Adaptively Secure Proxy Re-encryption

Georg Fuchsbauer[1], Chethan Kamath[2], Karen Klein[2], and Krzysztof Pietrzak[2]

[1] Inria and ENS Paris
georg.fuchsbauer@ens.fr
[2] IST Austria
{ckamath,karen.klein,pietrzak}@ist.ac.at

Abstract. A proxy re-encryption (PRE) scheme is a public-key encryption scheme that allows the holder of a key pk to derive a re-encryption key for any other key pk'. This re-encryption key lets anyone transform ciphertexts under pk into ciphertexts under pk' without having to know the underlying message, while transformations from pk' to pk should not be possible (unidirectional). Security is defined in a multi-user setting against an adversary that gets the users' public keys and can ask for re-encryption keys and can corrupt users by requesting their secret keys. Any ciphertext that the adversary cannot trivially decrypt given the obtained secret and re-encryption keys should be secure.

All existing security proofs for PRE only show *selective* security, where the adversary must first declare the users it wants to corrupt. This can be lifted to more meaningful *adaptive* security by guessing the set of corrupted users among the n users, which loses a factor exponential in n, rendering the result meaningless already for moderate n.

Jafargholi et al. (CRYPTO'17) proposed a framework that in some cases allows to give adaptive security proofs for schemes which were previously only known to be selectively secure, while avoiding the exponential loss that results from guessing the adaptive choices made by an adversary. We apply their framework to PREs that satisfy some natural additional properties. Concretely, we give a more fine-grained reduction for several unidirectional PREs, proving adaptive security at a much smaller loss. The loss depends on the graph of users whose edges represent the re-encryption keys queried by the adversary. For trees and chains the loss is quasi-polynomial in the size and for general graphs it is exponential in their depth and indegree (instead of their size as for previous reductions). Fortunately, trees and low-depth graphs cover many, if not most, interesting applications.

Our results apply e.g. to the bilinear-map based PRE schemes by Ateniese et al. (NDSS'05 and CT-RSA'09), Gentry's FHE-based scheme (STOC'09) and the LWE-based scheme by Chandran et al. (PKC'14).

Keywords: Proxy reencryption · adaptive security · tightness

The full version of this paper can be found on the IACR eprint archive: [FKKP18].

1 Introduction

A proxy re-encryption (PRE) scheme is a public-key encryption scheme with an additional functionality: Alice and Bob, who have key pairs $(\mathtt{pk}_A, \mathtt{sk}_A)$ and $(\mathtt{pk}_B, \mathtt{sk}_B)$, respectively, can generate a re-encryption key (re-key, for short) $\mathtt{rk}_{A,B}$ that allows its holder, say Peggy, to act as a proxy; that is, she can transform ciphertexts under $\mathtt{pk}_A$ to ciphertexts under $\mathtt{pk}_B$ without having to know the underlying message. A trivial way to accomplish this would be for Alice to hand her secret key $\mathtt{sk}_A$ to Peggy, who can then decrypt ciphertexts under $\mathtt{pk}_A$, encrypt them under $\mathtt{pk}_B$ and send them to Bob. Alice's secret key acts thus as the re-key and de- and encryption algorithms are used for re-encryption. However, this approach requires Alice to reveal her secret key to Peggy and therefore place complete trust on her. The more interesting cases are when the parties are mutually distrustful.

Bidirectional vs. unidirectional. In the above setting, if the re-key $\mathtt{rk}_{A,B}$ allows Peggy to also transform ciphertexts under $\mathtt{pk}_B$ to $\mathtt{pk}_A$, the PRE scheme is called "bidirectional". For such schemes the re-key is necessarily a function of both $\mathtt{sk}_A$ and $\mathtt{sk}_B$. In this paper we are interested in the more interesting case of "unidirectional" PRE schemes where the re-key $\mathtt{rk}_{A,B}$ can only transform ciphertexts from $\mathtt{pk}_A$ to $\mathtt{pk}_B$, and not vice-versa, and ciphertexts under $\mathtt{pk}_B$ remain secure even given $\mathtt{sk}_A$ and $\mathtt{rk}_{A,B}$. (Henceforth we will always assume PREs to be unidirectional.) As opposed to bidirectional PREs, the re-key generation algorithm in a unidirectional PRE takes as input "source" keys $(\mathtt{pk}_A, \mathtt{sk}_A)$ and only the "target" public key $\mathtt{pk}_B$.

Single hop vs. multiple hops. Suppose a third user, Charlie, holding keys $(\mathtt{pk}_C, \mathtt{sk}_C)$, enters the picture and suppose Peggy obtains the re-key $\mathtt{rk}_{B,C}$ that allows her to transform ciphertexts under Bob's public key to ciphertexts under Charlie's public key. Peggy can, by definition, transform a ciphertext c_A under $\mathtt{pk}_A$ to a ciphertext c_B under $\mathtt{pk}_B$ using her re-key $\mathtt{rk}_{A,B}$. If it allows Peggy to transform ciphertext c_B, which has already been re-encrypted once, to a ciphertext c_C under $\mathtt{pk}_C$ using the re-key $\mathtt{rk}_{B,C}$ then we say that the PRE scheme allows two "hops". In a similar manner, one can consider multiple hops of re-encryptions. Such a scheme is termed "multi-hop" as opposed to a "single-hop" scheme (which does not allow re-encryptions of already re-encrypted ciphertexts).

1.1 Modelling Security

The basic notion of security for unidirectional PREs is that of indistinguishability under chosen-plaintext attack (CPA). There are n users and, at the beginning of the game, the adversary gets their public keys $\mathtt{pk}_1, \ldots, \mathtt{pk}_n$ from the challenger. In the first phase, the adversary can corrupt users of its choice by requesting their secret keys; in the second phase, it can obtain re-keys $\mathtt{rk}_{i,j}$ and re-encryptions for ciphertexts of its choice. The scheme is CPA-secure if it is infeasible for

the adversary to distinguish encryptions of two messages under a key that the adversary has not corrupted either directly or indirectly (through a re-key or re-encryption query to a corrupted user).

Just as in standard public-key encryption, the above security definition can be strengthened to chosen-ciphertext attack (CCA) by allowing the adversary access to a decryption oracle which, on input a ciphertext and a public key $\mathtt{pk}_i$ returns the decryption of the ciphertext under $\mathtt{sk}_i$. The conditions to ensure non-triviality have to be altered accordingly.

We note that both definitions are *selective* in nature: the adversary must choose the set of players it corrupts before issuing any queries.

1.2 Prior Work

Bidirectional PREs were introduced as "atomic proxy cryptography" by Blaze, Bleumer and Strauss [BBS98], who constructed a multi-hop scheme under the decisional Diffie-Hellman assumption. Unidirectional PREs were introduced later by Ateniese et al. [AFGH05]. Their main motivation was to limit the amount of trust placed on the proxy, as required by their application to access control for distributed storage. Since the notion of security for unidirectional PRE is different from a bidirectional PRE, they also reformulated the notion of CPA (for the single-hop setting). Assuming hardness of certain problems on bilinear groups, they constructed CPA-secure schemes that are single-hop and unidirectional.

The definition of CCA security for single-hop bidirectional schemes is due to Canetti and Hohenberger [CH07] and is more involved than previous definitions, mainly because the adversary is allowed adaptive corruption. They gave a scheme satisfying their notion under the standard decisional bilinear Diffie-Hellman assumption. The definition of CCA security in the unidirectional setting is due to Libert and Vergnaud [LV08], who instantiate it under a slightly non-standard assumption on bilinear groups.

The earlier constructions of multi-hop, unidirectional schemes were based on program obfuscation [HRsV07, CCV12]. In his seminal paper, Gentry [Gen09] gave a generic construction of PREs from fully homomorphic encryption. The first construction (with succinct ciphertexts) based on a standard assumption is due to Chandran et al. [CCL+14]: their scheme is CPA-secure assuming decisional learning with errors. Phong et al. [PWA+16] followed up with a construction that, in addition, enjoys a security property called "key-privacy". The only construction of a CCA-secure multi-hop, unidirectional scheme is due to Fan and Liu [FL17]. In their paper, they also rigorously defined the security models (CPA and CCA) for the multi-hop setting.

Cohen [Coh17] has recently argued that CPA security might be too weak for some applications and introduced *indistinguishability against honest-reencryption attack* (HRA), a notion that lies between CPA and CCA. He also showed that if a PRE satisfies a property called "source-hiding", which several existing CPA-secure schemes do, then HRA security reduces to CPA security.

1.3 Our Contribution

Our starting point is the observation that, unlike bidirectional PREs, the security definitions for unidirectional PREs (that is, CPA, HRA and CCA) are all *selective* in nature: the adversary must choose the set of parties it corrupts before issuing any queries. A more meaningful notion would be *adaptive* security, where the adversary is allowed to corrupt users at any time during the game. However, modelling this turns out to be as tricky as in the bidirectional setting. In this paper, we lift the definitions for CPA and HRA to the adaptive setting.

1.3.1 First Contribution: Modelling Adaptive Corruption. The main problem that arises when we allow the adversary to adaptively corrupt users is that we must ensure that the adversary cannot trivially win the security game. For bidirectional PREs this was handled in [CH07] by defining a relation that keeps track of the dependency between the re-keys and re-encryptions that were issued during the game. Our approach is similar in spirit: the security game maintains a "recoding graph" that has n nodes, and whose edges are derived from the re-keys and re-encryptions issued to the adversary. The exact definitions of the recoding graph for adaptive CPA and for adaptive HRA differ slightly, but in both cases it is defined so that no corrupt key is reachable from the challenge key. That is, the adversary is forbidden from making *any* re-key or re-encryption queries to a corrupt user that is reachable from the challenge key. The recoding graph now allows to ensure non-triviality of the adversary's actions by checking a few basic graph properties.

1.3.2 Second Contribution: The Reduction. Proving adaptive security can be reduced to showing selective security by initially guessing the set of users that will be corrupted. However, this reduction loses an exponential factor in n, rendering the reduction meaningless already for moderate n. As our main contribution, we give a more fine-grained reduction from adaptive to selective security which in many practical settings and for several existing schemes (or minor variants) implies adaptive security at much smaller (quasi-polynomial, or even polynomial) loss. More precisely, the loss in our reduction depends on the structure of the recoding graph: for trees and chains we get a quasi-polynomial $n^{O(\log n)}$ loss, whereas for general graphs the loss is exponential in their depth. Fortunately, trees, chains, and low-depth graphs cover many, if not most, interesting applications.

Security assumptions. A key step in our search for a tighter reduction was the identification of the basic security assumptions on a PRE that we required in our arguments. For the case of CPA, it turned out to be ciphertext indistinguishability and *weak* key-privacy, both fairly standard security requirements already explored in some of the previous works.

As the name suggests, a PRE is ciphertext-indistinguishable (or, for short, indistinguishable) if the underlying encryption is. Since the syntax of the encryption algorithm for a PRE is slightly different from that of a standard public-key

encryption, the definition of indistinguishability has to be slightly changed. To be precise, the encryption algorithm for a PRE takes also a "level" as input, and we require that the ciphertexts are indistinguishable *on all levels*. It is not hard, therefore, to see that any selectively CPA-secure PRE has to trivially satisfy indistinguishability.

The notion of key-privacy was introduced in a strong form in [ABH09]. We require the PRE to satisfy a much weaker property, namely that a re-key $\mathtt{rk}_{A,B}$ looks pseudorandom given just the source and target public keys $\mathtt{pk}_A$ and $\mathtt{pk}_B$. Existing PRE schemes that satisfy the stronger key privacy as defined in [ABH09] are therefore candidates for our reduction.

To apply our reduction to HRA-secure PRE, we need a third assumption to hold: source-hiding. This is the same property that allowed Cohen [Coh17] to lift a CPA-secure PRE to a HRA-secure one. Informally, a PRE is source-hiding if ciphertexts that result from re-encryptions are distributed close to fresh encryptions (at the corresponding level).

For PRE satisfying these assumptions, we show that the framework of Jafargholi et al. [JKK+17], who gave an abstraction of the techniques from [FKPR14], can be applied. This framework has been used to show adaptive security of a variety of cryptographic protocols (e.g., secret sharing, garbled circuits etc.) in the "symmetric-key" setting while avoiding an exponential loss that typically results from the guessing step when going from selective to adaptive security. Its application to PREs in the work is the first in the "public-key" setting. We describe their framework in more detail below.

The [JKK+17] framework. A standard way to prove adaptive security is to first define a "selective" variant that requires the adversary to commit to some of its choices (e.g., whom to corrupt, or on what input to be challenged at the end) at the beginning of the game. Let $\mathcal{W}$ denote the set of all possible choices.

Consider a selective security notion defined as two games H^0 and H^1 being indistinguishable. A security proof often uses a hybrid argument: one defines a sequence of hybrid games $(\mathsf{H}_0, \ldots, \mathsf{H}_\tau)$ where the first and last games correspond to the original selective games (i.e., $\mathsf{H}^0 = \mathsf{H}_0$ and $\mathsf{H}^1 = \mathsf{H}_\tau$). One then proves that any two consecutive hybrids (H_t and H_{t+1}) are ϵ-indistinguishable. As indistinguishability satisfies the triangle inequality, the extreme games H_0 and H_τ are $(\epsilon \cdot \tau)$-indistinguishable.

Now to prove security against an adaptive adversary (who will not reveal its choices at the beginning), one defines a new reduction that just guesses the adversary's future choices at random from the set $\mathcal{W}$ and then follows the selective reduction. Conditioned on the guess being correct, this reduction has the same success probability as the selective one.

Unfortunately, the overall loss in security of this second step is as large as the size of $\mathcal{W}$, which is typically exponential (e.g., exponential in the number of parties that can be corrupted). Thus, if the selective reduction implied ϵ-indistinguishability (based on some underlying assumption), the adaptive reduction will only imply $(\epsilon \cdot |\mathcal{W}|)$-indistinguishability, which in most cases will be meaningless.

The key observation in [JKK+17] was that in many selective reductions as above, only a highly compressed version $h(w)$ of the information $w \in \mathcal{W}$ that the adversary commits to is actually used in the simulation of intermediate hybrids. Jafargholi et al. called these "partially selective" hybrids, as opposed to the original hybrids, which are "fully selective". They show that the security loss in such cases is only exponential in the length of $h(w)$ (its the longest value for any two consecutive hybrids), and not exponential in the length of the entire w.

In all the instances to which the [JKK+17] framework has been applied the simulation of the security game depends on some underlying graph (e.g., the access structure in secret sharing or the Boolean circuit in case of garbling) and the hybrid games involve incremental changes to the simulation *depending* on the structure of this graph. Jafargholi et al. managed to decouple the particulars of the simulation from the design of the hybrids by using a pebbling game on the graph (the graph must thus be directed and acyclic). To be more precise, they associated the simulation of a hybrid (H_t) to a pebbling configuration ($\mathcal{P}_t$), and therefore the incremental changes in the simulation to the pebbling sequence $(\mathcal{P}_0, \ldots, \mathcal{P}_\tau)$. In particular, if a vertex carries a pebble then the part of simulation of the hybrid that is dependent on the vertex is carried out in a different manner (e.g., in garbling using Yao's scheme the ciphertexts in the garbled table for a gate are all bogus). The rules of the simulation is what then determines the pebbling rules, i.e., when exactly a pebble can be placed on or removed from a vertex. The extreme hybrids correspond to the initial and final pebbling configurations, and the immediate goal is to show that two hybrids that differ by a pebble are indistinguishable to an adversary. Indistinguishability of the original games then follows by transitivity of indistinguishability.

In the fully selective games of the above examples, the adversary commits to the whole graph; but, as explained above, knowledge of the vertices that are pebbled suffices to simulate the intermediate hybrids. Therefore, in the partially selective game the adversary "commits" to some pebbling configuration. Since we have established a correspondence between the simulation and a pebbling configuration, the task of designing a better sequence of hybrids has been reduced to finding a better pebbling sequence. In particular, the fewer pebbles are on the graph at any particular time, the more concisely we can describe this configuration, and thus the smaller the incurred security loss.

Designing the hybrids. The graph that underlies the simulation in adaptive CPA and HRA is precisely the recoding graph. (Strictly speaking, it suffices to consider the subgraph that is reachable from the challenge vertex, which we will call the "challenge graph".) The presence (or not) of a pebble on a vertex dictates how the re-encryption and re-key queries outgoing from that vertex are simulated. Therefore in the fully selective games, the adversary commits to the recoding graph (which is different from the original selective game in which the adversary committed to the set of corrupt users), whereas in the partially selective games it "commits" just to a pebbling configuration.

Let us first consider adaptive CPA: the edges of the recoding graph correspond to the re-key and re-encryption queries made by the adversary during the

game. For simplicity, assume that the recoding graph has a single source vertex i^* that is also the vertex the adversary wants to be challenged on. Once it has made all the queries, the adversary receives its challenge, which is the encryption of either m_0^* or m_1^* under $\mathtt{pk}_{i^*}$; let CPA^0 and CPA^1 denote the respective games. In case there are no outgoing edges from i^*, indistinguishability of CPA^0 and CPA^1 follows from ciphertext indistinguishability (the first assumption): The reduction embeds the challenge public key (of the indistinguishability game) as the i^*-th key, relays (m_0^*, m_1^*) to its challenger and forwards the challenge ciphertext it receives to the adversary. As there are no outgoing re-keys from i^*, the simulation does not require the secret key $\mathtt{sk}_{i^*}$.

In case i^* does have outgoing edges, the idea is to use a sequence of hybrids to reach a game where knowledge of $\mathtt{sk}_{i^*}$ is not required for simulation, just like above. To argue indistinguishability of hybrids, we use weak key-privacy, which guarantees that a re-key looks pseudorandom given the source and target public keys. Weak key-privacy allows the simulator to fake the outgoing edges from a vertex, after which the secret key for this vertex is not required for simulation anymore. However, the simulator cannot fake edges right away: it has to fake all children of a vertex first, before it can rely on weak key-privacy. Consequently, the pebbling must obey the following rule: in a move, a pebble can be placed on or removed from a vertex only if all its children carry pebbles.

To be precise, in game H_t^b, for each pebbled vertex in $\mathcal{P}_t$ all queried re-keys outgoing from that vertex are faked. Observe that as the secret key corresponding to a vertex is used only for the generation of the re-keys outgoing from that vertex, the simulation of a hybrid can be carried out *without* knowledge of the secret key corresponding to the pebbled vertices. Thus, a pebbling sequence describes a sequence of hybrids.

Main result. Our main result bounds the security loss for arbitrary recoding graphs in terms of their space and time complexity, where a graph is said to have space complexity σ and time complexity τ if there exists a valid pebbling strategy for that graph that uses at most σ pebbles and requires at most τ moves. More generally, a class of graphs has space complexity σ and time complexity τ if this is the case for every graph in that class.

Theorem 1 (Informal Theorem 5). *Let $\mathcal{G}(n)$ denote a family of graphs on n vertices with space-complexity σ and time-complexity τ. Then a PRE that is ciphertext-indistinguishable and weakly key-private for computationally bounded adversaries is also adaptively CPA-secure against computationally bounded adversaries for recoding graphs in $\mathcal{G}$ with a loss in security of $\approx \tau \cdot n^{\sigma}$. If the PRE is also statistically source-hiding then it is also adaptively HRA-secure.*

In many applications, the underlying recoding graph is has a very particular structure like trees (or even paths) and low-depth graphs, which cover many interesting applications. For paths, or fixed-arity trees, our reduction only loses a quasi-polynomial factor. For low-depth graphs, the loss is exponential only in the depth (and thus polynomial for fixed depth-graphs). Below, we mention two such applications.

Table 1: PRE schemes we prove adaptively CPA and HRA secure (see §5 and the full version [FKKP18] for the definitions of the assumptions).

Scheme	Setting	Assumption(s)	Hops
[CCL+14] (Constr. 2, Section 5)	Lattices	LWE	Multiple
[AFGH05] ([FKKP18, Constr. 4])	Bilinear maps	eDBDH and XDH	Single
[ABH09] ([FKKP18, Constr. 6])	Bilinear maps	eDBDH and DLin	Single
[Gen09] ([FKKP18, Constr. 5])	–	FHE	Multiple

1. In *key rotation* for encrypted cloud storage, a client has its data encrypted on a server, and occasionally wants to re-encrypt it (say, to restore security after key leakage). As the client does not trust the server, it will not want to hand it the decryption key. When using PRE, the client can simply send a re-key to the server, which enables it to locally re-encrypt all ciphertexts to the new key. In this application the recoding graph is simply a chain.
2. Another common application is *forwarding of encrypted email* without involving the receiver, say, for delegation during vacation or for filtering spam emails. In most cases the underlying delegation structure will be captured by simple graphs. For example, if delegation only happens to subordinates, the depth of the recoding graph is bounded by the depth of the hierarchy of the organisation.

1.3.3 Third Contribution: Adaptively-Secure PREs. Finally, we show that the aforementioned three properties are satisfied by several existing constructions or by minor variants thereof, and thus Theorem 1 can be applied to them. An overview of these schemes is given in Table 1. We consider the most interesting corollary to our results the adaptive security of the LWE-based scheme by Chandran et al. [CCL+14]:

Theorem 2 (Informal Theorem 6). *The quasi-polynomially secure decisional LWE problem implies multi-hop, unidirectional adaptively CPA/HRA-secure PRE for chains or complete binary trees.*

2 Formal Definitions

Notation. We use $[a, b]$ to denote $\{a, a+1, \ldots, b\}$ and $[b]$ as a shorthand for $[1, b]$. We will only consider logarithms to the base 2 (i.e., $\log := \log_2$). For two sets $\mathcal{X}, \mathcal{Y}$ we write $\mathcal{X} \Delta \mathcal{Y}$ for the symmetric difference. We write $x \leftarrow \mathcal{X}$ for sampling an element x uniformly at random from the set $\mathcal{X}$; analogously, $x_1, \ldots, x_n \leftarrow \mathcal{X}$ denotes sampling $x_1, \ldots, x_n$ independently and uniformly at random from the set $\mathcal{X}$. To indicate sampling according to a distribution X on $\mathcal{X}$, we write $x \leftarrow X$.

By $[X]$ we denote the support of X, i.e., the values with positive probability. For two distributions X, Y, $\Delta(X, Y)$ denotes their statistical distance. We write $X \equiv Y$ if X has the same input/output distribution as Y. Two distributions $X = \{X_\kappa\}_{\kappa\in\mathbb{N}}$ and $\{Y_\kappa\}_{\kappa\in\mathbb{N}}$ are (s, ϵ)-indistinguishable, denoted $X \approx_{(s,\epsilon)} Y$, if for every adversary A of size at most s

$$|\mathbb{P}[\mathsf{A}(X) = 1] - \mathbb{P}[\mathsf{A}(Y) = 1]| \leq \epsilon.$$

Throughout the paper, we will repeatedly use the following lemma concerning the transitivity of the indistinguishability relation $\approx$:

Lemma 1. *Let X, Y, Z be distributions on a set $\mathcal{X}$. If $X \approx_{(s_1,\epsilon_1)} Y$ and $Y \approx_{(s_2,\epsilon_2)} Z$, then $X \approx_{(\min(s_1,s_2),\epsilon_1+\epsilon_2)} Z$.*

For indistinguishability-based security games, we use $\langle \mathsf{G}, \mathsf{A}\rangle$ to denote the bit output by the challenger G at the end of its interaction with the adversary A. We say that two games G^0 and G^1 are (s, ϵ)-indistinguishable, denoted $\mathsf{G}^0 \approx_{(s,\epsilon)} \mathsf{G}^1$, if for every adversary A of size at most s

$$|\mathbb{P}[\langle \mathsf{G}^0, \mathsf{A}\rangle = 1] - \mathbb{P}[\langle \mathsf{G}^1, \mathsf{A}\rangle = 1]| \leq \epsilon.$$

For an algorithm A, we use s_A to denote its size; in a similar manner, for a set $\mathcal{X}$, we use $s_\mathcal{X}$ to denote the complexity of sampling from $\mathcal{X}$ uniformly at random.

Notation for graphs. We let $G = (\mathcal{V}, \mathcal{E})$ denote a directed graph with vertices $\mathcal{V}$ (usually $\mathcal{V} = [n]$ for some $n \in \mathbb{N}$) and edges $\mathcal{E} \subseteq \mathcal{V}^2$. The indegree (resp., outdegree) of a vertex is defined as the number of edges coming in to (resp., going out of) that vertex. The indegree (resp., outdegree) of the graph is the maximum indegree (resp., outdegree) over all the vertices. A vertex with indegree (resp., outdegree) zero is called a source (resp., sink). A vertex i is *connected* to another vertex j (or alternatively j is reachable from i) if there is a directed path from i to j in G.

2.1 Proxy Reencryption: Formal Definitions

Definition 1 (Multi-hop, unidirectional PRE). *A multi-hop, unidirectional PRE scheme for a message space $\mathcal{M}$ consists of the six-tuple of algorithms $(\mathsf{S}, \mathsf{K}, \mathsf{RK}, \mathsf{E}, \mathsf{D}, \mathsf{RE})$, which are explained below.*

- $\mathsf{S}(1^\kappa, 1^\lambda) \to \mathtt{pp}$**:** *On input the security parameter κ and the maximum level λ (both in unary) supported by the scheme,* **setup** *outputs the public parameters* $\mathtt{pp}$*. We assume that* $\mathtt{pp}$ *is implicit in other function calls.*
- $\mathsf{K}(\mathtt{pp}) \to (\mathtt{pk}, \mathtt{sk})$**: Key generation** *returns a public key* $\mathtt{pk}$ *and the corresponding secret key* $\mathtt{sk}$*.*
- $\mathsf{RK}((\mathtt{pk}_i, \mathtt{sk}_i), \mathtt{pk}_j) \to \mathtt{rk}_{i,j}$**:** *On input a source key pair $(\mathtt{pk}_i, \mathtt{sk}_i)$ and a target public key $\mathtt{pk}_j$,* **re-key generation** *generates a unidirectional re-encryption key (rekey, for short) $\mathtt{rk}_{i,j}$.*

$\mathsf{E}(\mathtt{pk},(m,\ell)) \to (c,\ell)$**: Encryption** *takes as input the public key* $\mathtt{pk}$*, a message* m *and a level* $\ell \in [\lambda]$*, and outputs a level-*ℓ *ciphertext* (c,ℓ)*.*

$\mathsf{D}(\mathtt{sk},(c,\ell)) \to m$ *: On input a ciphertext* (c,ℓ) *and the secret key* $\mathtt{sk}$*,* **decryption** *outputs a message* m*, or the symbol* $\perp$ *(if the ciphertext is invalid).*

$\mathsf{RE}(\mathtt{rk}_{i,j},\mathtt{pk}_i,\mathtt{pk}_j,(c_i,\ell)) \to (c_j,\ell+1)$**: Reencryption** *takes a re-key* $\mathtt{rk}_{i,j}$*, a source public key* $\mathtt{pk}_i$*, a target public key* $\mathtt{pk}_j$ *and a level-*ℓ *ciphertext* c_i *under* $\mathtt{pk}_i$ *and transforms it to a level-*$(\ell+1)$ *ciphertext* c_j *under* $\mathtt{pk}_j$*. Only ciphertexts belonging to levels* $\ell \in [\lambda-1]$ *can be re-encrypted. In constructions where arguments* $\mathtt{pk}_i$ *and/or* $\mathtt{pk}_j$ *are optional, we simply drop them.*

Definition 1 differs slightly from the definition of multi-hop unidirectional PRE in [FL17]. Here, the re-keys are level-agnostic: the same re-key can be used to re-encrypt a ciphertext belonging to any level. In [FL17], however, a re-key associated to a level *cannot* be used to re-encrypt a ciphertext from a different level. We require the PRE to satisfy the following two correctness properties.

Definition 2 (Correctness [ABH09]). *A proxy re-encryption scheme (as in Definition 1) is correct w.r.t. the message space* $\mathcal{M}$ *if the following two properties hold:*

1. *Correctness of encryption:* $\forall \kappa,\lambda \in \mathbb{N}\ \forall\, \mathtt{pp} \in [\mathsf{S}(1^\kappa,1^\lambda)]\ \forall\,(\mathtt{pk},\mathtt{sk}) \in [\mathsf{K}(\mathtt{pp})]$ $\forall\,(m,\ell) \in \mathcal{M} \times [\lambda]$*:*

$$\mathbb{P}[\mathsf{D}\big(\mathtt{sk},\mathsf{E}(\mathtt{pk},(m,\ell))\big) \neq m] = \mathsf{negl}(\kappa,\lambda),$$

   where the probability is over the random coins of E*.*
2. *Correctness of re-encryption:* $\forall\,\kappa,\lambda \in \mathbb{N}\ \forall\, \mathtt{pp} \in [\mathsf{S}(1^\kappa,1^\lambda)]\ \forall\,(\mathtt{pk}_i,\mathtt{sk}_i)$, $(\mathtt{pk}_j,\mathtt{sk}_j) \in [\mathsf{K}(\mathtt{pp})]\ \forall\, \mathtt{rk}_{i,j} \in [\mathsf{RK}((\mathtt{pk}_i,\mathtt{sk}_i),\mathtt{pk}_j)]\ \forall\,(m,\ell) \in \mathcal{M} \times [\lambda-1]$*:*

$$\mathbb{P}[\mathsf{D}\big(\mathtt{sk}_j,\mathsf{RE}(\mathtt{rk}_{i,j},\mathtt{pk}_i,\mathtt{pk}_j,\mathsf{E}(\mathtt{pk}_i,(m,\ell)))\big) \neq m] = \mathsf{negl}(\kappa,\lambda),$$

   where the probability is over the random coins of E *and* RE*.*

2.2 Modelling Security

2.2.1 Selective Corruption. The selective security of a multi-hop, unidirectional PRE scheme against a chosen-plaintext attack is modelled using the security game given in Game 1. It is an extension of the security model for single-hop PRE from [ABH09] to the multi-hop setting.[3] The limiting feature of the model is that the adversary has to fix, beforehand in Phase 1 (see Game 1), the honest and corrupt public keys. Its goal is to distinguish an encryption of m_0 from an encryption of m_1 (for m_0, m_1 of its choice) under a key of its choice. The game aborts if the adversary does one of the following:

- query the challenge oracle on a corrupt public key ($\mathtt{abort}_1$);
- request a re-key from an honest key to a corrupt key ($\mathtt{abort}_2$); or
- query a re-encryption from an honest to a corrupt key ($\mathtt{abort}_3$).

Challenger $\mathsf{sCPA}^b(1^\kappa, 1^\lambda, n)$
1: Set $\mathcal{C} = \emptyset$ ▷ Stores the corrupt public keys
2: $\mathtt{pp} \leftarrow \mathsf{PRE.S}(1^\kappa, 1^\lambda)$, $(\mathtt{pk}_1, \mathtt{sk}_1), \ldots, (\mathtt{pk}_n, \mathtt{sk}_n) \leftarrow \mathsf{PRE.K}(\mathtt{pp})$ ▷ Generate keys
3: $\forall i, j \in [n], i \neq j : \mathtt{rk}_{i,j} \leftarrow \mathsf{PRE.RK}((\mathtt{pk}_i, \mathtt{sk}_i), \mathtt{pk}_j)$ ▷ Generate re-keys
4: $\mathsf{state} \leftarrow \mathsf{A}_1^{(\mathtt{corrupt},\cdot)}(\mathtt{pp})$ ▷ Phase 1
5: $b' \leftarrow \mathsf{A}_2^{(\mathtt{rekey},\cdot,\cdot),(\mathtt{reencrypt},\cdot,\cdot,\cdot),(\mathtt{challenge},\cdot,\cdot,\cdot)}(\mathtt{pk}_1, \ldots, \mathtt{pk}_n, \mathsf{state})$ ▷ Phase 2
6: **return** b'

Oracle $(\mathtt{corrupt}, i)$
1: Add i to $\mathcal{C}$
2: **return** $\mathtt{sk}_i$

Oracle $(\mathtt{rekey}, i, j)$
1: **if** $i \notin \mathcal{C}$ and $j \in \mathcal{C}$ **then** HALT **end if** ▷ $\mathtt{abort}_2$
2: **return** $\mathtt{rk}_{i,j}$

Oracle $(\mathtt{reencrypt}, i, j, (c_i, \ell))$
1: **if** $i \notin \mathcal{C}$ and $j \in \mathcal{C}$ **then** HALT **end if** ▷ $\mathtt{abort}_3$
2: **return** $(c_j, \ell + 1) \leftarrow \mathsf{PRE.RE}(\mathtt{rk}_{i,j}, \mathtt{pk}_i, \mathtt{pk}_j, (c_i, \ell))$

Oracle $(\mathtt{challenge}, i^*, (m_0^*, m_1^*), \ell^*)$ ▷ Single access
1: **if** $i^* \in \mathcal{C}$ **then** HALT **end if** ▷ $\mathtt{abort}_1$
2: **return** $(c_{i^*}, \ell^*) \leftarrow \mathsf{PRE.E}(\mathtt{pk}_{i^*}, (m_b^*, \ell^*))$

Game 1: sPRE-CPA

Definition 3 (sPRE-CPA-security). *A PRE scheme is (s, ϵ)-selectively secure against chosen-plaintext attack if $\mathsf{sCPA}^0 \approx_{(s,\epsilon)} \mathsf{sCPA}^1$, where sCPA^b is defined in Game 1.*

Security against honest-reencryption attack. A stronger security definition was introduced in [Coh17] to address some of the restrictions that sPRE-CPA imposes on the adversary. The idea is to allow re-encryptions from honest to corrupt keys, if the ciphertexts to re-encrypt were honestly generated. The adversary can obtain such honest ciphertexts via an `encrypt` oracle, which stores them in a list. The `reencrypt` oracle now takes the index of an honestly generated ciphertext. It was shown in [Coh17] that (selective) HRA-security implies (selective) CPA-security and also that if the PRE scheme is re-encryption-simulatable (a generalisation of Definition 9) then (selective) CPA-security implies (selective) HRA-security. In sPRE-HRA, which we formally define in Game 2, $\mathtt{abort}_3$ is relaxed to

– $\mathtt{abort}_3^*$: The adversary queries the re-encryption of a ciphertext that is the result of a chain of re-encryptions of the challenge ciphertext from an honest to a corrupt key.

Definition 4 (sPRE-HRA-security). *A PRE scheme is (s, ϵ)-selectively secure against honest-reencryption attack if $\mathsf{sHRA}^0 \approx_{(s,\epsilon)} \mathsf{sHRA}^1$, where sHRA^b is defined in Game 2.*

[3] [FL17] formalised security differently; we stick to the definition from [ABH09].

Challenger $\mathsf{sHRA}^b(1^\kappa, 1^\lambda, n)$
1: Set $\mathcal{C}, \mathcal{E} = \emptyset$ ▷ Stores corrupt keys and issued re-keys and re-encryptions
2: Set $C = 0$ ▷ Counts ciphertexts generated
3: Set $\mathcal{L}, \mathcal{L}^* = \emptyset$ ▷ Stores honest ciphertexts and which derived from challenge
4: $\mathsf{pp} \leftarrow \mathsf{PRE.S}(1^\kappa, 1^\lambda)$, $(\mathtt{pk}_1, \mathtt{sk}_1), \ldots, (\mathtt{pk}_n, \mathtt{sk}_n) \leftarrow \mathsf{PRE.K}(\mathsf{pp})$ ▷ Generate keys
5: $\forall i, j \in [n], i \neq j : \mathtt{rk}_{i,j} \leftarrow \mathsf{PRE.RK}((\mathtt{pk}_i, \mathtt{sk}_i), \mathtt{pk}_j)$ ▷ Generate re-keys
6: $\mathsf{state} \leftarrow \mathsf{A}_1^{(\mathtt{corrupt},\cdot)}(\mathsf{pp})$ ▷ Phase 1
7: $b' \leftarrow \mathsf{A}_2^{(\mathtt{encrypt},\cdot,\cdot),(\mathtt{rekey},\cdot,\cdot),(\mathtt{reencrypt},\cdot,\cdot,\cdot),(\mathtt{challenge},\cdot,\cdot,\cdot)}(\mathtt{pk}_1, \ldots, \mathtt{pk}_n, \mathsf{state})$ ▷ Phase 2
8: **return** b'

Oracles `corrupt` and `rekey` are defined like in Game 1.

Oracle $(\mathtt{encrypt}, i, (m, \ell))$
1: $(c, \ell) \leftarrow \mathsf{PRE.E}(\mathtt{pk}_i, (m, \ell))$
2: Increment C and add $(C, i, m, (c, \ell))$ to $\mathcal{L}$
3: **return** (c, ℓ)

Oracle $(\mathtt{reencrypt}, i, j, k)$
1: Retrieve $(k, i, m, (c_i, \ell))$ from $\mathcal{L}$ and increment C
2: $(c_j, \ell + 1) \leftarrow \mathsf{PRE.RE}(\mathtt{rk}_{i,j}, \mathtt{pk}_i, \mathtt{pk}_j, (c_i, \ell))$
3: **if** $k \in \mathcal{L}^*$ **then** ▷ The ciphertext is derived from the challenge
4: **if** $j \in \mathcal{C}$ **then** HALT **else** add C to $\mathcal{L}^*$ ▷ $\mathtt{abort}_3^*$ **end if**
5: **end if**
6: Add $(C, j, m, (c_j, \ell + 1))$ to $\mathcal{L}$
7: **return** $(c_j, \ell + 1)$

Oracle $(\mathtt{challenge}, i^*, (m_0^*, m_1^*), \ell^*)$ ▷ Single access
1: Compute $(c_{i^*}, \ell^*) \leftarrow \mathsf{PRE.E}(\mathtt{pk}_{i^*}, (m_b^*, \ell^*))$ and increment C
2: **if** $i^* \in \mathcal{C}$ **then** HALT **else** add C to $\mathcal{L}^*$ ▷ $\mathtt{abort}_1$ **end if**
3: Add $(C, i^*, m_b^*, (c_{i^*}, \ell^*))$ to $\mathcal{L}$
4: **return** (c_{i^*}, ℓ^*)

Game 2: sPRE-HRA

2.2.2 Modelling Adaptive Corruption. The adaptive security games corresponding to Games 1 and 2 are given in Games 3 and 4, respectively. To model adaptive corruption, we think of the game being played on a directed graph $G = (\mathcal{V}, \mathcal{E})$ called the "recoding" graph. The vertices of the recoding graph correspond to the public keys, i.e., $\mathcal{V} = [n]$. The edges are derived from the re-keys and re-encryptions issued to the adversary in the security game, and their purpose is to ensure that the adversary does not win the game in a trivial manner. In particular, the recoding graph is defined so that *no* corrupt key is reachable from the challenge key. To be precise, in CPA an edge (i, j) is added to $\mathcal{E}$ if the adversary made either a $(\mathtt{rekey}, i, j)$ or $(\mathtt{reencrypt}, i, j, \cdot)$ query (see Game 3 and Figure 1). Consequently, the adversary is forbidden from making *any* re-key or re-encryption queries to a corrupt user that is reachable from the challenge key.[4]

[4] The selective CPA notion (Game 1) is in fact more restrictive in that it does not allow re-keys and re-encryptions from *any* honest user to a corrupt user.

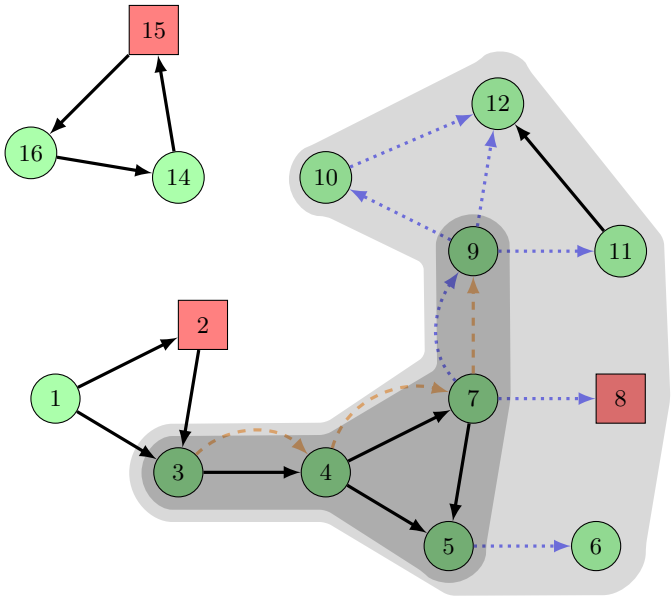

Fig. 1: Recoding graph. The (round) green nodes represent the honest users, whereas the (square) red nodes are the corrupted ones. The edges denote the recoding information. In particular, the (solid) black edges are the re-keys, the (dashed) orange edges are the re-encryptions related to the challenge ciphertext (therefore, 3 is the challenge vertex) and (dotted) blue edges represent the remaining re-encryptions. For CPA, all the edges are counted, but for HRA the blue (dotted) edges are *not* counted. The subgraph of the recoding graph that forms the challenge graph (cf. §4) is shaded: the darker inner shading for HRA, whereas the lighter outer shading is the challenge graph for CPA. Note that the edge $(7, 8)$ is valid in the case of HRA, but invalid for CPA (and therefore the CPA challenger would abort at the end of such an execution.)

For HRA, on the other hand, (i, j) is added to $\mathcal{E}$ if the adversary made either a $(\texttt{rekey}, i, j)$ query or a $(\texttt{reencrypt}, i, j, k)$ query where the k-th ciphertext is a re-encryption of the challenge ciphertext (see Game 4 and Figure 1). This is less restrictive than in CPA: the adversary can make re-encryption queries to a corrupt user that is reachable from the challenge key *unless* it is related to the challenge ciphertext.

For comparison we have reformulated the selective notions defined in Games 1 and 2 using a recoding graph instead of explicit aborts. Games 8 and 9 in the full version define the exact same notions as Games 1 and 2, respectively.

Definition 5 (PRE-CPA-security). *A PRE scheme is (s, ϵ)-adaptively secure against chosen-plaintext attack if $\mathsf{CPA}^0 \approx_{(s,\epsilon)} \mathsf{CPA}^1$, where CPA^b is defined in Game 3.*

Challenger $\mathsf{CPA}^b(1^\kappa, 1^\lambda, n)$
1: Set $\mathcal{C}, \mathcal{E} = \emptyset$ ▷ Stores corrupt keys and issued re-keys and re-encryptions
2: $\mathtt{pp} \leftarrow \mathsf{PRE.S}(1^\kappa, 1^\lambda)$, $(\mathtt{pk}_1, \mathtt{sk}_1), \ldots, (\mathtt{pk}_n, \mathtt{sk}_n) \leftarrow \mathsf{PRE.K}(\mathtt{pp})$ ▷ Generate keys
3: $\forall i, j \in [n], i \neq j : \mathtt{rk}_{i,j} \leftarrow \mathsf{PRE.RK}((\mathtt{pk}_i, \mathtt{sk}_i), \mathtt{pk}_j)$ ▷ Generate re-keys
4: $b' \leftarrow \mathsf{A}^{(\mathtt{corrupt},\cdot),(\mathtt{rekey},\cdot,\cdot),(\mathtt{reencrypt},\cdot,\cdot,\cdot),(\mathtt{challenge},\cdot,\cdot,\cdot)}(\mathtt{pp}, \mathtt{pk}_1, \ldots, \mathtt{pk}_n)$
5: **if** A made call $(\mathtt{challenge}, i^*, \cdot, \cdot)$ for some i^* **then** ▷ Check abort conditions
6: **if** $\exists\, i \in \mathcal{C} : i^*$ is connected to i in $([n], \mathcal{E})$ **then return** 0 **end if**
7: **end if**
8: **return** b'

Oracle $(\mathtt{corrupt}, i)$
1: Add i to $\mathcal{C}$
2: **return** $\mathtt{sk}_i$

Oracle $(\mathtt{rekey}, i, j)$
1: Add (i, j) to $\mathcal{E}$ ▷ Add to recoding graph
2: **return** $\mathtt{rk}_{i,j}$

Oracle $(\mathtt{reencrypt}, i, j, (c_i, \ell))$
1: Add (i, j) to $\mathcal{E}$ ▷ Add to recoding graph
2: **return** $(c_j, \ell + 1) \leftarrow \mathsf{PRE.RE}(\mathtt{rk}_{i,j}, \mathtt{pk}_i, \mathtt{pk}_j, (c_i, \ell))$

Oracle $(\mathtt{challenge}, i^*, (m_0^*, m_1^*), \ell^*)$ ▷ Single access
1: **return** $(c_{i^*}, \ell^*) \leftarrow \mathsf{PRE.E}(\mathtt{pk}_{i^*}, (m_b^*, \ell^*))$

Game 3: PRE-CPA

Definition 6 (PRE-HRA-security). *A PRE scheme is (s, ϵ)-adaptively secure against honest-reencryption attack if $\mathsf{HRA}^0 \approx_{(s,\epsilon)} \mathsf{HRA}^1$, where HRA^b is defined in Game 4.*

3 Preliminaries

This section provides the background necessary for the main results in §4. We start with the security assumptions on PREs that allow us to prove adaptive security (§3.1) and then give an overview of the framework of [JKK+17] (§3.2), the description of the pebbling game that is used in the design of the hybrids (§3.3).

3.1 Security Assumptions on PRE

We describe the three security properties of PRE schemes that allow us to prove adaptive security: indistinguishability, key-privacy and source-hiding.

Indistinguishability of ciphertexts. For proxy re-encryption, we require the notion of indistinguishability, as defined for public-key encryption, to hold on *all* levels:

Definition 7 (Indistinguishability). *A proxy re-encryption scheme PRE has (s, ϵ)-indistinguishable ciphertexts if $\mathsf{IND}^0 \approx_{(s,\epsilon)} \mathsf{IND}^1$ with IND as in Game 5.*

Challenger $\mathsf{HRA}^b(1^\kappa, 1^\lambda, n)$
1: Set $\mathcal{C}, \mathcal{L}, \mathcal{L}^* = \emptyset$ and $C = 0$ ▷ $\mathcal{L}$ stores honest enc's, $\mathcal{L}^*$ marks challenge reenc's
2: $\mathcal{E} = \emptyset$ ▷ The edges of the recoding graph
3: $\mathtt{pp} \leftarrow \mathsf{PRE.S}(1^\kappa, 1^\lambda)$, $(\mathtt{pk}_1, \mathtt{sk}_1), \ldots, (\mathtt{pk}_n, \mathtt{sk}_n) \leftarrow \mathsf{PRE.K}(\mathtt{pp})$ ▷ Generate keys
4: $\forall i, j \in [n], i \neq j : \mathtt{rk}_{i,j} \leftarrow \mathsf{PRE.RK}((\mathtt{pk}_i, \mathtt{sk}_i), \mathtt{pk}_j)$ ▷ Generate re-keys
5: $b' \leftarrow \mathsf{A}^{(\mathtt{corrupt},\cdot),(\mathtt{rekey},\cdot,\cdot),(\mathtt{encrypt},\cdot,\cdot),(\mathtt{reencrypt},\cdot,\cdot,\cdot),(\mathtt{challenge},\cdot,\cdot,\cdot)}(\mathtt{pp}, \mathtt{pk}_1, \ldots, \mathtt{pk}_n)$
6: **if** A made call $(\mathtt{challenge}, i^*, \cdot, \cdot)$ for some i^* **then** ▷ Check abort conditions
7: **if** $\exists\, i \in \mathcal{C} : i^*$ is connected to i **then return** 0 **end if**
8: **end if**
9: **return** b'

Oracle $(\mathtt{corrupt}, i)$
1: Add i to $\mathcal{C}$
2: **return** $\mathtt{sk}_i$

Oracle $(\mathtt{rekey}, i, j)$
1: Add (i, j) to $\mathcal{E}$ ▷ Add to recoding graph
2: **return** $\mathtt{rk}_{i,j}$

Oracle $(\mathtt{encrypt}, i, (m, \ell))$
1: $c \leftarrow \mathsf{PRE.E}(\mathtt{pk}_i, (m, \ell))$, increment C and add $(C, i, m, (c, \ell))$ to $\mathcal{L}$
2: **return** c

Oracle $(\mathtt{reencrypt}, i, j, k)$
1: Retrieve $(k, i, m, (c_i, \ell))$ from $\mathcal{L}$
2: $(c_j, \ell + 1) \leftarrow \mathsf{PRE.RE}(\mathtt{rk}_{i,j}, \mathtt{pk}_i, \mathtt{pk}_j, (c_i, \ell))$
3: Increment C and add $(C, j, m, (c_j, \ell + 1))$ to $\mathcal{L}$
4: **if** $k \in \mathcal{L}^*$ **then** ▷ c_j derived from challenge
5: Add C to $\mathcal{L}^*$ and add (i, j) to $\mathcal{E}$ ▷ Add to recoding graph
6: **end if**
7: **return** $(c_j, \ell + 1)$

Oracle $(\mathtt{challenge}, i^*, (m_0^*, m_1^*), \ell^*)$ ▷ Single access
1: Compute $(c_{i^*}, \ell^*) \leftarrow \mathsf{PRE.E}(\mathtt{pk}_{i^*}, (m_b^*, \ell^*))$
2: Increment C, add $(C, i^*, m_b^*, (c_{i^*}, \ell^*))$ to $\mathcal{L}$ and C to $\mathcal{L}^*$
3: **return** (c_{i^*}, ℓ^*)

Game 4: PRE-HRA

Challenger $\mathsf{IND}^b(1^\kappa, 1^\lambda)$
1: $\mathtt{pp} \leftarrow \mathsf{PRE.S}(1^\kappa, 1^\lambda)$, $(\mathtt{pk}, \mathtt{sk}) \leftarrow \mathsf{PRE.K}(\mathtt{pp})$
2: **return** $b' \leftarrow \mathsf{A}^{(\mathtt{challenge},\cdot,\cdot)}(\mathtt{pp}, \mathtt{pk})$

Oracle $(\mathtt{challenge}, (m_0^*, m_1^*), \ell^*)$
1: **return** $\mathsf{PRE.E}(\mathtt{pk}, (m_b^*, \ell^*))$

Game 5: Security game IND for ciphertext indistinguishability

Key-Privacy. The original notion of key-privacy for PREs, which we refer to as "strong" key-privacy, was introduced in [ABH09]. It is modelled by a security game similar to sPRE-CPA: the adversary has access to `corrupt`, `rekey` and `reencrypt` oracles, but as a challenge it has to distinguish a real re-key from a re-key sampled *uniformly at random* from the support of re-keys. We refer the readers to [ABH09] for the details.

We only need a weaker definition stating that re-keys should hide the source keys. That is, the re-key $\mathtt{rk}_{0,1}$ from source $(\mathtt{pk}_0, \mathtt{sk}_0)$ to a target key $\mathtt{pk}_1$ should

be indistinguishable from a random source to $\mathtt{pk}_1$. In addition, we need this property to hold with respect to multiple re-keys. More formally, the security game for weak key-privacy is given in Game 6 where the simulator RK^* is defined as

$$\mathsf{RK}^*(\mathtt{pp}, \mathtt{pk}_1) := \mathsf{RK}((\mathtt{pk}_0, \mathtt{sk}_0), \mathtt{pk}_1) : (\mathtt{pk}_0, \mathtt{sk}_0) \leftarrow \mathsf{K}(\mathtt{pp}).$$

Definition 8 (Weak key-privacy). *Let $\delta \in \mathbb{N}$. A proxy re-encryption scheme* PRE *is (s, ϵ, δ)-weakly key-private if $\mathsf{KP}^0 \approx_{(s,\epsilon)} \mathsf{KP}^1$ with* KP *as in Game 6.*

Source-hiding. Source-hiding is a special case of re-encryption-simulatability, a notion that was introduced in [Coh17]. It requires that re-encryptions can be simulated without knowledge of the secret key. In particular, the simulated re-encryptions should be indistinguishable from re-encrypted ciphertexts even when given the secret keys for the source and target public keys, as well as the re-key that was used for re-encryption (hence the notion of indistinguishability is at least that of statistical indistinguishability). A PRE scheme is called *source-hiding* if re-encrypted ciphertexts have the same distribution as "fresh" ciphertexts, i.e., the encryption algorithm can be used as a simulator for re-encryption.

Definition 9 (Source-hiding). *A proxy re-encryption scheme* PRE *is (s, ϵ)-source-hiding if $\mathsf{SH}^0 \approx_{(s,\epsilon)} \mathsf{SH}^1$, with* SH *as defined in Game 7.*

3.2 Overview of [JKK+17]

Random guessing. A standard way to prove adaptive security is to first show security in a "selective" version of the adaptive game, in which the adversary commits to some of its future choices, and then use random guessing of the adversary's commitment to reduce adaptive security to selective security. For instance, consider the indistinguishability game for identity-based encryption: in the selective counterpart the adversary commits to the challenge identity at the start of the game, and the adaptive to selective reduction then works by guessing the challenge identity. More formally, let G^0 and G^1 denote the two adaptive games. For some function $g\colon \{0,1\}^* \to \mathcal{W}$ we define below the selective games $\mathsf{H}^0 = \mathsf{SEL}_{\mathcal{W}}[\mathsf{G}^0, g]$ and $\mathsf{H}^1 = \mathsf{SEL}_{\mathcal{W}}[\mathsf{G}^1, g]$ where the adversary commits to some information $w \in \mathcal{W}$ – for the case of IBE, $\mathcal{W}$ is the set of all identities. Note

Challenger $\mathsf{KP}^b(1^\kappa, 1^\lambda)$
1: $\mathtt{pp} \leftarrow \mathsf{PRE.S}(1^\kappa, 1^\lambda)$, $(\mathtt{pk}_0, \mathtt{sk}_0), \ldots, (\mathtt{pk}_\delta, \mathtt{sk}_\delta) \leftarrow \mathsf{K}(\mathtt{pp})$
2: $\forall j \in [\delta] : \mathtt{rk}^{(0)}_{0,j} \leftarrow \mathsf{RK}((\mathtt{pk}_0, \mathtt{sk}_0), \mathtt{pk}_j)$
3: $\mathtt{rk}^{(1)}_{0,j} \leftarrow \mathsf{RK}^*(\mathtt{pp}, \mathtt{pk}_j)$
4: **return** $b' \leftarrow \mathsf{A}(\mathtt{pp}, \mathtt{pk}_0, \ldots, \mathtt{pk}_\delta, \mathtt{rk}^{(b)}_{0,1}, \ldots, \mathtt{rk}^{(b)}_{0,\delta})$

Game 6: Security game KP for weak key-privacy

Challenger $\mathsf{SH}^b(1^\kappa, 1^\lambda)$
1: $\mathtt{pp} \leftarrow \mathsf{PRE.S}(1^\kappa, 1^\lambda)$
2: $(\mathtt{pk}_0, \mathtt{sk}_0), (\mathtt{pk}_1, \mathtt{sk}_1) \leftarrow \mathsf{PRE.K}(\mathtt{pp})$
3: $\mathtt{rk}_{0,1} \leftarrow \mathsf{PRE.RK}((\mathtt{pk}_0, \mathtt{sk}_0), \mathtt{pk}_1)$
4: $b' \leftarrow \mathsf{A}^{(\mathtt{challenge},\cdot,\cdot)}(\mathtt{pp}, (\mathtt{pk}_0, \mathtt{sk}_0), (\mathtt{pk}_1, \mathtt{sk}_1), \mathtt{rk}_{0,1})$
5: **return** b'

Oracle $(\mathtt{challenge}, m^*, \ell^*)$ ▷ $\ell^* \in [\lambda - 1]$
1: $c_0 \leftarrow \mathsf{PRE.E}(\mathtt{pk}_0, (m^*, \ell^*))$
2: $c_1^{(0)} \leftarrow \mathsf{PRE.RE}(\mathtt{rk}_{0,1}, \mathtt{pk}_0, \mathtt{pk}_1, c_0)$ ▷ Real re-encryption
3: $c_1^{(1)} \leftarrow \mathsf{PRE.E}(\mathtt{pk}_1, (m^*, \ell^* + 1))$ ▷ Simulate re-encryption
4: **return** $(c_0, c_1^{(b)})$

Game 7: Security game SH for source hiding

that the selective game gets a commitment w from the adversary but essentially ignores it during the rest of the game. It checks that the commitment matches what actually happened during the game only at the very end of the game; whether w matches is defined via the function g.

Definition 10 (Fully selectivised game [JKK+17]). *Given an (adaptive) game* G *and some function* $g\colon \{0,1\}^* \to \mathcal{W}$*, the selectivised game* $\mathsf{H} = \mathsf{SEL}_{\mathcal{W}}[\mathsf{G}, g]$ *is defined as follows. The adversary* A *first sends a commitment* $w \in \mathcal{W}$ *to* H*. Then* H *runs the challenger* G *against* A*, at the end of which* G *outputs a bit* $\hat{b}$*. Let* $\mathtt{transcript}$ *denote all communication exchanged between* G *and* A*. If* $g(\mathtt{transcript}) = w$ *then* H *outputs the bit* $\hat{b}$ *and else it outputs* 0*.*

Next, suppose that the selective security is proved using a hybrid argument. That is, to show the indistinguishability of H^0 and H^1 suppose we have a sequence of $\tau + 1$ (selective) hybrid games $\mathsf{H}^0 = \mathsf{H}_0, \mathsf{H}_1, \ldots, \mathsf{H}_\tau = \mathsf{H}^1$ (see Figure 2). If we only assume that neighbouring hybrids $\mathsf{H}_i, \mathsf{H}_{i+1}$ are indistinguishable, then by combining the hybrid argument and random guessing we get that G^0 and G^1 are indistinguishable with a loss in distinguishing advantage of $\tau \cdot |\mathcal{W}|$. The factor of $|\mathcal{W}|$ is the cost of the random guessing, whereas the factor of τ is due to the hybrid argument. This is stated in the following (recall that $s_{\mathcal{W}}$ denotes the complexity of sampling from $\mathcal{W}$):

Theorem 3 ([BB04, JKK+17]). *Assume we have two games defined via (adaptive) challengers* G^0 *and* G^1 *respectively. Let* $g\colon \{0,1\}^* \to \mathcal{W}$ *be an arbitrary function and define the selectivised games* $\mathsf{H}^b = \mathsf{SEL}_{\mathcal{W}}[\mathsf{G}^b, g]$ *for* $b \in \{0,1\}$*. Also assume that for each* $i \in [\tau]$*, the games* $\mathsf{H}_{i-1}, \mathsf{H}_i$ *are* (s, ϵ)*-indistinguishable. Then,* G^0 *and* G^1 *are* $(s - s_{\mathcal{W}},\ \epsilon \cdot \tau \cdot |\mathcal{W}|)$*-indistinguishable.*

The framework. In some cases only *part* of the information that the adversary commits to is used in simulating the intermediate hybrids, but when considering all the hybrids the whole commitment is being used. For example, the simulation

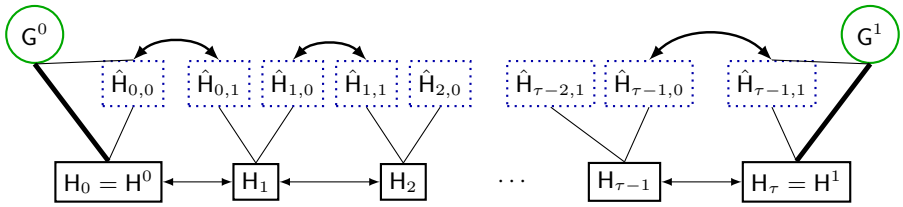

Fig. 2: A schematic diagram showing the relationship between adaptive, fully selective and partially selective hybrids. The adaptive games G^0 and G^1 are in green (circles); the fully selective games $\mathsf{H}_0, \ldots, \mathsf{H}_\tau$ are in solid black (boxes); and the partially selective games $\hat{\mathsf{H}}_{0,0}, \hat{\mathsf{H}}_{0,1}, \ldots, \hat{\mathsf{H}}_{\tau-1,0}, \hat{\mathsf{H}}_{\tau-1,1}$ are in (dotted) blue boxes. The arrows indicate indistinguishability.

of an intermediate hybrid in the case of IBE could rely on only certain bits of the challenge identity. It is shown in [JKK+17] that the security loss in such cases can be limited to the maximum size of the information used across any two successive hybrids.

More formally, [JKK+17] makes a stronger assumption: not only are neighbouring hybrids $\mathsf{H}_i, \mathsf{H}_{i+1}$ indistinguishable, but they are "selectivised" versions, of "partially" selective games $\hat{\mathsf{H}}_{i,0}, \hat{\mathsf{H}}_{i,1}$ which are already indistinguishable. In particular, for each pair of neighbouring hybrids $\mathsf{H}_i, \mathsf{H}_{i+1}$ there exists a pair of partially selective hybrids $\hat{\mathsf{H}}_{i,0}, \hat{\mathsf{H}}_{i,1}$ (see Figure 2) in which the adversary commits to much less information $h_i(w) \in \mathcal{U}$ instead of $w \in \mathcal{W}$. The selectivised game essentially ignores w and only relies on the partial information $u = h_i(w)$ during the course of the game but at the very end it still checks that the full commitment w matches what actually happened during the game.

Definition 11 (Partially selectivised game [JKK+17]). *Assume $\hat{\mathsf{H}}$ is a game which expects to receive some commitment $u \in \mathcal{U}$ from the adversary in the beginning. Given functions $g\colon \{0,1\}^* \to \mathcal{W}$ and $h\colon \mathcal{W} \to \mathcal{U}$ the partially selectivised game $\mathsf{H} = \mathsf{SEL}_{\mathcal{U}\to\mathcal{W}}[\hat{\mathsf{H}}, g, h]$ is defined as follows. The adversary A first sends a commitment $w \in \mathcal{W}$ to H and H begins running $\hat{\mathsf{H}}$ and passes it $u = h(w)$. It then continues running the game between $\hat{\mathsf{H}}$ and A at the end of which $\hat{\mathsf{H}}$ outputs a bit $\hat{b}$. Let* `transcript` *denote all communication exchanged between $\hat{\mathsf{H}}$ and A. If $g($*`transcript`*$) = w$ then H outputs the bit $\hat{b}$ and else it outputs* 0.

Note that different pairs of partially selective hybrids $\hat{\mathsf{H}}_{i,0}, \hat{\mathsf{H}}_{i,1}$ might rely on completely different partial information $h_i(w)$ about the adversary's choices. The partially selective hybrid associated to each H_i can thus be different when we compare $\mathsf{H}_{i-1}, \mathsf{H}_i$ (in which case it is $\hat{\mathsf{H}}_{i-1,1}$) and when we compare H_i and H_{i+1} (in which case it is $\hat{\mathsf{H}}_{i,0}$) – see Figure 2. The next theorem shows that we only incur a security loss proportional to $|\mathcal{U}|$ rather than $|\mathcal{W}|$ if we can define a sequence of partially selective hybrids which only require commitments from $\mathcal{U}$.

Theorem 4 ([JKK+17]). *Let* G^0 *and* G^1 *be two adaptive games. For some function* $g\colon \{0,1\}^* \to \mathcal{W}$ *we define the selectivised games* $\mathsf{H}^0 = \mathsf{SEL}_{\mathcal{W}}[\mathsf{G}^0, g]$, $\mathsf{H}^1 = \mathsf{SEL}_{\mathcal{W}}[\mathsf{G}^1, g]$. *Let* $\mathsf{H}^0 = \mathsf{H}_0, \mathsf{H}_1, \ldots, \mathsf{H}_\tau = \mathsf{H}^1$ *be some sequence of hybrid games. Assume that for each* $i \in [0, \tau - 1]$, *there exists a function* $h_i\colon \mathcal{W} \to \mathcal{U}$ *and games* $\hat{\mathsf{H}}_{i,0}, \hat{\mathsf{H}}_{i,1}$ *such that*

$$\mathsf{H}_i \equiv \mathsf{SEL}_{\mathcal{U}\to\mathcal{W}}[\hat{\mathsf{H}}_{i,0}, g, h_i] \quad \textit{and} \quad \mathsf{H}_{i+1} \equiv \mathsf{SEL}_{\mathcal{U}\to\mathcal{W}}[\hat{\mathsf{H}}_{i,1}, g, h_i]. \tag{1}$$

Furthermore, if $\hat{\mathsf{H}}_{i,0}, \hat{\mathsf{H}}_{i,1}$ *are* (s, ϵ)*-indistinguishable for all* $i \in [0, \tau - 1]$, *then* G^0 *and* G^1 *are* $(s - s_{\mathcal{U}}, \epsilon \cdot \tau \cdot |\mathcal{U}|)$*-indistinguishable.*

3.3 Pebbling Games

The reversible pebbling game on DAGs was introduced in [Ben89] to model reversible computation. We define a variant in which the pebbling rules have been adapted for application to PREs. In particular, the rule is the opposite of that in [Ben89]: a pebble can be placed on or removed from a vertex if all its *children* carry a pebble.[5]

Definition 12. *A reversible pebbling of a directed acyclic graph* $G = (\mathcal{V}, \mathcal{E})$ *with a unique source vertex* i^* *is a sequence* $\mathcal{P} := (\mathcal{P}_0, \ldots, \mathcal{P}_\tau)$ *of pebbling configurations* $\mathcal{P}_t \subseteq \mathcal{V}$. *Two subsequent configurations differ only in one vertex and the following rule is respected in a move: a pebble can be placed on or removed from a vertex iff all its children carry a pebble. That is,* $\mathcal{P}$ *is a valid sequence iff*

$$\forall t \in [\tau]\ \exists!\, i \in \mathcal{P}_{t-1} \triangle \mathcal{P}_t \quad \textit{and} \quad \mathsf{children}(i, G) \subseteq \mathcal{P}_{t-1}.$$

Starting with an empty graph (i.e., $\mathcal{P}_0 = \emptyset$*), the goal of the game is to place a pebble on the source (i.e.,* $i^* \in \mathcal{P}_\tau$*).*

For a DAG G, let $\mathcal{P}_G$ denote the set of all valid reversible pebbling sequences (as per Definition 12) for G. The *time complexity* of a *particular* sequence $\mathcal{P} = (\mathcal{P}_0, \ldots, \mathcal{P}_\tau)$ for a DAG G is defined as $\tau_G(\mathcal{P}) := \tau$, whereas its space complexity is defined as

$$\sigma_G(\mathcal{P}) := \max_{t \in [0,\tau]} |\mathcal{P}_t|.$$

Definition 13 (Space- and time-complexity of a class of DAGs). *We say that a class of DAGs* $\mathcal{G}$ *has time complexity* τ *and space complexity* σ *if*

$$\forall G \in \mathcal{G}\ \exists \mathcal{P} \in \mathcal{P}_G : \tau_G(\mathcal{P}) \leq \tau \wedge \sigma_G(\mathcal{P}) \leq \sigma.$$

Concrete Bounds. The pebbling complexity for the pertinent classes of *single-source* graphs on n vertices are listed in Table 2. These bounds are proved in Lemmata 2 through 4 in the full version.

[5] Alternatively, one can think of the pebbling game in Definition 12 as the classical reversible pebbling game played on a DAG whose edges have their direction flipped.

Table 2: Space and time complexity for different classes of DAGs. These bounds are proved in Lemmata 2 through 4 in the full version.

Family	Bounds	
	Space (σ)	Time (τ)
DAGs with outdegree δ and depth d $\mathcal{G}(n, \delta, d)$	$(\delta + 1) \cdot d$	$(2\delta)^d$
Complete binary trees of size n $\mathcal{B}(n)$	$\log n$	n^2
Chains of length $\mathcal{C}(n)$	$\log n + 1$	$3^{\log n}$

4 Framework for Adaptive Security

In this section we demonstrate, using the framework of [JKK+17], how adaptive security can be achieved for PREs. We focus on CPA, and the analogous result for HRA is given in the full version [FKKP18]. As for the applications given in [JKK+17], we use pebbling games on DAGs to design the hybrid games. Each pebbling configuration uniquely determines a hybrid game bridging the two real games CPA^0 and CPA^1. The DAG that we pebble in the proof is the subgraph of the recoding graph that is reachable from the challenge i^* (via the edges $\mathcal{E}$ defined during the game); it is thus a subgraph of the recoding graph with one unique source i^*, which we call the *challenge graph*. A pebble on a vertex allows the simulation of the hybrid to be carried out *without* the knowledge of the secret key associated with that vertex. The pebbling rules will ensure that hybrids corresponding to two successive pebbling configurations can be proven indistinguishable assuming key-privacy.

4.1 Adaptive Security Against Chosen-Plaintext Attack

We first show how a pebbling sequence on the challenge graph defines a sequence of fully selective hybrids (Lemma 2), and then prove that these hybrids are partially selectivised (Lemma 3).

4.1.1 Fully Selective Hybrids. In the fully selectivised version of PRE-CPA (Game 3), A first makes a commitment $\hat{G}$ to the challenge graph. Any correct commitment $\hat{G}$ must therefore have one unique source, which we denote by $\hat{\imath}$. The selective challenger is thus $\mathsf{SEL}_{\mathcal{G}}[\mathsf{CPA}^b, g]$, where g is the function that extracts the recoding graph G and the challenge user i^* from the transcript and returns the challenge graph, i.e., the subgraph of G reachable from i^*. Note that this is fundamentally different from the original selective game (i.e., sCPA in Game 1) where the adversary commits, beforehand, to *the set of corrupt public keys*.

Each hybrid is associated with a pebbling configuration $\mathcal{P}_t$ and a bit b, and we consider the sequence of hybrids $\mathsf{H}^0_0, \ldots, \mathsf{H}^0_\tau, \mathsf{H}^1_\tau, \ldots, \mathsf{H}^1_0$. The pebbling state of a vertex dictates how the outgoing re-key and re-encrypt queries are simulated,

whereas the bit determines how the challenge query is answered. To be precise, in game H_t^b, for each pebbled vertex in $\mathcal{P}_t$ all used re-keys outgoing from that vertex are faked, and the challenge query is answered by an encryption of m_b^*. (Rekeys outgoing from pebbled vertices that are not used for any queries are defined as real re-keys.) Observe that the secret key corresponding to a vertex is used only for the generation of the re-keys outgoing from that vertex; the simulation of a hybrid can thus be carried out *without* knowledge of the secret keys corresponding to the pebbled vertices (as the non-queried re-keys need not be generated).

Since the initial pebbling configuration is the empty set, H_0^0 and H_0^1 correspond to the (fully selectivised) games $\mathsf{SEL}_{\mathcal{G}}[\mathsf{CPA}^0, g]$ and $\mathsf{SEL}_{\mathcal{G}}[\mathsf{CPA}^1, g]$, respectively. Now, consider the middle hybrids H_τ^0 and H_τ^1: they are the same except for the response to the challenge query which is the encryption of m_0^* in the former and the encryption of m_1^* in the latter. Since the pebbling configuration $\mathcal{P}_\tau$, by definition, contains a pebble on the challenge vertex i^*, the simulation of this hybrid can be carried out without knowledge of the secret key corresponding to i^*. This means we can reduce indistinguishability of the PRE to the indistinguishability of these two hybrids. To be precise, the reduction embeds the challenge public key at $\hat{\imath}$, which is defined by the commitment $\hat{G}$ and replies to the challenge query (in the CPA game) by sending the challenge ciphertext (of the indistinguishability game). Note that if $i^* \neq \hat{\imath}$, that is, the commitment $\hat{G}$ doesn't coincide with the transcript of the CPA game, then the hybrid returns 0 anyway. The reduction is formally defined in Algorithm 2.

Next, consider any two hybrids H_t^b and H_{t+1}^b, $t \in [0, \tau - 1]$ and $b \in \{0, 1\}$. Also, assume $\mathcal{P}_{t+1}$ results from $\mathcal{P}_t$ by placing a pebble on the vertex i_0 (the case when a pebble is removed can be argued analogously). The simulation of H_t^b and H_{t+1}^b is the same except for the (used) re-keys outgoing from i_0: in H_t^b they are all real whereas in H_{t+1}^b they are all fake. By the rules of the pebbling game, the children of i_0 all carry pebbles in the configurations $\mathcal{P}_t$ and $\mathcal{P}_{t+1}$; therefore the simulation need not know the corresponding secret keys. This means that we can prove indistinguishability of H_t^b and H_{t+1}^b from weak key-privacy: the reduction embeds the (key-privacy) challenge public keys $\mathtt{pk}_0, \ldots, \mathtt{pk}_\delta$ at i_0 and its children, and uses the challenge re-keys $\mathtt{rk}_{0,1}, \ldots, \mathtt{rk}_{0,\delta}$ to simulate the re-key oracle for queries from i_0 to its children. The reduction is formally defined in Algorithm 3. (Note that the simulation of the reduction in Algorithm 3 is perfect: if the commitment $\hat{G}$ does not match with the transcript, it returns 0; else, we have $\hat{\imath} = i^*$ and by definition of the pebbling, i_0 is reachable from $\hat{\imath} = i^*$ and so are its children $i_1, \ldots, i_\delta$. If the adversary corrupts any of these, then the game returns 0.)

In summary, we get a sequence of hybrids $\mathsf{SEL}_{\mathcal{G}}[\mathsf{CPA}^0, g] = \mathsf{H}_0^0, \ldots, \mathsf{H}_\tau^0, \mathsf{H}_\tau^1, \ldots, \mathsf{H}_0^1 = \mathsf{SEL}_{\mathcal{G}}[\mathsf{CPA}^1, g]$, where each pair of subsequent hybrids can be proved indistinguishable. Security in the fully selectivised CPA game follows by Lemma 1. We state this formally in Lemma 2 below.

Lemma 2 (Security against fully selectivised PRE-CPA). *Consider the sequence of hybrids* $\mathsf{H}_0^0, \ldots, \mathsf{H}_\tau^0, \mathsf{H}_\tau^1, \ldots, \mathsf{H}_0^1$, *where* H_t^b *is defined in Algorithm 1*

Hybrid $\mathsf{H}^b_t(1^\kappa, 1^\lambda, n)$

1: Obtain the challenge graph $\hat{G} \in \mathcal{G}(n, \delta, d)$ from A
2: Compute $\mathcal{P}_t \leftarrow \mathsf{P}(\hat{G}, t)$ ▷ The t-th pebbling configuration
3: Set $\mathcal{C}, \mathcal{E} = \emptyset$ ▷ Stores corrupt keys and issued re-keys and re-encryptions
4: $\mathtt{pp} \leftarrow \mathsf{PRE.S}(1^\kappa, 1^\lambda)$, $(\mathtt{pk}_1, \mathtt{sk}_1), \ldots, (\mathtt{pk}_n, \mathtt{sk}_n) \leftarrow \mathsf{PRE.K}(\mathtt{pp})$
5: $\forall i \in \mathcal{P}_t, \forall j \in \mathsf{children}(i, \hat{G})$: $\mathtt{rk}_{i,j} \leftarrow \mathsf{RK}^*(\mathtt{pp}, \mathtt{pk}_j)$ ▷ Fake re-keys
6: $\forall i \in \mathcal{P}_t, \forall j \in [n] \setminus \{\mathsf{children}(i, \hat{G}) \cup i\}$: $\mathtt{rk}_{i,j} \leftarrow \mathsf{PRE.RK}((\mathtt{pk}_i, \mathtt{sk}_i), \mathtt{pk}_j)$ ▷ Real re-keys
7: $\forall i \in [n] \setminus \mathcal{P}_t, \forall j \neq i$: $\mathtt{rk}_{i,j} \leftarrow \mathsf{PRE.RK}((\mathtt{pk}_i, \mathtt{sk}_i), \mathtt{pk}_j)$ ▷ Real re-keys
8: $b' \leftarrow \mathsf{A}^{(\mathtt{corrupt},\cdot),(\mathtt{rekey},\cdot,\cdot),(\mathtt{reencrypt},\cdot,\cdot,\cdot),(\mathtt{challenge},\cdot,\cdot,\cdot)}(\mathtt{pp}, \mathtt{pk}_1, \ldots, \mathtt{pk}_n)$
9: **if** A made call $(\mathtt{challenge}, i^*, \cdot, \cdot)$ for some i^* **then** ▷ Check abort conditions
10: **if** $\exists\, i \in \mathcal{C} : i^*$ is connected to i in $([n], \mathcal{E})$ **then return** 0 **end if**
11: **end if**
12: **if** $\hat{G}$ is the subgraph of $([n], \mathcal{E})$ reachable from i^* **then return** b' **end if**
13: **return** 0

Algorithm 1: Template for generating fully selective PRE-CPA hybrids given a pebbling configuration. All the oracles are defined like in Game 3.

using the pebbling configuration $\mathcal{P}_t$. H^b_0 is the fully selectivised game of CPA^b: i.e., $\mathsf{H}^b_0 = \mathsf{SEL}_{\mathcal{G}}[\mathsf{CPA}^b, g]$ where g extracts the challenge graph (subgraph reachable from the challenge vertex) from a transcript. Moreover, if the adversary makes at most Q_{RE} re-encryption queries, then a PRE scheme that is (s_1, ϵ_1)-indistinguishable and $(s_2, \epsilon_2, \delta)$-weakly key-private is (s, ϵ)-secure against fully selectivised PRE-CPA restricted to challenge graphs in $\mathcal{G}(n, \delta, d)$ with

$$s := \min(s_1, s_2) - s_{\mathsf{CPA}} \quad \textit{and} \quad \epsilon := \epsilon_1 + 2\tau \cdot \epsilon_2,$$

where $s_{\mathsf{CPA}} \approx O(s_{\mathsf{P}} + n^2 \cdot s_{\mathsf{RK}} + Q_{\mathsf{RE}} \cdot s_{\mathsf{RE}})$ denotes the complexity of simulating the CPA game.

PRE-CPA-security follows from random guessing (Theorem 3) but with a security loss of 2^{n^2}, where n^2 is an upper bound on the number of bits required to encode the challenge subgraph:

Corollary 1 (PRE-CPA-security by random guessing). *A PRE scheme that is (s_1, ϵ_1)-indistinguishable and $(s_2, \epsilon_2, \delta)$-weakly key-private is (s, ϵ)-secure against PRE-CPA restricted to challenge graphs in $\mathcal{G}(n, \delta, d)$, where*

$$s := \min(s_1, s_2) - s_{\mathsf{CPA}} - s_{\mathcal{G}} \quad \textit{and} \quad \epsilon := (\epsilon_1 + 2\tau \cdot \epsilon_2) \cdot 2^{n^2}.$$

4.1.2 Partially Selective Hybrids. In hybrid H^b_t described in Algorithm 1, we observe that not all information on the committed recoding graph $\hat{G}$ is actually required for the simulation. In fact, only the pebbling configuration $\mathcal{P}_t$ is required to simulate the hybrid: re-keys are only required once a corresponding re-key or a re-encrypt query is issued; for a pebbled node, such queries lead to

Reduction $\mathsf{R}_\tau^{(\mathsf{IND.challenge},\cdot,\cdot)}(\mathtt{pp}^*, \mathtt{pk}^*)$ ▷ $\mathtt{pk}^*$ denotes the challenge public key
1: Obtain the challenge graph $\hat{G} \in \mathcal{G}(n, \delta, d)$ from A
2: Compute $\mathcal{P}_\tau \leftarrow \mathsf{P}(\hat{G}, \tau)$ ▷ The τ-th pebbling configuration
3: Set $\mathcal{C}, \mathcal{E} = \emptyset$ ▷ Stores corrupt keys and issued re-keys and re-encryptions
4: $(\mathtt{pk}_1, \mathtt{sk}_1), \ldots, (\mathtt{pk}_{\hat{i}-1}, \mathtt{sk}_{\hat{i}-1}), (\mathtt{pk}_{\hat{i}+1}, \mathtt{sk}_{\hat{i}+1}), \ldots, (\mathtt{pk}_n, \mathtt{sk}_n) \leftarrow \mathsf{PRE.K}(\mathtt{pp}^*)$
5: Let $\hat{i}$ be the source of $\hat{G}$, set $\mathtt{pk}_{\hat{i}} := \mathtt{pk}^*$ ▷ Embed challenge public key
6: $\forall i \in \mathcal{P}_\tau, \forall j \in \mathsf{children}(i, \hat{G})$: $\mathtt{rk}_{i,j} \leftarrow \mathsf{RK}^*(\mathtt{pp}, \mathtt{pk}_j)$ ▷ Fake re-keys
7: $\forall i \in [n] \setminus \mathcal{P}_\tau, \forall j \in \mathsf{children}(i, \hat{G})$: $\mathtt{rk}_{i,j} \leftarrow \mathsf{PRE.RK}((\mathtt{pk}_i, \mathtt{sk}_i), \mathtt{pk}_j)$ ▷ Real re-keys
8: $\forall i \in \mathcal{P}_\tau, \forall j \in [n] \setminus \{\mathsf{children}(i, \hat{G}) \cup i\}$: $\mathtt{rk}_{i,j} \leftarrow \mathsf{PRE.RK}((\mathtt{pk}_i, \mathtt{sk}_i), \mathtt{pk}_j)$ ▷ Real re-keys
9: $\forall i \in [n] \setminus \mathcal{P}_\tau, \forall j \neq i$: $\mathtt{rk}_{i,j} \leftarrow \mathsf{PRE.RK}((\mathtt{pk}_i, \mathtt{sk}_i), \mathtt{pk}_j)$ ▷ Real re-keys
10: $b' \leftarrow \mathsf{A}^{(\mathtt{corrupt},\cdot),(\mathtt{rekey},\cdot,\cdot),(\mathtt{reencrypt},\cdot,\cdot,\cdot),(\mathtt{challenge},\cdot,\cdot,\cdot)}(\mathtt{pp}^*, \mathtt{pk}_1, \ldots, \mathtt{pk}_n)$
11: **if** A made call $(\mathtt{challenge}, i^*, \cdot, \cdot)$ for some i^* **then** ▷ Check abort conditions
12: **if** $\exists\, i \in \mathcal{C} : i^*$ is connected to i in $([n], \mathcal{E})$ **then return** 0 **end if**
13: **end if**
14: **if** $\hat{G}$ is the subgraph of $([n], \mathcal{E})$ reachable from i^* **then return** b' **end if**
15: **return** 0

Oracles `rekey` and `reencrypt` are defined like in Game 3.

Oracle $(\mathtt{corrupt}, i)$
1: **if** $i = \hat{i}$ **then** HALT: R_τ returns 0 **end if** ▷ Commitment $\hat{G}$ doesn't match or...
2: Add i to $\mathcal{C}$ and **return** $\mathtt{sk}_i$ ▷ ... i^* corrupted

Oracle $(\mathtt{challenge}, i^*, (m_0^*, m_1^*), \ell^*)$ ▷ Single access
1: $(c_{i^*}, \ell^*) \leftarrow \mathsf{IND.challenge}((m_0^*, m_1^*), \ell^*)$ ▷ Embed challenge ciphertext
2: **return** (c_{i^*}, ℓ^*)

Algorithm 2: The reduction showing that the hybrids H_τ^0 and H_τ^1 are indistinguishable by indistinguishability of ciphertexts.

an edge added in $\mathcal{E}$; thus the re-key is simulated (while the "not-queried" re-keys are never used during the experiment).

In addition to the pebbling configuration $\mathcal{P}_\tau$, the reduction from ciphertext indistinguishability (cf. Algorithm 2) also needs to know the challenge vertex $\hat{i}$ in order to embed the challenge public key. The reduction from weak key-privacy (cf. Algorithm 3) requires, in addition to $\mathcal{P}_t$, the vertex that is pebbled or unpebbled in $\mathcal{P}_{t+1}$ (i.e., the vertex i_0) and its children, so it can embed its challenge public keys and re-keys.

To sum up, two consecutive hybrids H_t^b and H_{t+1}^b can be shown to be indistinguishable using a lot less information than what the adversary commits to. We thus have the following:

Lemma 3 (Partially selectivised hybrids). *Let $\mathcal{P}_0, \ldots, \mathcal{P}_\tau$ and $\mathsf{H}_0^0, \ldots, \mathsf{H}_\tau^0$, $\mathsf{H}_\tau^1, \ldots, \mathsf{H}_0^1$ be defined as in Lemma 2, and let σ denote the space complexity of the pebbling sequence. Then, for $t \in [0, \tau - 1]$ and $b, \beta \in \{0, 1\}$,*

$$\mathsf{H}_{t+\beta}^b \equiv \mathsf{SEL}_{\mathcal{U}\to\mathcal{G}}[\hat{\mathsf{H}}_{t,\beta}^b, g, h_t] \quad \textit{and} \quad \mathsf{H}_\tau^b \equiv \mathsf{SEL}_{\mathcal{U}\to\mathcal{G}}[\hat{\mathsf{H}}_{\tau,0}^b, g, h_\tau],$$

Reduction $\mathsf{R}_t^b(\mathtt{pp}^*, \mathtt{pk}_0^*, \ldots, \mathtt{pk}_\delta^*, \mathtt{rk}_{0,1}^*, \ldots, \mathtt{rk}_{0,\delta}^*)$

1: Obtain the challenge graph $\hat{G} \in \mathcal{G}(n, \delta, d)$ from A
2: Compute $\mathcal{P}_t \leftarrow \mathsf{P}(\hat{G}, t)$, $\mathcal{P}_{t+1} \leftarrow \mathsf{P}(\hat{G}, t+1)$ ▷ The t-th and $(t+1)$-th configurations
3: Set $\mathcal{C}, \mathcal{E} = \emptyset$ ▷ Stores corrupt keys and issued re-keys and re-encryptions
4: $i_0 := \mathcal{P}_t \Delta \mathcal{P}_{t+1}$, $i_1, \ldots, i_\delta := \mathsf{children}(i_0, \hat{G})$ ▷ i_0 denotes pebbled/unpebbled vertex
5: $\forall k \in [0, \delta]$: $\mathtt{pk}_{i_k} := \mathtt{pk}_k^*$ ▷ Embed the challenge public keys
6: $\forall k \in [n] \setminus \{i_0, \ldots, i_\delta\}$: $(\mathtt{pk}_k, \mathtt{sk}_k) \leftarrow \mathsf{PRE.K}(\mathtt{pp}^*)$ ▷ Real keys
7: $\forall k \in [\delta] : \mathtt{rk}_{i_0,i_k} := \mathtt{rk}_{0,k}^*$ ▷ Embed challenge re-keys
8: $\forall i \in \mathcal{P}_t \setminus \{i_0\}, \forall j \in \mathsf{children}(i, \hat{G})$: $\mathtt{rk}_{i,j} \leftarrow \mathsf{RK}^*(\mathtt{pp}^*, \mathtt{pk}_j)$ ▷ Fake re-keys
9: $\forall i \in \mathcal{P}_t, \forall j \in [n] \setminus \{\mathsf{children}(i, \hat{G}) \cup i\}$: $\mathtt{rk}_{i,j} \leftarrow \mathsf{PRE.RK}((\mathtt{pk}_i, \mathtt{sk}_i), \mathtt{pk}_j)$ ▷ Real re-keys
10: $\forall i \in [n] \setminus (\mathcal{P}_t \cup \{i_0\}), \forall j \neq i$: $\mathtt{rk}_{i,j} \leftarrow \mathsf{PRE.RK}((\mathtt{pk}_i, \mathtt{sk}_i), \mathtt{pk}_j)$ ▷ Real re-keys
11: $b' \leftarrow \mathsf{A}^{(\mathtt{corrupt},\cdot),(\mathtt{rekey},\cdot,\cdot),(\mathtt{reencrypt},\cdot,\cdot,\cdot),(\mathtt{challenge},\cdot,\cdot,\cdot)}(\mathtt{pp}^*, \mathtt{pk}_1, \ldots, \mathtt{pk}_n)$
12: **if** A made call $(\mathtt{challenge}, i^*, \cdot, \cdot)$ for some i^* **then** ▷ Check abort conditions
13: **if** $\exists\, i \in \mathcal{C} : i^*$ is connected to i in $([n], \mathcal{E})$ **then return** 0 **end if**
14: **end if**
15: **if** $\hat{G}$ is the subgraph of $([n], \mathcal{E})$ reachable from i^* **then return** b' **end if**
16: **return** 0

Oracles $\mathtt{rekey}$, $\mathtt{reencrypt}$ and $\mathtt{challenge}$ are defined like in Game 3.

Oracle $(\mathtt{corrupt}, i)$

1: **if** $i \in \{i_0, \ldots, i_\delta\}$ **then** HALT: R_τ returns 0 **end if** ▷ Commitment $\hat{G}$ doesn't match
2: Add i to $\mathcal{C}$ and **return** $\mathtt{sk}_i$ ▷ ... or i reachable from i^*

Algorithm 3: The reduction showing that the hybrids H_t^b and H_{t+1}^b, for $t \in [0, \tau - 1]$ and $b \in \{0, 1\}$, are indistinguishable by weak key-privacy.

where $\hat{\mathsf{H}}_{t,\beta}^b$ is defined in Algorithm 4 (see also Figure 3), g extracts the challenge graph from the transcript (as in Lemma 2). For $t \in [0, \tau - 1]$, h_t is the function that extracts the pebbling configuration $\mathcal{P}_t$, the pebbled/unpebbled vertex in $\mathcal{P}_{t+1}$ and its children; h_τ extracts the pebbling configuration $\mathcal{P}_\tau$ and the challenge node i^*. Thus, $\mathcal{U}$ corresponds to the set $\mathcal{V}^{\sigma+\delta+1}$.

The tighter bound for PRE-CPA-security now results by applying Theorem 4:

Theorem 5 (main, PRE-CPA security). *Let σ and τ denote, respectively, the pebbling space and time complexity for the class $\mathcal{G}(n, \delta, d)$. Then a PRE scheme that is (s_1, ϵ_1)-indistinguishable and $(s_2, \epsilon_2, \delta)$-weakly key-private is (s, ϵ)-PRE-CPA-secure restricted to challenge graphs in $\mathcal{G}(n, \delta, d)$, where*

$$s := \min(s_1, s_2) - s_{\mathsf{CPA}} - s_{\mathcal{G}} \quad \text{and} \quad \epsilon := (\epsilon_1 + 2\tau \cdot \epsilon_2) \cdot n^{\sigma+\delta+1}.$$

4.2 Corollaries

Finally, we calculate concrete bounds to Theorem 5 for the families of recoding graphs listed in Table 2. The approximate security loss (assuming $\epsilon_1 = \epsilon_2 = \epsilon'$)

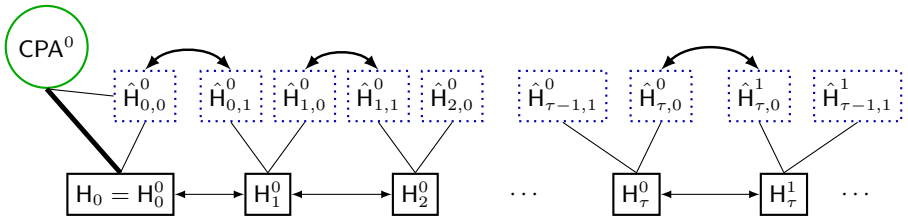

Fig. 3: Diagram showing the partially selectivised hybrids for PRE-CPA.

that results when substituting these bounds for CPA in Theorem 5 are $n^{O(d\cdot\delta)}$ for $\mathcal{G}(n,\delta,d)$, $n^{O(\log n)}$ for $\mathcal{B}(n)$ and $n^{O(\log n)}$ for $\mathcal{C}(n)$. The same bounds hold for HRA if one assumes that Q_{RE} and Q_{E} (i.e., number of queries) are polynomial and $\epsilon_3 = 2^{-\kappa}$ (i.e., the PRE scheme is *statistically* source-hiding): see the full version for the details.

5 Lattice-based Multi-Hop Scheme

Here, we describe the lattice-based unidirectional multi-hop PRE scheme from [CCL$^+$14]. The remainder of the schemes — viz. the pairing-based schemes from [AFGH05] and [ABH09] and Gentry's FHE-based construction — are given in the full version [FKKP18]. We note that being based directly on the decision LWE (DLWE) problem this scheme achieves better parameters than the construction from FHE.

5.1 [CCL$^+$14] Scheme.

In [CCL$^+$14], Chandran et al. propose two lattice-based unidirectional multi-hop proxy re-encryption schemes. The schemes are built upon Regev's encryption [Reg05] and its dual version [GPV07], respectively. Here, we will describe the former one, which is inspired by the fully homomorphic encryption scheme of [BV11]. Security can be proven assuming the hardness of the decisional learning with errors (DLWE) problem (cf. Definition 14 below).

We recall Regev's encryption scheme in Construction 1. We can now define the PRE scheme from [CCL$^+$14] using RGV in Construction 2.a. To achieve source-hiding, Chandran et al. propose the variant given as Construction 2.b.

In both schemes, the LWE error will grow with each re-encryption and the level bound λ needs to be chosen appropriately so that correctness of decryption is still guaranteed (with overwhelming probability). The second variant achieves the stronger notion of PRE-HRA-security (see below) at the cost of worse parameters; only a small number λ of re-encryptions is supported by this scheme and the underlying security assumption is very strong.

Security. The PRE scheme in Construction 2.a can be proven secure assuming the hardness of decisional learning with errors (DLWE). We will first show

Hybrid $\mathsf{H}^b_{t+\beta}$

Obtain the challenge graph $\hat{G} \in \mathcal{G}(n,\delta,d)$ from A and let $\hat{\imath}$ be its source
Compute $\mathcal{P}_t \leftarrow \mathsf{P}(\hat{G},t)$, $\mathcal{P}_{t+1} \leftarrow \mathsf{P}(\hat{G},t+1)$ ▷ The t-th and $(t+1)$-th configurations
$i_0 := \mathcal{P}_t \Delta \mathcal{P}_{t+1}$, $i_1,\dots,i_\delta := \mathsf{children}(i_0,\hat{G})$ ▷ i_0 denotes pebbled/unpebbled vertex
if $t < \tau$ **then** $\hat{b} \leftarrow \hat{\mathsf{H}}^b_{t,\beta}(\mathcal{P}_t,\{i_0,\dots,i_\delta\})$ ▷ Key-privacy hybrid
else $\hat{b} \leftarrow \hat{\mathsf{H}}^b_{\tau,0}(\mathcal{P}_\tau,\{\hat{\imath},\bot,\dots,\bot\})$ **end if** ▷ $t=\tau$, $\beta=0$: ind hybrid
if $\hat{G}$ is subgraph of $([n],\mathcal{E})$ reachable from i^* **then return** b' **end if**
return 0

$\hat{\mathsf{H}}^b_{t,\beta}(\mathcal{P}_t,\{i_0,\dots,i_\delta\})$

Set $\mathcal{C},\mathcal{E} = \emptyset$ ▷ Stores corrupt keys and issued re-keys and re-encryptions
if $t < \tau$ **then**
  if $i_0 \in \mathcal{P}_t$ **then** $\mathcal{P}_{t+1} := \mathcal{P}_t \setminus \{i_0\}$ **else** $\mathcal{P}_{t+1} := \mathcal{P}_t \cup \{i_0\}$ **end if**
end if
$\mathtt{pp} \leftarrow \mathsf{PRE.S}(1^\kappa,1^\lambda)$, $(\mathtt{pk}_1,\mathtt{sk}_1),\dots,(\mathtt{pk}_n,\mathtt{sk}_n) \leftarrow \mathsf{PRE.K}(\mathtt{pp})$
$\forall i,j \in [n], i \neq j : \mathtt{rk}_{i,j} = \bot$ ▷ Delay re-key generation till the query
$b' \leftarrow \mathsf{A}^{(\mathtt{corrupt},\cdot),(\mathtt{rekey},\cdot,\cdot),(\mathtt{reencrypt},\cdot,\cdot,\cdot),(\mathtt{challenge},\cdot,\cdot,\cdot)}(\mathtt{pp},\mathtt{pk}_1,\dots,\mathtt{pk}_n)$
if A made call $(\mathtt{challenge},i^*,\cdot,\cdot)$ for some i^* **then** ▷ Check abort conditions
  if $\exists\, i \in \mathcal{C} : i^*$ is connected to i in $([n],\mathcal{E})$ **then return** 0 **end if**
end if
return b'

Oracles `corrupt` and `challenge` are defined like in Game 3.

Oracle $(\mathtt{rekey},i,j)$

if $\mathtt{rk}_{i,j} = \bot$ **then** ▷ Re-key not generated
  if $i \in \mathcal{P}_{t+\beta}$ **then** $\mathtt{rk}_{i,j} \leftarrow \mathsf{RK}^*(\mathtt{pp},\mathtt{pk}_j)$ ▷ Fake re-key
  else $\mathtt{rk}_{i,j} \leftarrow \mathsf{RK}((\mathtt{pk}_i,\mathtt{sk}_i),\mathtt{pk}_j)$ **end if** ▷ Real re-key
end if
Add (i,j) to $\mathcal{E}$ ▷ Add to recoding graph
return $\mathtt{rk}_{i,j}$

Oracle $(\mathtt{reencrypt},i,j,(c_i,\ell))$

if $\mathtt{rk}_{i,j} = \bot$ **then** ▷ Re-key not generated
  if $i \in \mathcal{P}_{t+\beta}$ **then** $\mathtt{rk}_{i,j} \leftarrow \mathsf{RK}^*(\mathtt{pp},\mathtt{pk}_j)$ ▷ Fake re-key
  else $\mathtt{rk}_{i,j} \leftarrow \mathsf{RK}((\mathtt{pk}_i,\mathtt{sk}_i),\mathtt{pk}_j)$ **end if** ▷ Real re-key
end if
Add (i,j) to $\mathcal{E}$ ▷ Add to recoding graph
return $(c_j,\ell+1) \leftarrow \mathsf{PRE.RE}(\mathtt{rk}_{i,j},\mathtt{pk}_i,\mathtt{pk}_j,(c_i,\ell))$

Algorithm 4: Partially selectivised hybrids. For $t \in [0,\tau-1]$ and $b,\beta \in \{0,1\}$: $\mathsf{H}^b_{t+\beta} = \mathsf{SEL}_{\mathcal{U}\to\mathcal{G}}[\hat{\mathsf{H}}^b_{t,\beta},g,h_t]$ and $\mathsf{H}^b_\tau = \mathsf{SEL}_{\mathcal{U}\to\mathcal{G}}[\hat{\mathsf{H}}^b_{\tau,0},g,h_\tau]$. Moreover, $\mathcal{U}$ is the set $\mathcal{V}^{\sigma+\delta+1}$. Note that the sampling of the re-keys has been deferred to the actual calls.

PRE-CPA-security of Construction 2.a and then consider PRE-HRA-security of Construction 2.b.

1. $\mathsf{S}(1^\kappa)$: Pick lattice parameters N, M, $q \in \mathbb{N}$ and a B-bounded error distribution χ on $\mathbb{Z}_q^M$. Sample $\boldsymbol{A} \leftarrow \mathbb{Z}_q^{M \times N}$ uniformly at random and return the public parameters $\mathtt{pp} = (\boldsymbol{A}, N, M, q, \chi)$.
2. $\mathsf{K}(\mathtt{pp})$: Sample $\boldsymbol{s} \leftarrow \mathbb{Z}_q^N$ uniformly at random and compute $\boldsymbol{b} = \boldsymbol{A} \cdot \boldsymbol{s} + \boldsymbol{e}$, where $\boldsymbol{e} \leftarrow \chi$. Set $\mathtt{pk} := \boldsymbol{b}$ as the public key and $\mathtt{sk} := \boldsymbol{s}$ as the secret key. Return $(\mathtt{pk}, \mathtt{sk})$.
3. $\mathsf{E}(\mathtt{pk}, m)$: On input $\mathtt{pk} \in \mathbb{Z}_q^M$ and a message bit $m \in \{0,1\}$, sample $\boldsymbol{r} \leftarrow \{0,1\}^M$ and output
$$\boldsymbol{c} := \boldsymbol{r}^T(\boldsymbol{A}, \boldsymbol{b}) + (0^N, m \cdot \lceil q/2 \rfloor) \in \mathbb{Z}_q^{N+1}.$$
4. $\mathsf{D}(\mathtt{sk}, \boldsymbol{c})$: On input a secret key $\mathtt{sk} = \boldsymbol{s} \in \mathbb{Z}_q^N$ and a ciphertext $\boldsymbol{c} = (\boldsymbol{\alpha}, \beta) \in \mathbb{Z}_q^N \times \mathbb{Z}_q$, output 0 if $\beta - \langle \boldsymbol{\alpha}, \boldsymbol{s} \rangle$ is closer to 0 than to $\lceil q/2 \rfloor$, else output 1.

Construction 1: Regev's Encryption scheme RGV [Reg05].

1. $\mathsf{S}(1^\kappa)$: Get parameters $\mathtt{pp}' \leftarrow \mathsf{RGV.S}(1^\kappa)$, level bound λ and "blurring error" bound E_ℓ for each level $\ell \in [\lambda]$. Return the parameters $\mathtt{pp} = (\mathtt{pp}', \lambda, (E_\ell)_{\ell \in [\lambda]})$.
2. $\mathsf{K}(\mathtt{pp})$: Run $\mathsf{RGV.K}(\mathtt{pp}')$ and output the result.
3. $\mathsf{E}(\mathtt{pk}, (m, \ell))$: Compute $\boldsymbol{c} = \mathsf{RGV.E}(\mathtt{pk}, m) + (0^N, f_\ell)$, where $f_\ell \leftarrow [-E_\ell, E_\ell] \cap \mathbb{Z}$, and return the level-$\ell$ ciphertext $(\boldsymbol{c}, \ell)$.
4. $\mathsf{RK}((\mathtt{pk}_i, \mathtt{sk}_i), \mathtt{pk}_j)$: Parse $\mathtt{sk}_i$ as $\boldsymbol{s}_i = (s_{i,1}, \ldots, s_{i,N}) \in \mathbb{Z}_q^N$. For $k \in [N]$ and $l \in [\lceil \log q \rceil]$, compute $K_{k,l} \leftarrow \mathsf{RGV.E}(\mathtt{pk}_j, 0) + (0^N, s_{i,k} \cdot 2^l)$. Return the re-key:
$$\mathtt{rk}_{i,j} := \{K_{k,l}\}_{k \in [N], l \in [\lceil \log q \rceil]}.$$
5. $\mathsf{RE}(\mathtt{rk}_{i,j}, \mathtt{pk}_i, \mathtt{pk}_j, (\boldsymbol{c}_i, \ell)) \to (c_j, \ell + 1)$: If $\ell \geq \lambda$, abort. Otherwise, parse the level-ℓ ciphertext $\boldsymbol{c}_i$ as $(\boldsymbol{\alpha}, \beta) \in \mathbb{Z}_q^N \times \mathbb{Z}_q$ and $\mathtt{rk}_{i,j}$ as $\{K_{k,l}\}_{k \in [N], l \in [\lceil \log q \rceil]}$. Denote by α_k the k-th component of $\boldsymbol{\alpha}$, and denote the bit decomposition of α_k as $\{\alpha_{k,l}\}_{l \in [\lceil \log q \rceil]}$, i.e., $\alpha_k = \sum_{l \in [\lceil \log q \rceil]} \alpha_{k,l} 2^l$, where each $\alpha_{k,l} \in \{0,1\}$. Compute
$$\boldsymbol{c}_j = (0^N, \beta) + \textstyle\sum_{k,l} \alpha_{k,l} \cdot K_{k,l} + \mathsf{RGV.E}(\mathtt{pk}_j, 0) + (0^N, f_{\ell+1}),$$
where $f_{\ell+1} \leftarrow [-E_{\ell+1}, E_{\ell+1}] \cap \mathbb{Z}$, and return $(\boldsymbol{c}_j, \ell + 1)$.
6. $\mathsf{D}(\mathtt{sk}, (\boldsymbol{c}, \ell))$: Run $\mathsf{RGV.D}(\mathtt{sk}, \boldsymbol{c})$ and output the result.

Construction 2: source-hiding unidirectional multi-hop PRE from [CCL+14]. We refer to the construction without the blurring (ignoring the boxes) by Construction 2.a and the construction with blurring (including the boxes) by Construction 2.b.

Definition 14 (DLWE [Reg05]). *Let N, M, $q \in \mathbb{N}$. For a matrix $\boldsymbol{A} \leftarrow \mathbb{Z}_q^{M \times N}$ and a secret vector $\boldsymbol{s} \leftarrow \mathbb{Z}_q^N$, each sampled uniformly at random, and a vector $\boldsymbol{e} \leftarrow \chi$ for an error distribution χ on $\mathbb{Z}_q^M$, the decisional LWE problem*

$\mathrm{DLWE}_{N,M,q,\chi}$ *is to distinguish* $(\boldsymbol{A}, \boldsymbol{A}\cdot\boldsymbol{s}+\boldsymbol{e})$ *from* $(\boldsymbol{A},\boldsymbol{b})$ *for a uniformly random sample* $\boldsymbol{b}\leftarrow\mathbb{Z}_q^M$.

To prove adaptive security for the two variants of Construction 2, we will need the following lemma [BV11].

Lemma 4 (matrix-vector leftover hash lemma). *Let* κ, N, $q\in\mathbb{N}$, *and* $M\geq N\cdot\log q+2\kappa$. *For* $\boldsymbol{A}\leftarrow\mathbb{Z}_q^{M\times N}$, $\boldsymbol{r}\leftarrow\{0,1\}^M$, *and* $\boldsymbol{y}\leftarrow\mathbb{Z}_q^N$ *each sampled uniformly at random, it holds* $\Delta((\boldsymbol{A},\boldsymbol{A}^T\boldsymbol{r}),(\boldsymbol{A},\boldsymbol{y}))\leq 2^{-\kappa}$.

Assuming the computational hardness of $\mathrm{DLWE}_{N,M,q,\chi}$ for appropriate parameters, by the above lemma we get for any $\mathtt{pp}=(\boldsymbol{A},N,M,q,\chi,\lambda)$, $\mathtt{pk}=\boldsymbol{b}$ and $m\in\{0,1\}$: $\mathsf{RGV.E}(\mathtt{pk},m)=\boldsymbol{r}^T(\boldsymbol{A},\boldsymbol{b})+(0^N,m\cdot\lceil q/2\rfloor)$ is computationally indistinguishable from $\boldsymbol{r}^T(\boldsymbol{A},\boldsymbol{b}')+(0^N,m\cdot\lceil q/2\rfloor)$, where $\boldsymbol{r}\leftarrow\{0,1\}^M$ and $\boldsymbol{b}'\leftarrow\mathbb{Z}_q^M$. The latter distribution is, in turn, statistically close to the uniform distribution on $\mathbb{Z}_q^{N+1}$. Informally, since $\mathsf{RGV.E}(\mathtt{pk},0)$ is computationally indistinguishable from uniformly random, ciphertexts, re-keys and re-encrypted ciphertexts all look uniformly random; in particular Construction 2.a satisfies indistinguishability of ciphertexts as well as δ-weak key privacy.

Lemma 5. *Assuming* $\mathrm{DLWE}_{N,M,q,\chi}$ *is* (s_1,ϵ_1)*-hard for parameters* N,M,q *as in Lemma 4, Construction 2.a satisfies* $(s_1-s_{\mathsf{E}},2(\epsilon_1+2^{-\kappa}))$*-indistinguishability and* $(s_1-O(\delta\,N\lceil\log q\rceil\,(s_{\mathbb{Z}_q^{N+1}}+s_{\mathsf{RGV.E}})),\delta N\lceil\log q\rceil\epsilon_1,\delta)$*-weak key-privacy.*

Theorem 6 (PRE-CPA-security of Construction 2.a). *Let* σ *and* τ *denote the space and time complexity for the class* $\mathcal{G}=\mathcal{G}(n,\delta,d)$. *Assume the* $\mathrm{DLWE}_{N,M,q,\chi}$ *problem is* (s_1,ϵ_1)*-hard for parameters* N,M,q *as in Lemma 4. Then Construction 2.a is* (s,ϵ)*-PRE-CPA-secure restricted to challenge graphs in* $\mathcal{G}$, *where*

$$s:=s_1-O(\delta\,N\lceil\log q\rceil\,(s_{\mathbb{Z}_q^{N+1}}+s_{\mathsf{RGV.E}}))-s_{\mathsf{CPA}}-s_{\mathcal{G}}\quad\text{and}$$
$$\epsilon:=(2\tau\cdot\delta N\lceil\log q\rceil+1)\cdot\epsilon_1\cdot n^{\sigma+\delta+1}.$$

Construction 2.a clearly does not satisfy source-hiding and, thus cannot be proven PRE-HRA-secure using our results. Fortunately, Construction 2.b solves this issue, but at the cost of only allowing for a constant level bound λ. The additional uniform error $f_\ell\leftarrow[-E_\ell,E_\ell]\cap\mathbb{Z}$ added in E and RE in Construction 2.b is used to "blur out" the different errors caused by encryption or re-encryption, respectively. Choosing the error bounds E_ℓ appropriately guarantees the source-hiding property of the scheme while still preserving correctness.[6] Chandran et al. refer to this rerandomisation technique as *strong blurring*; a more detailed analysis can be found in [DS16, Section 4.1], where the same method for rerandomization of Regev ciphertexts is used to discuss sanitizability of the FHE scheme from [BV11].

[6] In fact, we need to choose the error bounds $(E_\ell)_{\ell\in[\lambda]}$ exponentially large, eg., $E_1\geq(M+1)B2^\kappa$. Thus, to provide correctness of the scheme, one needs to choose the modulus q to be of size $\exp(O(\kappa))$ and the level bound λ of size $O(1)$.

To prove PRE-HRA-security of Construction 2.b, note that, as above, semantic security and δ-weak key-privacy of (E, D) directly follow by the security of Regev's encryption scheme. We get a result similar to Lemma 5.

Lemma 6. *For large enough (see Footnote 6) error ranges E_ℓ, $\ell \in [\lambda]$, Construction 2.b is (statistically) source-hiding.*

Theorem 7 (PRE-HRA-security of Construction 2.b). *Let ϵ be as in Theorem 6 and let σ and τ denote the space and time complexity for $\mathcal{G} = \mathcal{G}(n, \delta, d)$. If $\mathrm{DLWE}_{N,M,q,\chi}$ is (s_1, ϵ_1)-hard for parameters N, M, q as in Lemma 4 and E_ℓ ($\ell \in [\lambda]$), λ are chosen appropriately, then Construction 2.b is (s', ϵ')-PRE-HRA-secure restricted to challenge graphs in $\mathcal{G}$, where*

$$s' := s_1 - O(\delta\, N\, \lceil \log q \rceil\, (s_{\mathbb{Z}_q^{N+1}} + s_{\mathsf{RGV.E}})) - s_{\mathsf{HRA}} - s_{\mathcal{G}}, \quad \textit{and}$$
$$\epsilon' := 2n(n-1)(Q_{\mathsf{E}} + Q_{\mathsf{RE}})Q_{\mathsf{RE}} \cdot 2^{-\kappa} + \epsilon.$$

6 Open Problems

We leave as open problems to find adaptively secure PREs (either via the [JKK+17] framework or using a new technique) for more general settings, which includes unidirectional PREs on general graphs, bidirectional PREs and CCA-secure PRE (the schemes above only satisfy CPA, and the slightly stronger HRA security notion).

Acknowledgements. The first author is supported by the French ANR EfTrEC project (ANR-16-CE39-0002). The remaining authors are supported by the European Research Council, ERC consolidator grant TOCNeT (682815).

References

ABH09. Giuseppe Ateniese, Karyn Benson, and Susan Hohenberger. Key-private proxy re-encryption. In Marc Fischlin, editor, *Topics in Cryptology – CT-RSA 2009*, pages 279–294, Springer, 2009.

AFGH05. Giuseppe Ateniese, Kevin Fu, Matthew Green, and Susan Hohenberger. Improved proxy re-encryption schemes with applications to secure distributed storage. In *Proceedings of the Network and Distributed System Security Symposium, NDSS 2005, San Diego, California, USA*. The Internet Society, 2005.

BB04. Dan Boneh and Xavier Boyen. Secure identity based encryption without random oracles. In Matt Franklin, editor, *Advances in Cryptology – CRYPTO 2004*, volume 3152 of LNCS, pages 443–459. Springer, 2004.

BBS98. Matt Blaze, Gerrit Bleumer, and Martin Strauss. Divertible protocols and atomic proxy cryptography. In Kaisa Nyberg, editor, *Advances in Cryptology — EUROCRYPT'98*, pages 127–144, 1998. Springer.

Ben89. Charles H. Bennett. Time/space trade-offs for reversible computation. *SIAM Journal on Computing*, 18(4):766–776, 1989.

BV11. Z. Brakerski and V. Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 97–106, Oct 2011.

CCL+14. Nishanth Chandran, Melissa Chase, Feng-Hao Liu, Ryo Nishimaki, and Keita Xagawa. Re-encryption, functional re-encryption, and multi-hop re-encryption: A framework for achieving obfuscation-based security and instantiations from lattices. In *Public-Key Cryptography - PKC 2014*, volume 8383, pages 95–112. Springer, March 2014.

CCV12. Nishanth Chandran, Melissa Chase, and Vinod Vaikuntanathan. Functional re-encryption and collusion-resistant obfuscation. In Ronald Cramer, editor, *Theory of Cryptography*, pages 404–421, 2012. Springer.

CH07. Ran Canetti and Susan Hohenberger. Chosen-ciphertext secure proxy re-encryption. In *Proceedings of the 14th ACM Conference on Computer and Communications Security*, CCS '07, pages 185–194, 2007. ACM.

Coh17. Aloni Cohen. What about Bob? The inadequacy of CPA security for proxy reencryption. Cryptology ePrint Report 2017/785, 2017. https://ia.cr/2017/785.

DS16. Léo Ducas and Damien Stehlé. Sanitization of FHE ciphertexts. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology – EUROCRYPT 2016*, pages 294–310, 2016. Springer.

FKKP18. Georg Fuchsbauer and Chethan Kamath and Karen Klein and Krzysztof Pietrzak. Adaptively secure proxy re-encryption. Cryptology ePrint Archive, Report 2018/426, https://ia.cr/2018/426.

FKPR14. Georg Fuchsbauer and Momchil Konstantinov and Krzysztof Pietrzak and Vanishree Rao. Adaptive security of constrained PRFs. In Palash Sarkar and Tetsu Iwata, editors, *Advances in Cryptology – ASIACRYPT 2014*, pages 82–101, 2013. Springer.

FL17. Xiong Fan and Feng-Hao Liu. Proxy re-encryption and re-signatures from lattices. Cryptology ePrint Report 2017/456, https://ia.cr/2017/456.

Gen09. Craig Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing*, STOC '09, pages 169–178, 2009. ACM.

GPV07. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. *Electronic Colloquium on Computational Complexity (ECCC)*, 14(133), 2007.

HRsV07. Susan Hohenberger, Guy N. Rothblum, abhi shelat, and Vinod Vaikuntanathan. Securely obfuscating re-encryption. In Salil P. Vadhan, editor, *Theory of Cryptography*, pages 233–252, 2007. Springer.

JKK+17. Zahra Jafargholi, Chethan Kamath, Karen Klein, Ilan Komargodski, Krzysztof Pietrzak, and Daniel Wichs. Be adaptive, avoid overcommitting. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology – CRYPTO 2017*, pages 133–163. Springer, 2017.

LV08. Benoît Libert and Damien Vergnaud. Unidirectional chosen-ciphertext secure proxy re-encryption. In Ronald Cramer, editor, *Public Key Cryptography – PKC 2008*, volume 4939 of LNCS, pages 360–379. Springer 2008.

PWA+16. Le Trieu Phong, Lihua Wang, Yoshinori Aono, Manh Ha Nguyen, and Xavier Boyen. Proxy re-encryption schemes with key privacy from LWE. Cryptology ePrint Report 2016/327, 2016. https://ia.cr/2016/327.

Reg05. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing*, STOC '05, pages 84–93, 2005. ACM.