 975–982. IEEE, 2006.

[82] Tat Wing Chim, Siu-Ming Yiu, Lucas CK Hui, and Victor OK Li. Mlas: Multiple level authentication scheme for vanets. *Ad Hoc Networks*, 10(7):1445–1456, 2012.

[83] Changyu Dong and Naranker Dulay. Longitude: a privacy-preserving location sharing protocol for mobile applications. In *Trust Management V*, pages 133–148. Springer, 2011.

[84] Matthew M Lucas and Nikita Borisov. Flybynight: mitigating the privacy risks of social networking. In *Proceedings of the 7th ACM workshop on Privacy in the electronic society*, pages 1–8. ACM, 2008.

[85] Junzhou Luo, Xiaogang Wang, and Ming Yang. A resilient p2p anonymous routing approach employing collaboration scheme. *J. UCS*, 15(9):1797–1811, 2009.

[86] Mihaela Ion, Giovanni Russello, and Bruno Crispo. Design and implementation of a confidentiality and access control solution for publish/subscribe systems. *Computer networks*, 56(7):2014–2037, 2012.

[87] Hwajeong Seo and Howon Kim. Zigbee security for visitors in home automation using attribute based proxy re-encryption. In *Consumer Electronics (ISCE), 2011 IEEE 15th International Symposium on*, pages 304–307. IEEE, 2011.

[88] Masaaki Miki, Eiji Hayashi, and Hideki Shingai. Highly reliable and highly secure online storage platform supporting "timeon" regza cloud service. *Toshiba Review*, 68(5):25–27, 2013. In Japanese.

[89] Michael Egorov and MacLane Wilkison. Zerodb white paper. *CoRR*, abs/1602.07168, 2016.

[90] Felix Hörandner, Stephan Krenn, Andrea Migliavacca, Florian Thiemer, and Bernd Zwattendorfer. CREDENTIAL: A Framework for Privacy-Preserving Cloud-Based Data Sharing. In *SECPID@ARES 2016*, pages 742–749. IEEE, 2016.

[91] M. Shimano, G. Fujino, M. Miki, Y. Tsuzuki, I. TAKEYASU, and E. TOKITA. Key change management apparatus and key change management method, April 3 2014. US Patent App. 13/928,112.

[92] L. Xu and X. Wu. Proxy-based encryption method, proxy-based decryption method, network equipment, network device and system, October 28 2014. US Patent 8,873,754.

[93] G. Zhang. Data protection method and data protection system, September 17 2013. US Patent 8,539,606.

[94] D. BISWAS. Methods and apparatus for sharing real-time user context information, February 4 2014. US Patent 8,645,682.

[95] P. Paillier and A. Gouget. Data providing process based on an ibpe scheme, February 4 2014. US Patent 8,644,509.

[96] A.J. Farrugia, N. Sullivan, G. Fasoli, and M. Ciet. Encryption method and apparatus using composition of ciphers, March 25 2014. US Patent 8,681,975.

[97] S.R. Hohenberger, K. Fu, G. Ateniese, and M. Green. Unidirectional proxy re-encryption, January 10 2012. US Patent 8,094,810.