

Registration-Based Encryption from Standard Assumptions

Sanjam Garg* Mohammad Hajiabadi† Mohammad Mahmoody‡
Ahmadreza Rahimi§ Sruthi Sekar¶

October 23, 2018

Abstract

The notion of *Registration-Based Encryption* (RBE) was recently introduced by Garg, Hajiabadi, Mahmoody, and Rahimi [TCC'18] with the goal of removing the private-key generator (PKG) from IBE. Specifically, RBE allows encrypting to identities using a (compact) master public key, like how IBE is used, with the benefit that the PKG is substituted with a weaker entity called “key curator” who has no knowledge of any secret keys. Here individuals generate their secret keys on their own and then publicly *register* their identities and their corresponding public keys to the key curator. Finally, individuals obtain “rare” decryption-key updates from the key curator as the population grows. In their work, they gave a construction of RBE schemes based on the combination of indistinguishability obfuscation and somewhere statistically binding hash functions. However, they left open the problem of constructing RBE schemes based on standard assumptions.

In this work, we resolve the above problem and construct RBE schemes based on standard assumptions (e.g., CDH or LWE). Furthermore, we show a new application of RBE in a novel context. In particular, we show that anonymous variants of RBE (which we also construct under standard assumptions) can be used for realizing abstract forms of anonymous messaging tasks in simple scenarios in which the parties communicate by writing messages on a shared board in a synchronized way.

*sanjamg@berkeley.edu, Berkeley. Research supported in part from DARPA/ARL SAFEWARE Award W911NF15C0210, AFOSR Award FA9550-15-1-0274, AFOSR YIP Award, DARPA and SPAWAR under contract N66001-15-C-4065, a Hellman Award and research grants by the Okawa Foundation, Visa Inc., and Center for Long-Term Cybersecurity (CLTC, UC Berkeley). The views expressed are those of the author and do not reflect the official policy or position of the funding agencies.

†mdhajiabadi@berkeley.edu, Berkeley and University of Virginia. Supported by NSF award CCF-1350939 and AFOSR Award FA9550-15-1-0274.

‡mohammad@virginia.edu, University of Virginia. Supported by NSF CAREER award CCF-1350939, and two University of Virginia’s SEAS Research Innovation Awards.

§ahmadreza@virginia.edu, University of Virginia. Supported by NSF award CCF-1350939.

¶sruthisekar@iisc.ac.in, Indian Institute of Science, Bangalore.

Contents

1	Introduction	3
1.1	Technical Overview	4
1.2	Potential Applications of Anonymous RBE	8
2	Preliminary Definitions	9
2.1	Previous Definitions about RBE	9
2.2	New Definitions about (Anonymous) RBE	11
2.3	Blind Public Key Encryption	12
3	Blind Hash Garbling	14
3.1	Definition of Blind Hash Garbling	14
3.2	Construction of a Blind Hash Garbling Scheme	15
4	Efficient Blind T-RBE	16
4.1	Definition	16
4.2	Construction of an Efficient Blind T-RBE scheme	17
4.3	Proofs of Completeness, Compactness and Efficiency of the T-RBE Construction	19
4.4	Proof of Security of the T-RBE construction	19
4.5	Proof of Blindness of the T-RBE Construction	21
5	Anonymous Registration-based Encryption	22
5.1	Construction of an Efficient Anonymous RBE scheme	22
5.2	Completeness, Compactness and Efficiency of the RBE Construction	25
5.3	Proof of Anonymity and Security of the RBE construction	25

1 Introduction

Identity based encryption, first introduced by Shamir [Sha84], and then realized based on pairings by Boneh and Franklin [BF01], allows a set of remote parties to communicate secretly by only knowing one single public key and the name of the recipient identity. Despite being a milestone in foundations of cryptography and a powerful tool for simplifying key-management, real-world uses of IBE schemes come with a major caveat: IBE schemes require a *private-key generator* (PKG) who holds the master key and uses it to generate decryption keys for the identities. Therefore, the PKG has the ability to decrypt all ciphertexts. This issue, inherent to IBE by design, is known as the *key escrow* problem.

Many previous works tried to rectify the key escrow problem in IBE. These efforts include making the trust de-centralized using multiple PKGs [BF01], making the PKG *accountable* for distributing the decryption keys to unauthorized users [Goy07, GLSW08], making it hard for PKG to find out the receiver identity in a large set of identities [CCV04, Cho09, WQT18], or using Certificateless Public Key Cryptography [ARP03] as a hybrid of IBE and public-key directories. However, none of these efforts resolve the key escrow problem completely.

Motivated by entirely removing PKGs from IBE schemes, recently Garg, Hajiabadi, Mahmoody, and Rahimi [GHMR18] introduced the notion of *registration-based encryption* (RBE for short). In an RBE scheme, the PKG entity is substituted by a much weaker entity called the *key curator* (KC for short). The KC will not possess any secret keys, and all it does is to manage the set of public keys of the *registered* identities. More specifically, in an RBE scheme identities (or, rather the users corresponding to the identities) generate their own public and secret keys, and then they will register their public keys to the KC who maintains and updates a public parameter \mathbf{pp}_n where n is the number of parties who have joined the system so far. This public parameter \mathbf{pp}_n can be used (now and in the future) to encrypt messages to any of the n identities who have registered so far. The first key efficiency requirement of RBE schemes is that \mathbf{pp}_n is *compact*; i.e., $\text{poly}(\kappa, \log n)$ in size where κ is the security parameter. Moreover, RBE requires that the process of “identity registration” is also efficient; i.e., runs in time $\text{poly}(\kappa, \log n)$. In order to connect an updated public parameter \mathbf{pp}_n to the previously registered identities, RBE allows the identities to obtain *updates* from the KC, which (together with their own secret keys) can be used as decryption keys. The second efficiency requirement of RBE is that such updates are only needed at most $O(\log n)$ times over the lifetime of the system. In summary, RBE schemes are required to perform both “identity registration” and “update generation” in time *sublinear* in n . In particular, these two operations are required to run in time $\text{poly}(\kappa, \log n)$ where κ is the security parameter.

The work of [GHMR18] showed how to construct RBE schemes based on the combination of indistinguishability obfuscation (IO) [BGI⁺01, GGH⁺13] and somewhere statistically binding hash functions (SSBH) [HW15]. Towards the goal of basing RBE schemes on more standard assumptions, [GHMR18] also showed how to construct *weakly efficient* RBE schemes with $\text{poly}(\kappa, n)$ identity registration time based on standard assumptions such as CDH and LWE. The work of [GHMR18] left open the question of constructing RBE schemes (with the required registration time $\text{poly}(\kappa, \log n)$) from standard cryptographic assumptions. This gap leads us to the following question, which is the main question studied in this work:

Can we base registration-based encryption on standard cryptographic assumptions?

Our results. In this work, we resolve the above question affirmatively. Namely, as the main result of this work, we construct RBE schemes with all the required compactness and efficiency requirements based on standard assumptions such as LWE, CDH, or Factoring. In particular, in our RBE scheme (based on CDH or LWE assumptions) the time it takes to register any new identity into the system is only $\text{poly}(\kappa, \log n)$ where κ is the security parameter and n is the number of identities registered into the system so far.

In addition to resolving the question above, in this work we show the usefulness of RBE by demonstrating a connection between an *anonymous* variant of RBE (defined similarly to how anonymous IBE [BDCOP04] is defined) to an abstract anonymous messaging primitive that we call *anonymous board communication* (ABC for short). At a high level, (anonymous) IBE fails to achieve ABC, exactly because of the key escrow problem, which does not exist in RBE.

1.1 Technical Overview

In this subsection, we will first describe the high level ideas behind our RBE scheme based on standard assumptions. We will then describe how to add the extra property of anonymity to RBE, allowing it to be used for realizing ABC as described above.

Figure 1 shows the high level structure and the roadmap of the primitives that we use (and construct along the way) for achieving RBE from standard assumptions. The features in parentheses (i.e., “blind” and “anonymous”) can be added to or removed from the figure. When they are added, Figure 1 demonstrates the way we obtain *anonymous* RBE.




Figure 1: Roadmap

The big picture. We construct our RBE scheme based on the primitive *hash garbling* which was formally defined in [GHMR18] but was used implicitly in some prior works [CDG⁺17, DG17, DG17, DGHM18, BLSV18], and a new primitive “time-stamp” RBE (T-RBE for short) that we introduce in this work. T-RBE is a special case of RBE, where we use the time-stamp t_{id} of the registration time of each identity instead of their (arbitrary string identity) id . T-RBE also requires the same efficiency and compactness requirements of RBE. Since T-RBE is a special case of RBE, achieving T-RBE from standard assumptions is potentially easier, and indeed we leverage this fact in our approach.

In particular, as depicted in Figure 1, we first show how T-RBE can be constructed from public-key encryption and hash garbling schemes. Then, having T-RBE, hash garbling, and a red-black Merkle tree we show how to construct an RBE scheme. This resolves the open question of [GHMR18] where they obtained *weakly* efficient RBE schemes. Finally, we show that by substituting

each of the used primitives with a “blind” version of them as defined in [BLSV18] we can bootstrap our new RBE construction to make it anonymous. Below, we describe each of these steps, their corresponding challenges, and how we resolve them.

The weakly efficient RBE of [GHMR18]. Before describing our RBE construction, we describe the main challenge in achieving the required registration efficiency, which made the construction of [GHMR18] (based on standard assumptions) *weakly efficient*. For that, we need to quickly recall the high level structure of the (weakly-efficient) construction of [GHMR18]. At a high level, the registration algorithm in the construction of [GHMR18] leads to an auxiliary information stored at the KC that consists of some Merkle trees $\text{Tree}_1, \dots, \text{Tree}_n$ where each of these trees in their leaves contains the ids of the already registered identities, along with their public keys pk . The encryption algorithm of this scheme requires to use the public key corresponding to the specific identity to encrypt the message. To do so, it requires to do a binary search on the tree containing the id (so as to access the corresponding public key), for which it is required that the leaves of the trees are *sorted* according to their labels (i.e., identities) from left to right. This binary search was captured by generation of a sequence of garbled circuits. Now, going back to the registration algorithm, if the execution of the registration algorithm ever leads to two trees with the same size, those two will be merged into one tree, which *again* needs to have the identities sorted. (Merging the trees is necessary to keep the public parameter compact, because the public parameter is basically the concatenation of the roots of these trees, and so the number of trees shall remain bounded.) Hence, every time a tree merge operation is done in their registration process, the entire tree needs to be restructured based on the *newly sorted* list of the identities in the two trees. This very sorting made the construction of [GHMR18] weakly efficient.

Step one: weakening the functionality for sake of efficiency. To overcome the challenge of achieving the required efficiency for RBE, in this work we first introduce a new primitive that we call T-RBE (short form of *time-stamp* RBE) that weakens RBE for the sake of achieving efficient registration. (Looking ahead, we will later bootstrap T-RBE to RBE.) This new primitive has the same functionality as RBE, except that in a T-RBE, the identities do not register with their actual ids, but rather with the corresponding time-stamps of their registration moments. This has the immediate obvious effect that the time-stamp strings already arrive in sorted order, hence removing the need of re-structuring the trees for sorting purposes. Note that since time-stamps are taking the role of the names of identities, all the algorithms (including the encryption algorithm) shall use the time-stamp as the identity’s label (instead of the id).

We now describe how T-RBE can indeed overcome the efficiency challenge left open about RBE. Notice that the time-stamps used for registrations are already sorted based on the arrival times. One useful consequence of this phenomenon is that, if we apply the same approach of [GHMR18] for weakly-efficient RBE, the resulting T-RBE will have leaves of the trees *automatically sorted* at all times. Hence, when we want to merge two trees, we may simply hash their roots into a new root, with a guarantee that all the leaves in one sub-tree will be larger than all those in the other sub-tree. Hence, by restricting the problem to T-RBE, we would not require to restructure the trees when we merge them. Hence, the T-RBE registration overcomes the main reason for inefficiency in the RBE construction of [GHMR18] and turns out to be efficient as required. The other algorithms of the T-RBE scheme are also similar to the corresponding algorithms of the RBE in [GHMR18], with some natural changes.

Step two: bootstrapping T-RBE to RBE (using Red-Black Merkle Tree). Recall that T-RBE uses the *time-stamps* of the identities as if they were the actual ids. So, to come up with a construction of RBE with the same efficiency as T-RBE, we somehow need to find a way to connect the actual identities of the parties to their corresponding registration time-stamps. In particular, the encryption algorithm of the RBE scheme, which now takes as input an identity id as opposed to a time-stamp, would need to first obtain the corresponding time-stamp of the given id , and then run the T-RBE encryption algorithm using this time-stamp. Further, we also need to ensure that the registration and update algorithms of this RBE scheme are efficient. At a high level, we achieve this efficiency by using a red-black Merkle tree in addition to the auxiliary input maintained by our T-RBE scheme. Such a tree allows us to (indirectly) obtain the time-stamp corresponding to an identity id in an efficient way (details of which are described below), without having to store the tree in the public parameter (which is prohibitive due to the tree size).

In order to enable this indirect access to the red-black Merkle tree, we will make further use of the hash-garbling primitive. Here we give a high-level description of our RBE scheme which uses T-RBE (and hash garbling) as subroutines together with the help of a red-black Merkle tree. In this scheme, the auxiliary information \mathbf{aux} , stored by the key curator consists of η full binary Merkle trees $\text{Tree}_1, \dots, \text{Tree}_\eta$, as was in the construction of [GHMR18]. In addition, \mathbf{aux} contains a red-black Merkle tree TimeTree , whose leaves contain pairs of identities/time-stamps, sorted according to the identities. TimeTree is a key part in enabling the efficiency of our registration part.

- **The description of TimeTree .** It is a close variant of red-black Merkle tree, where each leaf is of the form (id, t_{id}) and every non-leaf node in this tree contains the hash of its left child, the largest identity in its left sub-tree and the hash of the right child (all hashes use a hash key generated at the setup and described by a CRS). It differs from just red-black tree in the sense that each internal node also contains the largest identity in its left sub-tree. The choice of this specific data-structure is crucial for our construction. Notice that the leaves are sorted in ascending order of the identities. In addition, every node has another bit of information representing its color, which would be helpful in keeping the tree balanced using the red-black Merkle tree rotation algorithms.
- **How to register (Updating \mathbf{aux}).** As we use the T-RBE as a subroutine, \mathbf{aux} consists of the auxiliary information of T-RBE, \mathbf{aux}_T , the TimeTree and a list of already registered identities. \mathbf{aux}_T itself comprises of the Merkle trees $\mathcal{T} = \{\text{Tree}_1, \dots, \text{Tree}_\eta\}$, with the time-stamps and their corresponding public keys at the leaf nodes. So to update the \mathbf{aux} when somebody registers, we insert their identity id as well as their time-stamp t_{id} to TimeTree , update \mathbf{aux}_T using the T-RBE subroutine and add id to the list of registered identities. Recall that the T-RBE registration process involves creating a new Merkle tree with leaf nodes t_{id} and its corresponding public key \mathbf{pk} and the merging the trees in \mathcal{T} which are of same depth (merging only requires hashing the roots of the two trees to obtain a new tree).
- **How to encrypt.** The encryption algorithm takes as input the public parameter \mathbf{pp} , a message \mathbf{m} , and a receiver’s identity id , and outputs a ciphertext, which is obtained by encrypting \mathbf{m} using the T-RBE encryption under the time-stamp corresponding to id . To do this, the encryption algorithm requires to first look up id in TimeTree to obtain its time-stamp and then use it to encrypt \mathbf{m} under the T-RBE encryption. However, the encryption algorithm only takes \mathbf{pp} as public information, which is too small to contain TimeTree , and so a “direct search” is impossible. To get around this problem, the encryption algorithm “defers” this

search process to the time of decryption: Specifically, the encryption algorithm constructs a sequence of (garbled) programs, which enable one to do a binary search on `TimeTree` (during the decryption time) to obtain the time-stamp corresponding to `id` and then to use it to encrypt the message using the T-RBE encryption. The ciphertext then consists of the hash garbling of these programs.

- **How to decrypt.** The decryption algorithm takes as input two paths $u = \text{pth}_1, \text{pth}_2$, a secret key `sk` and a ciphertext, which contains the garbled programs, and outputs a message (or aborts). Here, `pth1` is the path from the root of `TimeTree` to the node that contains the `id` of the decryptor and its time-stamp t_{id} , and `pth2` is the updated path (obtained using the update algorithm of T-RBE) required for running the T-RBE decryption algorithm. Using `pth1`, the decryptor runs the hash garbled programs to obtain the ciphertext under the T-RBE scheme with time-stamp t_{id} as the output of the last program. Then, it runs the decryption algorithm of T-RBE with inputs `pth2`, `sk` and the obtained ciphertext to get the message `m`.
- **How to update u (the auxiliary information required by decryptor).** With the registration as described above, we can guarantee the efficiency. The updation algorithm requires to only read a path in the `TimeTree`, which leads to `id` and its time-stamp t_{id} and further use the updation algorithm of the T-RBE scheme. This would also be efficient. But we also need to guarantee that the number of times an `id` calls the updation algorithm is at most $\log(n)$ (where n is the the number of identities registered so far). This is not guaranteed if only have a single variant of the `TimeTree`. This is because, each time an `id` registers, its addition to the `TimeTree` modifies the root hash, changing the root-to-leaf path for every other identity registered. We resolve this issue by maintaining a variant of the same `TimeTree`, at the times corresponding to which each Merkle tree in \mathcal{T} was last updated (which means we maintain at most $\log n$ variants of `TimeTree`). This guarantees two things: firstly that an identity contained in `Treei` of \mathcal{T} will definitely be contained in the corresponding i -th variant of `TimeTree` and secondly that this identity only requires to update its root-to-leaf path in the `TimeTree`, when the tree `Treei` is modified. This would guarantee that the number of updates required by every `id` is at most $\log(n)$.

Adding anonymity. Adding the anonymity feature to our RBE scheme involves techniques which are in essence same as those used in [BLSV18]. We build an RBE scheme achieving the stronger notion of “blindness,” which in turn implies the required anonymity property of an anonymous RBE scheme. While the notion of anonymity guarantees that the identity `id` is hidden along with the message being encrypted (similar to an anonymous IBE), the property of “blindness” gives a stronger guarantee that the ciphertext generated on a uniformly random message looks uniform. The fact that this stronger guarantee is achieved by our scheme and that it implies anonymity is shown in 5.3.

As shown in the Figure 1, we can get the blind T-RBE based on blind PKE and blind hash garbling schemes. The construction of the blind T-RBE is exactly the same as the regular T-RBE, except that instead of using a regular PKE scheme and a hash garbling scheme, we will use blind variants of these primitives and separate the corresponding ciphertexts and hash inputs into two parts. Then using blind T-RBE and blind hash garbling, we can get the desired anonymous RBE scheme (which is in fact a blind RBE).

1.2 Potential Applications of Anonymous RBE

Here we describe a possible application of anonymous RBE, which demonstrates the power of (anonymous) RBE in scenarios where other seemingly similar and powerful primitives (such as anonymous IBE [BDCOP04, BW06] or anonymous PKE [BBDP01]) seem incapable of.

Anonymous board communication (ABC). In an ABC scheme, a dynamic set of (semi-honest) parties $\{\text{id}_1, \dots, \text{id}_n\}$ anonymously communicate by writing and reading on a single shared board \mathbf{B} in a synchronized way. More formally, whenever a party joins the system, they update some information on the board \mathbf{B} . Also, the communication between $\{\text{id}_1, \dots, \text{id}_n\}$ is done in synchronized cycles. In cycle t , each identity id_i has a list of messages $\mathbf{m}_{i,1}, \mathbf{m}_{i,2}, \dots$ to be delivered to some parties $\text{id}_{i,1}, \text{id}_{i,2}, \dots$. Then at the beginning of the cycle t , all of the parties write some arbitrary messages on the common board \mathbf{B} . After all the parties are done with writing, in the second half of the cycle t , all the parties read (their selected parts or all of) the content of the board \mathbf{B} . By the end of cycle t , all the parties should be able to obtain the messages intended to be sent to them. Namely, if the message \mathbf{m} was sent to identity id (by another identity id'), by the end of cycle t , \mathbf{m} should be obtained by id . The security goal is to keep all the senders and receivers anonymous from the perspective of any adversary who can read all the information written on the common board \mathbf{B} . The efficiency goal here is have parties write their messages in “small” time. In particular, we require the write time of each party to be $\text{poly}(\kappa, \log n) \cdot k$ where k is the total length of the message to be sent out and κ is the security parameter.

ABC from anonymous IBE? One might try to get ABCs from anonymous IBE as follows. The public key of the anonymous IBE scheme is written on the board, and the parties use it to compose their messages for the desired receivers. During the read part of each cycle, each party will read the whole board \mathbf{B} and try to decrypt messages that are sent to them. This approach provides receiver anonymity and sender efficiency, but the main issue is that the master secret key should be stored somewhere and there is only one place for storing it: on the board \mathbf{B} . So, anyone who has full access to the board can decrypt all the messages. On the other hand, one might try to avoid IBE and use a public-key directory storing public-keys of an (anonymous) PKE scheme at the board \mathbf{B} . Then the parties can try to use this information and write their messages to the desired recipients on the same board. The problem with this approach is that the writing parties need to read the whole set of public-keys from the board to obtain the desired target public keys.

ABC from anonymous RBE. Interestingly, anonymous RBE directly enables ABC and achieves what anonymous IBE and PKE seem incapable of. In particular, a specific part of the common board \mathbf{B} would be dedicated to store the information required for maintaining the public parameters and the auxiliary information of the KC of the anonymous RBE scheme (and note that no master secret key exists). This way, by reading the public parameter pp_n (after n people have joined), an (anonymous-receiver) message can be composed to any recipient. In particular, if $\mathbf{m}_{i,1}, \mathbf{m}_{i,2}, \dots$ are to be sent to parties $\text{id}_{i,1}, \text{id}_{i,2}, \dots$ by party id_i , all id_i does is to generate anonymous-receiver ciphertexts containing the messages and encrypted to the right identities. The anonymity of the sender stems from the fact that RBE can be used like an IBE and anyone (including those in the system) can compose messages to any other identify in a secret way. Finally, the fact that the key curator is transparent (and all it does is curating the keys) allows the parties to join the system freely one by one and update the public parameter as required.

Relation to other works on secure messaging. We emphasize that our ABC primitive is not by any means aiming to capture practical scenarios of secure messaging that have been a source of intense work over recent years [RMS18, CGF10, CGBM15, BSJ⁺17, AKTZ17, UDB⁺15,

CGCD⁺17, CB95]. For example a large body of work (e.g., see [CGF10, CGBM15, CB95]) explore realistic messaging settings in which a set of distributed parties communicate over a network and aim to message each other in a way that senders, receivers, (and some more specific relations) remain secret despite the messages being sent over the network. Some of these works achieve privacy by using non-colluding servers, while ABC only uses one “server”. On the other hand, our ABC occurs in a centralized setting in which all the messages by the parties are directly written to and read from the shared board. In a different direction, many works (e.g., see [BSJ⁺17, CGCD⁺17, PR18, JS18, BGB04] in the context of what is now known as “Ratcheted Key Exchange” study ways to expand shared keys to secure refreshed keys to be used in the future, and so they fundamentally differ from ABC simply because in ABC there are no keys shared between the parties. Thus, we note that the main goal of introducing ABC is to demonstrate a basic abstract messaging scenario with challenges that can be resolved immediately using (anonymous) RBE, while other powerful tools do not seem to be capable of doing the same.

2 Preliminary Definitions

In this section we describe the needed definitions. We separate the definitions of variants of RBE into those borrowed from previous work and those that are introduced in this work.

2.1 Previous Definitions about RBE

In this subsection we recall the definition of RBE, taken verbatim from [GHMR18].

Definition 2.1 (Syntax of RBE). A *registration-based encryption* (RBE for short) scheme consists of PPT algorithms (Gen, Reg, Enc, Upd, Dec) working as follows. The Reg and Upd algorithms are performed by the key curator, which we call KC for short.

- **Generating common random string.** Some of the subroutines below will need a common random string crs , which could be sampled publicly using some public randomness beacon. crs of length $\text{poly}(\kappa)$ is sampled at the beginning, for the security parameter κ .
- **Key generation.** $\text{Gen}(1^\kappa) \rightarrow (\text{pk}, \text{sk})$: The randomized algorithm Gen takes as input the security parameter 1^κ and outputs a pair of public/secret keys (pk, sk) . Note that these are only *public* and *secret* keys, not the *encryption* or *decryption* keys. The key generation algorithm is run by any honest party locally who wants to register itself into the system.
- **Registration.** $\text{Reg}^{\text{aux}}(\text{crs}, \text{pp}, \text{id}, \text{pk}) \rightarrow \text{pp}'$: The deterministic algorithm Reg takes as input the common random sting crs , current public parameter pp , a registering identity id and a public key pk (supposedly for the identity id), and it outputs pp' as the updated public parameters. The Reg algorithm uses *read and write* oracle access to aux which will be updated into aux' during the process of registration. (The system is initialized with public parameters pp and auxiliary information aux set to \perp .)
- **Encryption.** $\text{Enc}(\text{crs}, \text{pp}, \text{id}, \text{m}) \rightarrow \text{ct}$: The randomized algorithm Enc takes as input the common random sting crs , a public parameter pp , a recipient identity id and a plaintext message m and outputs a ciphertext ct .
- **Update.** $\text{Upd}^{\text{aux}}(\text{pp}, \text{id}) \rightarrow \text{u}$: The deterministic algorithm Upd takes as input the current information pp stored at the KC and an identity id , has *read only* oracle access to aux and generates an *update* information u that can help id to decrypt its messages.

- **Decryption.** $\text{Dec}(\text{sk}, \text{u}, \text{ct})$: The deterministic decryption algorithm Dec takes as input a secret key sk , an update information u , and a ciphertext ct , and it outputs a message $\text{m} \in \{0, 1\}^*$ or in $\{\perp, \text{GetUpd}\}$. The special symbol \perp indicates a syntax error, while GetUpd indicates that more recent update information (than u) might be needed for decryption.

Definition 2.2 (Completeness, compactness, and efficiency of RBE). For any interactive *computationally unbounded* adversary Adv that still has a limited $\text{poly}(\kappa)$ round complexity, consider the following game $\text{Comp}_{\text{Adv}}(\kappa)$ between Adv and a challenger Chal .

1. **Initialization.** Chal sets $\text{pp} = \perp$, $\text{aux} = \perp$, $\text{u} = \perp$, $\text{ID} = \emptyset$, $\text{id}^* = \perp$, $t = 0$, $\text{crs} \leftarrow U_{\text{poly}(\kappa)}$ and sends the sampled crs to Adv .
 2. Till Adv continues (which is at most $\text{poly}(\kappa)$ steps), proceed as follows. At every iteration, Adv chooses exactly one of the actions below to be performed.
 - (a) **Registering new (non-target) identity.** Adv sends some $\text{id} \notin \text{ID}$ and pk to Chal . Chal registers (id, pk) by letting $\text{pp} := \text{Reg}^{\text{aux}}(\text{crs}, \text{pp}, \text{id}, \text{pk})$ and $\text{ID} := \text{ID} \cup \{\text{id}\}$.
 - (b) **Registering the target identity.** If id^* was chosen by Adv already (i.e., $\text{id}^* \neq \perp$), skip this step. Otherwise, Adv sends some $\text{id}^* \notin \text{ID}$ to Chal . Chal then samples $(\text{pk}^*, \text{sk}^*) \leftarrow \text{Gen}(1^\kappa)$, updates $\text{pp} := \text{Reg}^{\text{aux}}(\text{crs}, \text{pp}, \text{id}^*, \text{pk}^*)$, $\text{ID} := \text{ID} \cup \{\text{id}^*\}$, and sends pk^* to Adv .
 - (c) **Encrypting for the target identity.** If $\text{id}^* = \perp$ then skip this step. Otherwise, Chal sets $t = t + 1$, then Adv sends some $\text{m}_t \in \{0, 1\}^*$ to Chal who then sets $\text{m}'_t := \text{m}_t$ and sends back a corresponding ciphertext $\text{ct}_t \leftarrow \text{Enc}(\text{crs}, \text{pp}, \text{id}^*, \text{m}_t)$ to Adv .
 - (d) **Decryption by target identity.** Adv sends a $j \in [t]$ to Chal . Chal then lets $\text{m}'_j = \text{Dec}(\text{sk}^*, \text{u}, \text{ct}_j)$. If $\text{m}'_j = \text{GetUpd}$, then Chal obtains the update $\text{u} = \text{Upd}^{\text{aux}}(\text{pp}, \text{id}^*)$ and then lets $\text{m}'_j = \text{Dec}(\text{sk}^*, \text{u}, \text{ct}_j)$.
 3. The adversary Adv wins the game if there is some $j \in [t]$ for which $\text{m}'_j \neq \text{m}_j$.
- Let $n = |\text{ID}|$ be the number of identities registered till a specific moment. We require the following properties to hold for any Adv (as specified above) and for *all* the moments (and so for all the values of ID and $n = |\text{ID}|$ as well) during the game $\text{Comp}_{\text{Adv}}(\kappa)$.

- **Completeness.** $\Pr[\text{Adv} \text{ wins in } \text{Comp}_{\text{Adv}}(\kappa)] = \text{negl}(\kappa)$.
- **Compactness of public parameters and updates.** $|\text{pp}|, |\text{u}|$ are both $\leq \text{poly}(\kappa, \log n)$.
- **Efficiency of runtime of registration and update.** The running time of each invocation of Reg and Upd algorithms is at most $\text{poly}(\kappa, \log n)$. (This implies the compactness property.)
- **Efficiency of the number of updates.** The *total* number of invocations of Upd for identity id^* in Step 2d of the game $\text{Comp}_{\text{Adv}}(\kappa)$ is at most $O(\log n)$ for every n during $\text{Comp}_{\text{Adv}}(\kappa)$.

Definition 2.3 (WE-RBE). A *weakly efficient RBE* (or WE-RBE for short) is defined similarly to Definition 2.2, where the specified $\text{poly}(\kappa, \log n)$ runtime efficiency of the registration algorithm is not required anymore, but instead we require the registration time to be $\text{poly}(\kappa, n)$.

Definition 2.4 (Security of RBE). For any interactive PPT adversary Adv , consider the following game $\text{Sec}_{\text{Adv}}(\kappa)$ between Adv and a challenger Chal . (Steps that are different from the completeness definition are denoted with purple stars (**)). Specifically, Steps 2c and 2d from Definition 2.2 are replaced by Step 3 below. Additionally, Step 3 from Definition 2.2 is replaced by Step 4 below.)

1. **Initialization.** Chal sets $\text{pp} = \perp$, $\text{aux} = \perp$, $\text{ID} = \emptyset$, $\text{id}^* = \perp$, $\text{crs} \leftarrow U_{\text{poly}(\kappa)}$ and sends the sampled crs to Adv .
2. Till Adv continues (which is at most $\text{poly}(\kappa)$ steps), proceed as follows. At every iteration, Adv chooses exactly one of the actions below to be performed.

- (a) **Registering new (non-target) identity.** Adv sends some $\text{id} \notin \text{ID}$ and pk to Chal. Chal registers (id, pk) by letting $\text{pp} := \text{Reg}^{\text{aux}}(\text{crs}, \text{pp}, \text{id}, \text{pk})$ and $\text{ID} := \text{ID} \cup \{\text{id}\}$.
 - (b) **Registering the target identity.** If id^* was chosen by Adv already (i.e., $\text{id}^* \neq \perp$), skip this step. Otherwise, Adv sends some $\text{id}^* \notin \text{ID}$ to Chal. Chal then samples $(\text{pk}^*, \text{sk}^*) \leftarrow \text{Gen}(1^\kappa)$, updates $\text{pp} := \text{Reg}^{\text{aux}}(\text{crs}, \text{pp}, \text{id}^*, \text{pk}^*)$, $\text{ID} := \text{ID} \cup \{\text{id}^*\}$, and sends pk^* to Adv.
 3. (**★★**) **Encrypting for the target identity.** If no id^* was chosen by Adv before (i.e., $\text{id}^* = \perp$) then Adv first sends some $\text{id}^* \notin \text{ID}$ to Chal. Next, Chal generates $\text{ct} \leftarrow \text{Enc}(\text{crs}, \text{pp}, \text{id}^*, b)$, where $b \leftarrow \{0, 1\}$ is a random bit, lets $\text{ID} = \text{ID} \cup \{\text{id}^*\}$, and sends ct to Adv.
 4. (**★★**) The adversary Adv outputs a bit b' and wins the game if $b = b'$.
- We call an RBE scheme secure if $\Pr[\text{Adv wins in } \text{Sec}_{\text{Adv}}(\kappa)] < \frac{1}{2} + \text{negl}(\kappa)$ for any PPT Adv.

2.2 New Definitions about (Anonymous) RBE

In this section we define an anonymity feature for the notion of RBE, and we will show later how to build efficient anonymous RBE from standard assumptions.

Definition 2.5 (Anonymous RBE). An anonymous RBE scheme has the same syntax as that of an RBE scheme, with PPT algorithms $(\text{Gen}, \text{Reg}, \text{Enc}, \text{Upd}, \text{Dec})$. It satisfies the properties of completeness, compactness and efficiency as a RBE scheme and has the following stronger notion of security: For any interactive PPT adversary Adv, consider the game $\text{EXP}_{\text{Adv}}^{\text{Anon}}(\kappa)$ between Adv and a challenger Chal as follows.

1. **Initialization.** Chal sets $\text{pp} = \perp$, $\text{aux} = \perp$, $\text{ID} = \emptyset$, $\text{id}_0 = \perp$, $\text{id}_1 = \perp$, $\text{crs} \leftarrow U_{\text{poly}(\kappa)}$ and sends the sampled crs to Adv.
2. Till Adv continues (which is at most $\text{poly}(\kappa)$ steps), proceed as follows. At every iteration, Adv can perform exactly one of the following actions.
 - (a) **Registering new (non-target) identity.** Adv sends some $\text{id} \notin \text{ID}$ and pk to Chal. Chal registers (id, pk) by getting $\text{pp} = \text{Reg}^{\text{aux}}(\text{crs}, \text{pp}, (\text{id}, \text{pk}))$ and lets $\text{ID} = \text{ID} \cup \{\text{id}\}$.
 - (b) **Registering new target identity pair.** If id_0 or id_1 was chosen by Adv already (i.e., $\text{id}_0 \neq \perp$ or $\text{id}_1 \neq \perp$), skip this step. Otherwise, Adv sends challenges $\text{id}_0, \text{id}_1 \notin \text{ID}$ to Chal. Chal first samples $(\text{pk}_0, \text{sk}_0) \leftarrow \text{Gen}(1^\kappa)$, registers id_0 by setting $\text{pp} = \text{Reg}^{\text{aux}}(\text{crs}, \text{pp}, (\text{id}_0, \text{pk}_0))$ and then samples $(\text{pk}_1, \text{sk}_1) \leftarrow \text{Gen}(1^\kappa)$, registers id_1 by setting $\text{pp} = \text{Reg}^{\text{aux}}(\text{crs}, \text{pp}, (\text{id}_1, \text{pk}_1))$. Next, Chal lets $\text{ID} = \text{ID} \cup \{\text{id}_0, \text{id}_1\}$ and sends pk_0, pk_1 to Adv.
3. **Encrypting for the challenge identity.** If id_0, id_1 was not chosen by Adv already (i.e., $\text{id}_0 = \perp, \text{id}_1 = \perp$), then Adv first sends some $\text{id}_0, \text{id}_1 \notin \text{ID}$ to Chal before continuing this step. Next, Chal samples a bit $b \in \{0, 1\}$ and generates the challenge ciphertext $\text{ct} \leftarrow \text{Enc}(\text{crs}, \text{pp}, \text{id}_b, b)$. Further, Chal sets $\text{ID} = \text{ID} \cup \{\text{id}_0, \text{id}_1\}$ and sends ct to Adv.
4. The adversary Adv outputs a bit b' and wins the game if $b' = b$.

We call an RBE scheme an Anonymous-RBE if $\Pr[\text{Adv wins in } \text{EXP}_{\text{Adv}}^{\text{Anon}}(\kappa)] < \frac{1}{2} + \text{negl}(\kappa)$ for any PPT Adv.

As a step-stone toward building efficient (anonymous) RBE, we will first show how to build a primitive which we call *timestamp-RBE*. We define this notion formally below.

Definition 2.6 (T-RBE). A *timestamp-RBE* (or T-RBE for short) has syntax exactly similar to Definition 2.1, except for one difference: we now consider the registration algorithm Reg , the encryption algorithm Enc and the update algorithm Upd to take as input the timestamp t_{id} of an identity id (binary representation of the time at which an identity registers) as input instead of the identity id . The completeness, compactness of public parameters, the efficiency of runtime of registration and update and the efficiency of the number of updates is exactly similar to Definition 2.2 and the security guarantee is similar to Definition 2.4, replacing identity id with its timestamp t_{id} in all the appropriate places. For a T-RBE, we define the notion of anonymity exactly as in Definition 2.5, replacing identity id with its timestamp t_{id} in all the appropriate places.

2.3 Blind Public Key Encryption

In this subsection we define blindness features for several cryptographic primitives, which will be used in our main constructions. We first start with the notion of blind PKE. The notion of blindness for PKE is well-studied with a few prior definitions; see, e.g., [BBDP01, Moh10]. Here we give a tailored version of this definition suitable for our later constructions.

Definition 2.7 (Blind Public Key Encryption). A blind public key encryption scheme (with public parameters) has algorithms (Params, G, E, D) which is IND-CPA secure and satisfies the following additional security property: the function $E(\text{pp}, \text{pk}, m; r)$ can be expressed as $E_1(\text{pp}; r) \parallel E_2(\text{pp}, \text{pk}, m; r)$ such that the distribution of $\{(\text{pp}, \text{pk}, \text{sk}, \text{Enc}(\text{pp}, \text{pk}, m; r)) : \text{pp} \leftarrow \text{Params}(1^\lambda), (\text{pk}, \text{sk}) \leftarrow G(\text{pp}), m \xleftarrow{\$} \mathcal{M}, r \xleftarrow{\$} \{0, 1\}^*\}$ is computationally indistinguishable from $\{(\text{pp}, \text{pk}, \text{sk}, E_1(\text{pp}; r), \text{subct}_2) : \text{pp} \leftarrow \text{Params}(1^\lambda), (\text{pk}, \text{sk}) \leftarrow G(\text{pp}), m \xleftarrow{\$} \mathcal{M}, r \xleftarrow{\$} \{0, 1\}^*, \text{subct}_2 \xleftarrow{\$} \{0, 1\}^L\}$, where $L = |E_2(\text{pp}, \text{pk}, m; r)|$.

We now define a blindness notion for garbled circuits. Our blindness requirement is the same as that introduced and used by [BLSV18].

Definition 2.8 (Blind Garbled Circuits [BLSV18]). A garbling scheme consists of PPT algorithms $(\text{Garble}, \text{Eval})$ and a simulator $G.\text{Sim}$ where:

1. $\text{Garble}(1^\lambda, 1^l, 1^m, C; \text{state}) := \text{Garble}_1(1^\lambda, 1^l, 1^m; \text{state}) \parallel \text{Garble}_2(1^\lambda, 1^l, 1^m, C; \text{state})$. Garble_1 takes as input the security parameter λ , the input length l and output length m for circuit C and a random value $\text{state} \in \{0, 1\}^\lambda$ and outputs the labels for input wire of the Garbled circuit $\{\text{lab}_{j,b}\}_{j \in [l], b \in \{0,1\}}$, where $\text{lab}_{j,b} \in \{0, 1\}^\lambda$ and Garble_2 takes the circuit C in addition, and outputs the garbled circuit \tilde{C} .
2. $\text{Eval}(1^\lambda, \tilde{C}, \tilde{\text{lab}})$ is a deterministic algorithm that takes as input the garbled circuit \tilde{C} , along with a set of l labels $\tilde{\text{lab}} = \{\text{lab}_j\}_{j \in [l]}$ and outputs a string $y \in \{0, 1\}^m$.
3. $G.\text{Sim}(1^\lambda, 1^{|\text{C}|}, 1^l, y)$ takes as input the security parameter λ , the description length of circuit C , the input length l and a string $y \in \{0, 1\}^m$ and outputs a simulated garbled circuit \tilde{C} and labels $\tilde{\text{lab}}$.

A blind garbling scheme must satisfy the following properties:

1. **Correctness.** For all circuits C , inputs x and all $(\tilde{C}, \{\text{lab}_{j,b}\}_{j \in [l], b \in \{0,1\}}) \leftarrow \text{Garble}(C)$ and $\tilde{\text{lab}} = \{\text{lab}_{j,x_j}\}_{j \in [l]}$, we have $\text{Eval}(\tilde{C}, \tilde{\text{lab}}) = C(x)$.

2. **Simulation Security.** For all circuits $C : \{0,1\}^l \rightarrow \{0,1\}^m$ and all inputs $x \in \{0,1\}^l$, the following distributions are computationally indistinguishable:

$$\begin{aligned} \{(\tilde{C}, \tilde{\text{lab}}) : (\tilde{C}, \{\text{lab}_{j,b}\}_{j \in [l], b \in \{0,1\}}) \leftarrow \text{Garble}(C), \tilde{\text{lab}} = \{\text{lab}_{j,x_j}\}_{j \in [l]}\} \\ \approx \{(\tilde{C}, \tilde{\text{lab}}) : (\tilde{C}, \tilde{\text{lab}}) \leftarrow \text{G.Sim}(1^\lambda, 1^{|\text{C}|}, 1^l, C(x))\} \end{aligned}$$

3. **Blindness.** $\text{G.Sim}(1^\lambda, 1^{|\text{C}|}, 1^l, U_m) \stackrel{\approx}{\approx} U$. The output of the simulator on a completely uniform output is indistinguishable from a uniform bit string.

We now review the notion of blind batch encryption from [BLSV18]. The notion of batch encryption is in turn similar to some notions such as hash encryption and laconic oblivious transfer [DG17, CDG⁺17].

Definition 2.9 (Blind Single Batch Encryption [BLSV18]). A blind single batch encryption scheme consists of PPT algorithms ($\text{Setup}, \text{H}, \text{HEnc}, \text{HDec}$):

1. $\text{Setup}(1^\lambda, 1^l)$ takes as input the security parameter λ , a length parameter l and outputs a hash key hk .
2. $\text{H}(\text{hk}, x)$ takes as input a hash key hk and $x \in \{0,1\}^l$ and deterministically outputs $h \in \{0,1\}^\lambda$.
3. $\text{HEnc}(\text{hk}, h, i, M)$ takes as input the hash key hk , hash value h and a message matrix $M \in \{0,1\}^{1 \times 2}$ and outputs a ciphertext ct , which can be written as a concatenation of two parts $\text{ct} = (\text{subct}_1, \text{subct}_2)$.
4. $\text{HDec}(\text{hk}, x, i, \text{ct})$ takes as input the ciphertext ct and outputs an $m \in \{0,1\}$.

A blind single batch encryption must satisfy the following properties:

1. **Correctness.** Let $\text{hk} \leftarrow \text{Setup}(1^\lambda, 1^l)$. For all x, i, M , taking $h = \text{H}(\text{hk}, x)$, $\text{ct} = \text{HEnc}(\text{hk}, h, i, M)$, it holds that $\text{HDec}(\text{hk}, x, i, \text{ct}) = M_{x_i}$, with probability at least $1/2 + 1/\text{poly}(\lambda)$ over the randomness of HEnc .
2. **Semantic Security.** For any PPT adversary Adv the probability of winning in the following game between Adv and a challenger Chal is $1/2 + \text{negl}(\lambda)$:
 - (a) Adv takes as input 1^λ and sends $1^l, x \in \{0,1\}^l, i \in [l]$ to Chal .
 - (b) Chal generates $\text{hk} = \text{Setup}(1^\lambda, 1^l)$ and sends hk to Adv .
 - (c) Adv sends a pair $M^{(0)}, M^{(1)}$, such that $M_{x_i}^{(0)} = M_{x_i}^{(1)}$, to Chal .
 - (d) Chal computes $h = \text{H}(\text{hk}, x)$, chooses $b \in_R \{0,1\}$ and sends $\text{ct} = \text{HEnc}(\text{hk}, h, i, M^{(b)})$ to Adv .
 - (e) Adv outputs a bit b' and wins if $b' = b$.
3. **Blindness.** The encryption $\text{HEnc}(\text{hk}, h, i, M; r)$ can be considered as a concatenation of $\text{HEnc}_1(\text{hk}; r) \parallel \text{HEnc}_2(\text{hk}, h, i, M; r)$. Further, any PPT adversary Adv the probability of winning in the following game with a Challenger Chal is at most $1/2 + \text{negl}(\lambda)$:

- (a) Adv takes as input 1^λ and sends $1^l, x, i$ to Chal.
- (b) Chal generates $\text{hk} = \text{Setup}(1^\lambda, 1^l)$ and computes $h = \text{H}(\text{hk}, x)$. Further it samples a random $b \in \{0, 1\}$, a random message matrix $M \in \{0, 1\}^{1 \times 2}$ and encrypts $(\text{subct}_1, \text{subct}_2) \leftarrow \text{HEnc}(\text{hk}, h, i, M)$. It generates ct as:
- If $b = 0$, the $\text{ct} = (\text{subct}_1, \text{subct}_2)$.
 - If $b = 1$, then pick a random subct'_2 of same length as subct_2 and set $\text{ct} = (\text{subct}_1, \text{subct}'_2)$.
- Chal sends hk, ct to Adv.
- (c) Adv outputs a bit b' and wins if $b' = b$.

3 Blind Hash Garbling

In this section we introduce and build a primitive which we call blind hash garbling, which will later be used as an ingredient in the construction of anonymous T-RBE schemes. We first define this notion below and will then show how to build it using tools defined in the previous sections.

3.1 Definition of Blind Hash Garbling

The notion of hash garbling was defined in [GHMR18]; here we review this notion and define a blindness feature for it.

Definition 3.1 (Blind Hash Garbling). A blind hash garbling scheme has the following polynomial time algorithms HGen, Hash, HObf, HInp:

- $\text{HGen}(1^\kappa, 1^l) \rightarrow \text{hk}$. It takes as input the security parameter κ and an output length parameter 1^l for $l \leq \text{poly}(\kappa)$, and outputs a hash key hk .
- $\text{Hash}(\text{hk}, x) = y$. It takes as input hk and $x \in \{0, 1\}^l$ and deterministically outputs $y \in \{0, 1\}^\kappa$.
- $\text{HObf}(\text{hk}, C, \text{state}) \rightarrow \tilde{C}$. It takes as input hk , a circuit C , and a secret state $\text{state} \in \{0, 1\}^\kappa$ and outputs a circuit \tilde{C} .
- $\text{HInp}(\text{hk}, y, \text{state}) \rightarrow \tilde{y}$. This takes as input hk , a value $y \in \{0, 1\}^\kappa$, and secret state state and outputs \tilde{y} . Consider \tilde{y} as concatenation of two parts $\tilde{y}_1 || \tilde{y}_2$.

A blind hash garbling scheme must satisfy the following properties:

- **Correctness.** For all κ, l , hash key $\text{hk} \leftarrow \text{HGen}(1^\kappa, 1^l)$, circuit C , input $x \in \{0, 1\}^l$, $\text{state} \in \{0, 1\}^\kappa, \tilde{C} \leftarrow \text{HObf}(\text{hk}, C, \text{state})$ and $\tilde{y} \leftarrow \text{HInp}(\text{hk}, \text{Hash}(\text{hk}, x), \text{state})$, then $\tilde{C}(\tilde{y}, x) = C(x)$.
- **Security.** There exists a PPT simulator Sim such that for all κ, l and PPT (in κ) Adv we have that

$$(\text{hk}, x, \tilde{y}_1, \tilde{C}, \tilde{y}_2) \stackrel{\text{c}}{\approx} (\text{hk}, x, \tilde{y}_1, \text{Sim}(\text{hk}, x, 1^{|\tilde{C}|}, C(x))),$$

where $\text{hk} \leftarrow \text{HGen}(1^\kappa, 1^l)$, $(C, x) \leftarrow \text{Adv}(\text{hk}), \text{state} \leftarrow \{0, 1\}^\kappa, \tilde{C} \leftarrow \text{HObf}(\text{hk}, C, \text{state})$ and $(\tilde{y}_1, \tilde{y}_2) \leftarrow \text{HInp}(\text{hk}, \text{Hash}(\text{hk}, x), \text{state})$.

- **Blindness.** The function $\tilde{y} = \text{HInp}(\text{hk}, y, \text{state}; r)$ can be expressed as the concatenation $\text{HInp}_1(\text{hk}; r) \parallel \text{HInp}_2(\text{hk}, y, \text{state}; r) = \tilde{y}_1 \parallel \tilde{y}_2$ such that

$$(\text{hk}, x, \text{HInp}_1(\text{hk}; r), \text{Sim}(\text{hk}, x, 1^{|\text{C}|}, U_{|\text{C}(x)|})) \stackrel{\epsilon}{\approx} (\text{hk}, x, \text{HInp}_1(\text{hk}; r), U_{|\tilde{\text{C}}|+|\tilde{y}_2|})$$

where $1^l, x \leftarrow \text{Adv}(1^\kappa)$, $\text{hk} \leftarrow \text{HGen}(1^\kappa, 1^l)$. (It is clear that the distinguisher should not know about the random output value that was used for simulation)

3.2 Construction of a Blind Hash Garbling Scheme

We require the following building blocks to construct blind hash garbling:

- Blind single batch encryption scheme ($\text{Setup}, \text{H}, \text{HEnc}, \text{HDec}$) as Definition 2.9.
- Blind garbled circuit scheme ($\text{Garble}, \text{Eval}$) as in Definition 2.8.

The blind hash garbling scheme is as follows:

1. $\text{HGen}(1^\kappa, 1^l)$: Generate $\text{hk} \leftarrow \text{Setup}(1^\kappa, 1^l)$ and output hk .
2. $\text{Hash}(\text{hk}, x)$: Generate $\text{H}(\text{hk}, x) = y$ and output y .
3. $\text{HObf}(\text{hk}, \text{C}, \text{state})$: Generate $\tilde{\text{C}} \leftarrow \text{Garble}_2(1^\kappa, 1^l, 1^m, \text{C}; \text{state})$ and output $\tilde{\text{C}}[\text{hk}]$ (circuit $\tilde{\text{C}}$ hardwired with hk)
4. $\text{HInp}(\text{hk}, y, \text{state})$:
 - Generate $\{\text{lab}_{j,b}\}_{j \in [l], b \in \{0,1\}} \leftarrow \text{Garble}_1(1^\kappa, 1^l, 1^m; \text{state})$
 - Generate $\tilde{y} = \{\text{HEnc}(\text{hk}, y, j, [\text{lab}_{j,0} \text{lab}_{j,1}])\}_{j \in [l]} = \{\text{subct}_{1,j}, \text{subct}_{2,j}\}_{j \in [l]}$.

Let $\tilde{y}_1 = \{\text{subct}_{1,j}\}_{j \in [l]}$ and $\tilde{y}_2 = \{\text{subct}_{2,j}\}_{j \in [l]}$. Output $\tilde{y} = (\tilde{y}_1, \tilde{y}_2)$

Theorem 3.2. *The above construction ($\text{HGen}, \text{Hash}, \text{HObf}, \text{HInp}$) satisfies the correctness, security and blindness properties as given in Definition 3.1.*

Proof. We now prove the above theorem.

1. **Correctness.** Consider the circuit $\tilde{\text{C}}[\text{hk}](\tilde{y}, x)$:
 - Recovers $\text{lab}_{j,x_j} := \text{HDec}(\text{hk}, x, j, \tilde{y}_j)$ for each $j \in [l]$, where $\tilde{y} = \{\tilde{y}_j\}_{j \in [l]}$.
 - Outputs $\text{Eval}(\tilde{\text{C}}, \{\text{lab}_{j,x_j}\}_{j \in [l]})$.

Then clearly by correctness of the blind garbling circuit and the correctness of the blind single batch encryption scheme, $\tilde{\text{C}}[\text{hk}](\tilde{y}, x) = \text{C}(x)$

2. **Security.** We define the simulator Sim as below: $\text{Sim}(\text{hk}, x, 1^{|\text{C}|}, \text{C}(x))$:
 - Evaluate $(\tilde{\text{C}}, \{\text{lab}_j\}_{j \in [l]}) \leftarrow \text{G.Sim}(1^\kappa, 1^{|\text{C}|}, 1^l, \text{C}(x))$.
 - For $j \in [l]$, let $M_j = [M_{j,0} \ M_{j,1}]$, where $M_{j,x_j} = \text{lab}_j$, $M_{j,1-x_j} \in_R \{0, 1\}^\kappa$
 - Evaluate $\tilde{y} = \{\text{HEnc}(\text{hk}, \text{H}(\text{hk}, x), j, M_j)\}_{j \in [l]}$. As expressed in the protocol, $\tilde{y} = (\tilde{y}_1, \tilde{y}_2)$.

- Output (\tilde{C}, \tilde{y}_2)

Then, by simulation security of the blind garbled circuit scheme and the semantic security of the blind single batch encryption scheme, it can be shown through a sequence of hybrids that:

$$(\text{hk}, x, \tilde{y}_1, \tilde{C}, \tilde{y}_2) \stackrel{\text{c}}{\approx} (\text{hk}, x, \tilde{y}_1, \text{Sim}(\text{hk}, x, 1^{|\text{C}|}, \text{C}(x))),$$

where $\text{hk} \leftarrow \text{HGen}(1^\kappa, 1^l)$, $(\text{C}, x) \leftarrow \text{Adv}(\text{hk}), \text{state} \leftarrow \{0, 1\}^\kappa$, $\tilde{C} \leftarrow \text{HObf}(\text{hk}, \text{C}, \text{state})$ and $(\tilde{y}_1, \tilde{y}_2) \leftarrow \text{HInp}(\text{hk}, \text{Hash}(\text{hk}, x), \text{state})$.

3. **Blindness.** For the simulator Sim described above, consider the distribution of $\text{Sim}(\text{hk}, x, 1^{|\text{C}|}, U_{|\text{C}(x)|})$ for a uniformly generated output.

- By the blindness of blind garbled circuit, $\text{G.Sim}(1^\kappa, 1^{|\text{C}|}, 1^l, U_{|\text{C}(x)|}) \stackrel{\text{c}}{\approx} U$.
- Hence for each $j \in [l]$, $M_j \stackrel{\text{c}}{\approx} U$. Thus, by blindness of the blind single batch encryption, $\tilde{y} = (\tilde{y}_1, \tilde{y}_2)$ (as described in the protocol), where \tilde{y}_1 can be expressed as $\text{HInp}_1(\text{hk}; r)$ and $(\text{hk}, x, \tilde{y}_1, \tilde{y}_2) \stackrel{\text{c}}{\approx} (\text{hk}, x, \tilde{y}_1, U)$.

Hence, it follows that $(\text{hk}, x, \tilde{y}_1, \text{Sim}(\text{hk}, x, 1^{|\text{C}|}, U_{|\text{C}(x)|})) \stackrel{\text{c}}{\approx} (\text{hk}, x, \tilde{y}_1, U)$

4 Efficient Blind T-RBE

We first define and construct an efficient blind T-RBE and then use it to construct an efficient Anonymous RBE scheme.

4.1 Definition

Definition 4.1 (Blind T-RBE). A T-RBE scheme $(\text{TGen}, \text{TReg}, \text{TEnc}, \text{TUpd}, \text{TDec})$ is said to be blind if, in addition to completeness, compactness, efficiency and security properties, as in Definition 2.6, it also satisfies the following blindness property: the function $\text{TEnc}(\text{crs}, \text{pp}, t_{\text{id}}, \text{m}; r)$ can be expressed as the concatenation $\text{TEnc}_1(\text{crs}, \text{pp}; r) \parallel \text{TEnc}_2(\text{crs}, \text{pp}, t_{\text{id}}, \text{m}; r)$ such that for any PPT adversary Adv , the probability of winning in the following game with a challenger Chal is at most $1/2 + \text{negl}(\kappa)$:

1. **Initialization.** Chal sets $\text{pp} = \perp$, $\text{aux} = \perp$, $\text{ID} = \emptyset$, $t = 1$, $t^* = \perp$, $\text{id}^* = \perp$, $\text{crs} \leftarrow U_{\text{poly}(\kappa)}$ and sends the sampled crs to Adv .
2. Till Adv continues (which is at most $\text{poly}(\kappa)$ steps), proceed as follows. At every iteration, Adv can perform exactly one of the following actions.
 - (a) **Registering new (non-target) identity.** Adv sends some $\text{id} \notin \text{ID}$ and pk to Chal . Chal registers (id, pk) by getting $\text{pp} = \text{TReg}^{\text{aux}}(\text{crs}, \text{pp}, (t, \text{pk}))$ and sets $\text{ID} = \text{ID} \cup \{\text{id}\}$ and $t = t + 1$.
 - (b) **Registering new target identity pair.** If id^* was chosen by Adv already (i.e., $\text{id}^* \neq \perp$), skip this step. Otherwise, Adv sends challenge identities $\text{id}^* \notin \text{ID}$ to Chal . Chal first samples $(\text{pk}^*, \text{sk}^*) \leftarrow \text{TGen}(1^\kappa)$, registers id^* by setting $\text{pp} = \text{TReg}^{\text{aux}}(\text{crs}, \text{pp}, (t, \text{pk}^*))$. Next, Chal lets $\text{ID} = \text{ID} \cup \{\text{id}^*\}$, $t = t^*$ and $t = t + 1$, and sends t^* , pk^* , sk^* to Adv . (Note that unlike the security property, the secret key *is* given to Adv .)

3. **Encrypting for the challenge identity.** If id^* was not chosen by Adv already (i.e., $\text{id}^* \neq \perp$), then Adv first sends some $\text{id}^* \notin \text{ID}$ to Chal before continuing this step. Next, Chal samples a random message $m \in_R \mathcal{M}$ and generates $(\text{subct}_1, \text{subct}_2) \leftarrow \text{TEnc}(\text{crs}, \text{pp}, t^*, m; r)$. It generates a bit $b \in_R \{0, 1\}$ and:

- if $b = 0$, set $\text{ct} = (\text{subct}_1, \text{subct}_2)$.
- if $b = 1$, generate a random subct'_2 of same length as subct_2 and set $\text{ct} = (\text{subct}_1, \text{subct}'_2)$.

Chal sends ct to Adv (Note that the Adv does not know the random message m being encrypted).

4. The adversary Adv outputs a bit b' and wins the game if $b' = b$.

4.2 Construction of an Efficient Blind T-RBE scheme

We construct a blind T-RBE scheme $(\text{TGen}, \text{TReg}, \text{TUpd}, \text{TEnc}, \text{TDec})$, as in Definition 4.1, using the following building blocks:

- Blind Hash Garbling scheme $(\text{HGen}, \text{Hash}, \text{HObf}, \text{HInp})$, where the function HInp is expressible as concatenation of function outputs of HInp_1 and HInp_2 as in Definition 3.1.
- Blind Public-key Encryption scheme (G, E, D) , where the function E is expressible as a concatenation of outputs of E_1 and E_2 , as in Definition 2.7.

The subroutines of the T-RBE scheme are defined as follows:

- $\text{TGen}(1^\kappa)$:
 1. $(\text{pk}, \text{sk}) \leftarrow G(1^\kappa)$
 2. Output (pk, sk) .
- $\text{TReg}^{\text{aux}}(\text{pp}, t_{\text{id}}, \text{pk})$:
 1. aux consists of a family of Merkle trees $\mathcal{T} = \{\text{Tree}_1, \dots, \text{Tree}_\eta\}$ which are constructed through the process of registration described below. It also consists of a list of timestamps corresponding to each tree in \mathcal{T} , TID , arranged in ascending order of timestamps of the identities.
 2. Parse $\text{pp} = (\text{hk}, (\text{rt}_1, d_1), \dots, (\text{rt}_\eta, d_\eta))$, where $\text{hk} \leftarrow \text{HGen}(1^\kappa, 1^{2\kappa + \log n})$ and (rt_i, d_i) represent the root and depth of Tree_i in \mathcal{T} .
 3. Updating aux :
 - (a) Create $\text{Tree}_{\eta+1}$ with leaves t_{id} and pk and root $\text{rt}_{\eta+1} = \text{Hash}(\text{hk}, t_{\text{id}} || \text{pk} || 0^\kappa)$.
 - (b) If there are two trees Tree_L and Tree_R in \mathcal{T} of same depth d then proceed as follows:
 - Let t_L and t_R denote the largest timestamps of Tree_L and Tree_R respectively (can be obtained by reading the last leaf of each tree). WLOG, suppose $t_L < t_R$.
 - Merge the trees, with Tree_L on the left and Tree_R on the right ¹, with corresponding roots rt_L and rt_R , to obtain Tree with root $\text{rt} = \text{Hash}(\text{hk}, \text{rt}_L || \text{rt}_R || t_L)$.

¹This will guarantee that the leaf nodes are sorted in ascending order of the timestamps of the identities.

- Remove Tree_L and Tree_R from \mathcal{T} and add Tree to it.
- (c) $\text{TID} := \text{TID} \cup \{t_{\text{id}}\}$.
- 4. Set $\text{pp}' = (\text{hk}, (\text{rt}_1, d_1), \dots, (\text{rt}_\zeta, d_\zeta))$, where (rt_i, d_i) represent the root and depth of Tree_i in updated \mathcal{T} .
- 5. Output pp' .
- $\text{TUpd}^{\text{[aux]}}(\text{pp}, t_{\text{id}})$:
 1. Parse $\text{pp} = (\text{hk}, (\text{rt}_1, d_1), \dots, (\text{rt}_\eta, d_\eta))$
 2. Let $u = \text{pth}$, the Merkle opening from the leaf node t_{id} and its sibling pk to the root rt_i (in Tree_i containing the timestamp t_{id}).
 3. Output u .
- $\text{TEnc}(\text{pp}, t_{\text{id}}, m)$:
 1. Parse $\text{pp} = (\text{hk}, (\text{rt}_1, d_1), \dots, (\text{rt}_\eta, d_\eta))$.
 2. For each $i = 1, \dots, \eta$:
 - (a) For each $j = 1, \dots, d_i$:
 - Sample $\text{state}_{i,j} \leftarrow \{0, 1\}^\kappa$
 - Generate $\tilde{\text{P}}_{i,j} \leftarrow \text{HObf}(\text{hk}, \text{P}_{i,j}, \text{state}_{i,j})$
 - (b) Obtain $\tilde{y}_{i,1} \leftarrow \text{HInp}(\text{hk}, \text{rt}_i, \text{state}_{i,1})$, where $\tilde{y}_{i,1} = \tilde{y}_{i,1}^{(0)} \parallel \tilde{y}_{i,1}^{(1)}$
 - (c) For each $j = 2, \dots, d_i$, obtain $\tilde{y}_{i,j}^{(0)} = \text{HInp}_1(\text{hk}; r_{i,j})$
 3. Output $\text{ct} = (\text{pp}, \{\tilde{\text{P}}_{i,j}\}_{i,j}, \{\tilde{y}_{i,j}^{(0)}\}_{i,j}, \{\tilde{y}_{i,1}^{(1)}\}_i, \text{E}_1(r))$. Let $\text{subct}_1 = (\text{pp}, \{\tilde{y}_{i,j}^{(0)}\}_{i,j}, \text{E}_1(r))$ and $\text{subct}_2 = (\{\tilde{\text{P}}_{i,j}\}_{i,j}, \{\tilde{y}_{i,1}^{(1)}\}_i)$. Then $\text{ct} = (\text{subct}_1, \text{subct}_2)$.

The program $\text{P}_{i,j}$ is as defined below:

Hardwired: $t_{\text{id}}, \text{hk}, \text{state}_{i,j+1}, m, r_{i,j+1}$ (where $\text{state}_{i,d_i+1} = \perp, r_{i,d_i+1} = r$)

Input: $a \parallel b \parallel t^*$

1. If $t^* = 0^\kappa$ and $a = t_{\text{id}}$, output $\text{E}_2(b, m; r_{i,j+1})$
 2. If $t^* = 0^\kappa$ and $a \neq t_{\text{id}}$, output \perp .
 3. If $t_{\text{id}} > t^*$, output $\text{HInp}_2(\text{hk}, b, \text{state}_{i,j+1}; r_{i,j+1})$
Else, output $\text{HInp}_2(\text{hk}, a, \text{state}_{i,j+1}; r_{i,j+1})$
- $\text{TDec}(\text{sk}, u, \text{ct})$:
 1. Parse ct as $(\text{pp}, \{\tilde{\text{P}}_{i,j}\}_{i,j}, \{\tilde{y}_{i,j}^{(0)}\}_{i,j}, \{\tilde{y}_{i,1}^{(1)}\}_i, \text{E}_1(r))$.
 2. Parse u as $\text{pth} = (z_0 = \text{rt}_{i^*}, z_1, \dots, z_{d_{i^*}} = t_{\text{id}} \parallel \|\text{pk}\|_{0^\kappa})$, the Merkle opening from leaf t_{id} to the root rt_{i^*} of Tree_{i^*} (containing t_{id}).
 3. For each $j = 1, \dots, d_{i^*} - 1$ evaluate:
 - $\tilde{y}_{i^*,j+1}^{(1)} \leftarrow \tilde{\text{P}}_{i^*,j}(\tilde{y}_{i^*,j}^*, z_j)$
 4. Let $c_2 = \tilde{\text{P}}_{i^*,d_{i^*}}(\tilde{y}_{i^*,d_{i^*}}^*, z_{d_{i^*}}^*)$. If $c_2 = \perp$, set $c = \perp$, else set $c = \text{E}_1(r) \parallel c_2$.

5. If $c = \perp$, output `GetUpd`, else output $D(\text{sk}, c)$.

Theorem 4.2. *The T-RBE construction 4.2 satisfies the completeness, compactness, efficiency, security and blindness (Definition 4.1) properties.*

In the following subsections, we prove Theorem 4.2.

4.3 Proofs of Completeness, Compactness and Efficiency of the T-RBE Construction

Completeness. By the correctness of the hash garbling scheme and the Public-key Encryption scheme, completeness property follows.

Compactness of public parameters and update. Consider the public parameter $\text{pp} = (\text{hk}, (\text{rt}_1, d_1), \dots, (\text{rt}_\eta, d_\eta))$. We observe that:

- The number of Merkle trees, η , in \mathcal{T} at any time is at most $\log(n)$. This is because the trees are full binary trees and the size of the trees are always different (as we keep merging in the registration process).
- The hash key hk , the root and the depth of each tree are all of size κ each.

Hence, the size of pp is $O(\kappa \cdot \log n)$.

Consider the update $\text{u} = \text{pth}$, the Merkle opening from leaf node t_{id} and its sibling pk to the root rt_i . The depth of the tree is d_i and hence, there are at most $2 \cdot d_i + 1$ nodes in the Merkle opening, where $d_i \leq \kappa$. Hence, the size of u is at most $O((2 \cdot \kappa + \log n) \cdot \kappa) = O(\text{poly}(\kappa, \log(n)))$.

Efficiency of runtime of registration and update. The registration process involves evaluating a hash value to create a new tree with the new identity and then merging the trees of same depth after this. We observe that:

- The number of merge operations is $O(\log n)$ as the number of trees is always logarithmic.
- Computing each hash costs $O(\kappa)$.

Hence, each invocation of the registration process takes time $O(\kappa \cdot \log n)$. Consider a single invocation of the update. This just involves reading aux to output the Merkle opening required and this takes time $O((2 \cdot \kappa + \log n) \cdot \kappa) = O(\text{poly}(\kappa, \log n))$.

Efficiency of the number of updates. Each identity would require to invoke `Upd`, whenever the Merkle opening for the id gets modified. This in turn happens whenever two trees are merged. Since the number of merges is at most $O(\log n)$, the total number of invocations of `Upd` by each identity is at most $O(\log n)$.

4.4 Proof of Security of the T-RBE construction

We prove the security assuming that there is only one tree at the time of encryption. The proof for the case of multiple trees will be the same.

Proof. Suppose that at the time of encryption, the underlying tree has root rt and depth d . For simplicity, for each $j \in [d]$, we denote the circuits $P_{1,j}$ by P_j and the state used for obfuscation, $\text{state}_{1,j}$ by state_j , i.e., for each $j \in [d]$

$$P_j \equiv P_{1,j}[t_{\text{id}}, \text{hk}, \text{state}_{j+1}, \text{m}, r_{1,j+1}]$$

where all the variables are as in the encryption algorithm of the construction.

As in the construction, let the Merkle opening from the leaf node t_{id} and its sibling \mathbf{pk} to the root \mathbf{rt} be denoted by:

$$\mathbf{pth} = ((t_{id}, \mathbf{pk}, 0^\kappa), (a_1, b_1, t_1), \dots, (a_{d-1}, b_{d-1}, t_{d-1}), \mathbf{rt})$$

As in the decryption algorithm of the construction, we denote the hash-obfuscation of the inputs of the circuits by $\tilde{y}_j \equiv \tilde{y}_{1,j} = (\tilde{y}_{1,j}^{(0)}, \tilde{y}_{1,j}^{(1)})$ for each $j \in [d]$. Then, in the actual game, the output of the encryption algorithm is $\mathbf{ct}_0 := (\mathbf{subct}_{0,1}, \mathbf{subct}_{0,2})$, where $\mathbf{subct}_{0,1} = (\mathbf{pp}, \tilde{y}_1^{(0)}, \dots, \tilde{y}_d^{(0)}, E_1(r))$ and $\mathbf{subct}_{0,2} = (\tilde{P}_1, \dots, \tilde{P}_d, \tilde{y}_1^{(1)})$ (as in the protocol we have two parts of the ciphertext).

We describe the following sequence of hybrids, where we first replace the garbled versions of the programs P_j and the corresponding garbled inputs \tilde{y}_j by their simulated variants, which do not use \mathbf{state}_j , one by one.

- **Hybrid₀ (encryption in real game)**: The ciphertext here will be $\mathbf{ct}_0 := (\mathbf{subct}_{0,1}, \mathbf{subct}_{0,2})$, where $\mathbf{subct}_{0,1} = (\mathbf{pp}, \tilde{y}_1^{(0)}, \dots, \tilde{y}_d^{(0)}, E_1(r))$ and $\mathbf{subct}_{0,2} = (\tilde{P}_1, \dots, \tilde{P}_d, \tilde{y}_1^{(1)})$, are as described above.
- **Hybrid₁**: We replace the first obfuscated program \tilde{P}_1 with its simulated form. $\tilde{P}_2, \dots, \tilde{P}_d$ are sampled as in the construction. Let $\tilde{P}_{1,\text{sim}}$ and $\tilde{y}_{1,\text{sim}}^{(1)}$ be sampled as:

$$(\tilde{P}_{1,\text{sim}}, \tilde{y}_{1,\text{sim}}^{(1)}) \leftarrow \text{Sim}(\mathbf{hk}, (a_{d-1}, b_{d-1}, t_{d-1}), 1^{|\mathbf{P}_1|}, \tilde{y}_2^{(1)})$$

Then, the ciphertext in this hybrid is $\mathbf{ct}_1 := (\mathbf{subct}_{1,1}, \mathbf{subct}_{1,2})$, where $\mathbf{subct}_{1,1} = (\mathbf{pp}, \tilde{y}_1^{(0)}, \dots, \tilde{y}_d^{(0)}, E_1(r))$ and $\mathbf{subct}_{1,2} = (\tilde{P}_{1,\text{sim}}, \tilde{P}_2, \dots, \tilde{P}_d, \tilde{y}_{1,\text{sim}}^{(1)})$.

- **Hybrid_i**, for each $i \in [d-1]$: The ciphertext is $\mathbf{ct}_i := (\mathbf{subct}_{i,1}, \mathbf{subct}_{i,2})$, where $\mathbf{subct}_{i,1} = (\mathbf{pp}, \tilde{y}_1^{(0)}, \dots, \tilde{y}_d^{(0)}, E_1(r))$ and $\mathbf{subct}_{i,2} = (\tilde{P}_{1,\text{sim}}, \dots, \tilde{P}_{i,\text{sim}}, \tilde{P}_{i+1}, \dots, \tilde{P}_d, \tilde{y}_{1,\text{sim}}^{(1)})$ where for each $j \in [i]$:

$$(\tilde{P}_{j,\text{sim}}, \tilde{y}_{j,\text{sim}}^{(1)}) \leftarrow \text{Sim}(\mathbf{hk}, (a_{d-j}, b_{d-j}, t_{d-j}), 1^{|\mathbf{P}_j|}, \tilde{y}_{j+1}^{(1)})$$

- **Hybrid_d**: The ciphertext is $\mathbf{ct}_d := (\mathbf{subct}_{d,1}, \mathbf{subct}_{d,2})$, where $\mathbf{subct}_{d,1} = (\mathbf{pp}, \tilde{y}_1^{(0)}, \dots, \tilde{y}_d^{(0)}, E_1(r))$ and $\mathbf{subct}_{d,2} = (\tilde{P}_{1,\text{sim}}, \dots, \tilde{P}_{d,\text{sim}}, \tilde{y}_{1,\text{sim}}^{(1)})$ where for each $j \in [d-1]$

$$(\tilde{P}_{j,\text{sim}}, \tilde{y}_{j,\text{sim}}^{(1)}) \leftarrow \text{Sim}(\mathbf{hk}, (a_{d-j}, b_{d-j}, t_{d-j}), 1^{|\mathbf{P}_j|}, \tilde{y}_{j+1}^{(1)})$$

and

$$(\tilde{P}_{d,\text{sim}}, \tilde{y}_{d,\text{sim}}^{(1)}) \leftarrow \text{Sim}(\mathbf{hk}, (t_{id}, \mathbf{pk}, 0^\kappa), 1^{|\mathbf{P}_d|}, E_2(\mathbf{pk}, \mathbf{m}; r))$$

By the simulation security of the hash garbling scheme, we know that, for each $j \in [d-1]$,

$$(\mathbf{hk}, (a_{d-j}, b_{d-j}, t_{d-j}), \tilde{y}_j^{(0)}, \tilde{P}_j, \tilde{y}_j^{(1)}) \stackrel{\epsilon}{\approx} (\mathbf{hk}, (a_{d-j}, b_{d-j}, t_{d-j}), \tilde{y}_j^{(0)}, \tilde{P}_{j,\text{sim}}, \tilde{y}_{j,\text{sim}}^{(1)})$$

and

$$(\mathbf{hk}, (t_{id}, \mathbf{pk}, 0^\kappa), \tilde{y}_d^{(0)}, \tilde{P}_d, \tilde{y}_d^{(1)}) \stackrel{\epsilon}{\approx} (\mathbf{hk}, (t_{id}, \mathbf{pk}, 0^\kappa), \tilde{y}_d^{(0)}, \tilde{P}_{d,\text{sim}}, \tilde{y}_{d,\text{sim}}^{(1)})$$

Hence, it follows that for each $i = 0, \dots, d-1$, **Hybrid_i** $\stackrel{\epsilon}{\approx}$ **Hybrid_{i+1}**.

Now, let **Hybrid_i⁰** denote the hybrids described above with use of underlying message \mathbf{m}_0 and

Hybrid_{*i*}¹ for message m_1 . By semantic security of the underlying public-key encryption scheme, we get:

$$\text{Sim}(\text{hk}, (t_{\text{id}}, \text{pk}, 0^\kappa), 1^{|\text{P}_d|}, \text{E}_2(\text{pk}, m_0; r)) \stackrel{\text{c}}{\approx} \text{Sim}(\text{hk}, (t_{\text{id}}, \text{pk}, 0^\kappa), 1^{|\text{P}_d|}, \text{E}_2(\text{pk}, m_1; r))$$

Hence, **Hybrid**_{*d*}⁰ $\stackrel{\text{c}}{\approx}$ **Hybrid**_{*d*}¹. Then, it follows that **Hybrid**₀⁰ $\stackrel{\text{c}}{\approx}$ **Hybrid**₀¹, which represent the actual security game with use of messages m_0 and m_1 in respective hybrids. Hence, the security of the T-RBE scheme is proved. \square

4.5 Proof of Blindness of the T-RBE Construction

We prove the blindness of the scheme assuming that there is only one tree at the time of encryption for simplicity. The proof for the case of multiple trees will be the same.

Proof. Suppose that at the time of encryption, the underlying tree has root rt and depth d . For simplicity, for each $j \in [d]$, we denote the circuits $\text{P}_{1,j}$ by P_j and the state used for obfuscation, $\text{state}_{1,j}$ by state_j , i.e., for each $j \in [d]$

$$\text{P}_j \equiv \text{P}_{1,j}[t_{\text{id}}, \text{hk}, \text{state}_{j+1}, \text{m}, r_{1,j+1}]$$

where all the variables are as in the encryption algorithm of the construction.

As in the construction, let the Merkle opening from the leaf node t_{id} and its sibling pk to the root rt be denoted by:

$$\text{pth} = ((t_{\text{id}}, \text{pk}, 0^\kappa), (a_1, b_1, t_1), \dots, (a_{d-1}, b_{d-1}, t_{d-1}), \text{rt}) .$$

As in the decryption algorithm of the construction, we denote the hash-obfuscation of the inputs of the circuits by $\tilde{y}_j \equiv \tilde{y}_{1,j} = (\tilde{y}_{1,j}^{(0)}, \tilde{y}_{1,j}^{(1)})$ for each $j \in [d]$. Then, in the actual game, where now the message m is chosen at random, the output of the encryption algorithm is $\text{ct}_0 := (\text{subct}_{0,1}, \text{subct}_{0,2})$, where $\text{subct}_{0,1} = (\text{pp}, \tilde{y}_1^{(0)}, \dots, \tilde{y}_d^{(0)}, \text{E}_1(r))$ and $\text{subct}_{0,2} = (\tilde{\text{P}}_1, \dots, \tilde{\text{P}}_d, \tilde{y}_1^{(1)})$ (as in the protocol we have two parts of the ciphertext). Clearly $\text{subct}_{0,1}$ is expressible as $\text{TEnc}_1(\text{pp}; r)$.

We describe the following sequence of hybrids, where we first follow the sequence of hybrids in the security proof of Section 4.4 to replace all the garbled versions of the programs P_j and the corresponding garbled inputs \tilde{y}_j by their simulated variants. Then we use the blindness property of the blind hash garbling scheme and the blind public key encryption scheme to replace the simulated garbled circuits with uniform.

- **Hybrid**₀ (**encryption in real blindness game**): The ciphertext in this hybrid will be $\text{ct}_0 := (\text{subct}_{0,1}, \text{subct}_{0,2})$, where $\text{subct}_{0,1} = (\text{pp}, \tilde{y}_1^{(0)}, \dots, \tilde{y}_d^{(0)}, \text{E}_1(r))$ and $\text{subct}_{0,2} = (\tilde{\text{P}}_1, \dots, \tilde{\text{P}}_d, \tilde{y}_1^{(1)})$, as described above.
- **Hybrid**₁: The ciphertext is $\text{ct}_1 := (\text{subct}_{1,1}, \text{subct}_{1,2})$, where $\text{subct}_{1,1} = (\text{pp}, \tilde{y}_1^{(0)}, \dots, \tilde{y}_d^{(0)}, \text{E}_1(r))$ and $\text{subct}_{1,2} = (\tilde{\text{P}}_{1,\text{sim}}, \dots, \tilde{\text{P}}_{d,\text{sim}}, \tilde{y}_{1,\text{sim}}^{(1)})$ where for each $j \in [d-1]$

$$(\tilde{\text{P}}_{j,\text{sim}}, \tilde{y}_{j,\text{sim}}^{(1)}) \leftarrow \text{Sim}(\text{hk}, (a_{d-j}, b_{d-j}, t_{d-j}), 1^{|\text{P}_j|}, \tilde{y}_{j+1}^{(1)})$$

and

$$(\tilde{\text{P}}_{d,\text{sim}}, \tilde{y}_{d,\text{sim}}^{(1)}) \leftarrow \text{Sim}(\text{hk}, (t_{\text{id}}, \text{pk}, 0^\kappa), 1^{|\text{P}_d|}, \text{E}_2(\text{pk}, \text{m}; r)) .$$

- **Hybrid_i** for each $i = 2, \dots, d$: The ciphertext is $\text{ct}_i := (\text{subct}_{i,1}, \text{subct}_{i,2})$, where $\text{subct}_{i,1} = (\text{pp}, \tilde{y}_1^{(0)}, \dots, \tilde{y}_d^{(0)}, E_1(r))$ and $\text{subct}_{i,2} = (\tilde{P}_{1,\text{sim}}, \dots, \tilde{P}_{d-(i-1),\text{sim}}, U_{|\tilde{P}_{d-(i-2)}|}, \dots, U_{|\tilde{P}_d|}, \tilde{y}_{1,\text{sim}}^{(1)})$.
- **Hybrid_{d+1}**: The ciphertext is $\text{ct}_{d+1} := (\text{subct}_{d+1,1}, \text{subct}_{d+1,2})$, where $\text{subct}_{d+1,1} = (\text{pp}, \tilde{y}_1^{(0)}, \dots, \tilde{y}_d^{(0)}, E_1(r))$ and $\text{subct}_{d+1,2} = (U_{|\tilde{P}_1|+|\tilde{y}_1^{(1)}|}, U_{|\tilde{P}_2|}, \dots, U_{|\tilde{P}_d|})$.

By the proof of security in Section 4.4, we know that **Hybrid₀** $\stackrel{\text{c}}{\approx}$ **Hybrid₁**. Now, by blindness property of the blind public key encryption scheme, we know that:

$$(\text{pk}, \text{sk}, E_1(r), E_2(\text{pk}, \text{m}; r)) \stackrel{\text{c}}{\approx} (\text{pk}, \text{sk}, E_1(r), U) .$$

Hence, by the blindness property of the hash garbling scheme, it follows that:

$$(\text{hk}, (t_{\text{id}}, \text{pk}, 0^\kappa), \text{sk}, \tilde{y}_d^{(0)}, \tilde{P}_{d,\text{sim}}, \tilde{y}_{d,\text{sim}}^{(1)}) \stackrel{\text{c}}{\approx} (\text{hk}, (t_{\text{id}}, \text{pk}, 0^\kappa), \text{sk}, \tilde{y}_d^{(0)}, U_{|\text{P}_d|+|\tilde{y}_d^{(1)}|}) .$$

Hence, it follows that **Hybrid₁** $\stackrel{\text{c}}{\approx}$ **Hybrid₂** even given the secret key sk .

By consecutive use of blindness property of the hash garbling scheme, it would follow that for each $i = 2, \dots, d$, **Hybrid_i** $\stackrel{\text{c}}{\approx}$ **Hybrid_{i+1}** which holds even given the secret key sk .

Hence, we have that **Hybrid₀** $\stackrel{\text{c}}{\approx}$ **Hybrid_{d+1}**, even given the secret key sk , which exactly represents the blindness game. This completes the proof of blindness of the T-RBE scheme. \square

5 Anonymous Registration-based Encryption

We now construct an efficient anonymous RBE scheme as in Definition 2.5.

5.1 Construction of an Efficient Anonymous RBE scheme

We now construct a RBE scheme (Gen, Reg, Enc, Upd, Dec) using the following building blocks:

- Blind Hash Garbling scheme (HGen, Hash, HObf, HInp), where the function HInp is expressible as concatenation of function outputs of HInp₁ and HInp₂, as in Definition 3.1.
- Blind T-RBE scheme (TGen, TReg, TEnc, TUpd, TDec), where the function TEnc is expressible as the concatenation of function outputs of TEnc₁ and TEnc₂, as in Definition 4.1.

The subroutines of our RBE scheme are defined as follows:

- Gen(1^κ) :
 1. $(\text{pk}, \text{sk}) \leftarrow \text{TGen}(1^\kappa)$
 2. Output (pk, sk) .
- Reg^[aux](pp, id, pk) :
 1. aux consists of a Red-black Merkle tree **TimeTree** which has the identities (along with their timestamps) at the leaf nodes, sorted according to the identities. **TimeTree** is constructed through the process of registration, as described below. aux also consists of the database aux_T required by the T-RBE scheme and a list **ID** of identities registered so far, in ascending order. Let the Merkle trees contained in aux_T be $\mathcal{T} = \{\text{Tree}_1, \dots, \text{Tree}_\eta\}$.

2. Parse \mathbf{pp} as $(\mathbf{hk}, \mathbf{pp}_T, (\mathbf{rt}_1, d_1), \dots, (\mathbf{rt}_\eta, d_\eta))$, where $\mathbf{hk} \leftarrow \text{HGen}(1^\kappa, 1^{3\kappa})$, \mathbf{pp}_T corresponds to the public parameter of the T-RBE scheme and $(\mathbf{rt}_1, d_1), \dots, (\mathbf{rt}_\eta, d_\eta)$ are the roots and depth of the same TimeTree corresponding to the time when $\text{Tree}_1, \dots, \text{Tree}_\eta$ where last updated respectively.²
 3. Updating TimeTree:
 - (a) TimeTree is a Red-Black Merkle tree with each non-leaf node containing a hash of its left child, hash of its right child and the largest identity on the left subtree. The leaves of the tree contain the identities (with their timestamps, i.e., a binary representation of when the identity registered). Further, each node has an additional bit of information, indicating if it's colored red or black. The color helps in keeping the tree "approximately" balanced after each insertion³.
 - (b) Evaluate $h_{\text{id}} = \text{Hash}(\mathbf{hk}, \text{id} || 0^{2\kappa - \log n} || t_{\text{id}})$.
 - (c) To insert the new id at the right location, we first parse the root of TimeTree as $\mathbf{rt} = h_1 || \text{id}^* || h_2$. If $\text{id} > \text{id}^*$, read the right child, else read the left child. Continue traversing the path down the tree to figure out the correct insertion point of id .
 - (d) The parent node of $\text{id} || t_{\text{id}}$ contains h_{id} along with the largest identity on its left subtree and the hash of its other child. Re-order the tree, recolor, to keep it balanced. This will take at most $\log n$ time (as it's a red-black tree). Every node that is not a leaf or a parent node of a leaf node, is of the form $h_L || \text{id}^* || h_R$, where $h_L = \text{Hash}(\mathbf{hk}, h_1 || \text{id}_L || h_2)$ and $h_R = \text{Hash}(\mathbf{hk}, h_3 || \text{id}_R || h_4)$ and id^* is the largest identity on the left subtree of this node.
 4. Evaluate $\mathbf{pp}'_T \leftarrow \text{TReg}^{\text{[aux]}}(\mathbf{pp}_T, t_{\text{id}}, \mathbf{pk})$, where t_{id} will be the binary representation of the timestamp corresponding to id (can be obtained by checking ID to see the current number of identities).
 5. Suppose the registration process of T-RBE above results in a merge of trees Tree_L and Tree_R and the corresponding root hashes of TimeTree at the time of last update of these trees be \mathbf{rt}_L and \mathbf{rt}_R . Remove (\mathbf{rt}_L, d_L) and (\mathbf{rt}_R, d_R) from \mathbf{pp} and add (\mathbf{rt}, d) , which is the root hash and the depth of TimeTree at time t_{id} . Suppose the updated root hashes and depth of TimeTree after all the merges in \mathcal{T} be $(\mathbf{rt}_1, d_1), \dots, (\mathbf{rt}_\zeta, d_\zeta)$.
 6. Set $\mathbf{pp}' = (\mathbf{hk}, \mathbf{pp}'_T, (\mathbf{rt}_1, d_1), \dots, (\mathbf{rt}_\zeta, d_\zeta))$.
 7. Output \mathbf{pp}' .
- $\text{Upd}^{\text{[aux]}}(\mathbf{pp}, \text{id})$:
 1. Parse $\mathbf{pp} = (\mathbf{hk}, \mathbf{pp}_T, (\mathbf{rt}_1, d_1), \dots, (\mathbf{rt}_\eta, d_\eta))$.
 2. Evaluate $\mathbf{u}_1 \leftarrow \text{TUpd}^{\text{[aux]}}(\mathbf{pp}_T, t_{\text{id}})$.
 3. Set pth to be the path from the leaf node $\text{id} || t_{\text{id}} || 0^{2\kappa - \log n}$ to the root hash \mathbf{rt}_i of TimeTree at the time of last modification of the Merkle tree, Tree_i , containing id .⁴

²Looking ahead, we need to store the root hashes of TimeTree at multiple times in order to ensure that the number of updates required by each person remains $\log n$.

³The main advantage of having a Red-Black Merkle tree is that after each insertion, the depth of the tree does not increase beyond $\log n$, where n is the number of people registered in the system. The balancing is not perfect, but ensures that further insertions, rearrangement after insertion to balance, searches, all take time $O(\log n)$.

⁴Note that we must store the versions of the same TimeTree at times corresponding to last update of each Tree_i in \mathcal{T} . But there would only be $\log n$ such versions.

4. Set $\mathbf{u} = (\mathbf{u}_1, \text{pth})$.
 5. Output \mathbf{u} .
- $\text{Enc}(\text{pp}, \text{id}, \mathbf{m})$:
 1. Parse pp as $(\text{hk}, \text{pp}_T, (\text{rt}_1, d_1), \dots, (\text{rt}_\eta, d_\eta))$.
 2. For each $i = 1, \dots, \eta$:
 - (a) For each $j = 1, \dots, d_\eta$:
 - i. Sample $\text{state}_{i,j} \leftarrow \{0, 1\}^\kappa$.
 - ii. Generate $\tilde{\mathbf{Q}}_{i,j} \leftarrow \text{HObf}(\text{hk}, \mathbf{Q}_{i,j}, \text{state}_{i,j})$
 - (b) Parse rt_i as $h_1 \parallel \text{id}^* \parallel h_2$.
 - (c) If $\text{id} > \text{id}^*$, obtain $\tilde{y}_{i,1} \leftarrow \text{HInp}(\text{hk}, h_2, \text{state}_1)$, where $\tilde{y}_{i,1} = \tilde{y}_{i,1}^{(0)} \parallel \tilde{y}_{i,1}^{(1)}$
Else obtain $\tilde{y}_{i,1} \leftarrow \text{HInp}(\text{hk}, h_1, \text{state}_1)$, where $\tilde{y}_{i,1} = \tilde{y}_{i,1}^{(0)} \parallel \tilde{y}_{i,1}^{(1)}$.
 - (d) For each $j = 2, \dots, d_i$, obtain $\tilde{y}_{i,j}^{(0)} \leftarrow \text{HInp}_1(\text{hk}; r_{i,j})$.
 3. Output $\text{ct} = (\text{pp}, \{\tilde{\mathbf{Q}}_{i,j}\}_{i,j}, \{\tilde{y}_{i,j}^{(0)}\}_{i,j}, \{\tilde{y}_{i,1}^{(1)}\}_i, \text{TEnc}_1(\text{pp}_T; \mathbf{r}))$.
Let $\text{subct}_1 = (\text{pp}, \{\tilde{y}_{i,j}^{(0)}\}_{i,j}, \text{TEnc}_1(\text{pp}_T; \mathbf{r}))$ and $\text{subct}_2 = (\{\tilde{\mathbf{Q}}_{i,j}\}_{i,j}, \{\tilde{y}_{i,1}^{(1)}\}_i)$. Then $\text{ct} = (\text{subct}_1, \text{subct}_2)$.

The program $\mathbf{Q}_{i,j}$ is defined as:

Hardwired: $\text{hk}, \text{state}_{i,j+1}, \text{id}, \mathbf{m}, r_{i,j+1}, \mathbf{r}, \text{pp}_T$ (where $\text{state}_{i,d_i+1} = \perp, r_{i,d_i+1} = \perp$)

Input: $a \parallel \text{id}^* \parallel b$

1. If $\text{id}^* = 0^{2\kappa - \log n}$ and $a = \text{id}$, output $\text{TEnc}_2(\text{pp}_T, b, \mathbf{m}; \mathbf{r})$
 2. If $\text{id}^* = 0^{2\kappa - \log n}$ and $a \neq \text{id}$, output \perp
 3. If $\text{id} > \text{id}^*$, output $\text{HInp}_2(\text{hk}, b, \text{state}_{i,j+1}; r_{i,j+1})$
Else, output $\text{HInp}_2(\text{hk}, a, \text{state}_{i,j+1}; r_{i,j+1})$
- $\text{Dec}(\text{sk}, \mathbf{u}, \text{ct})$:
 1. Parse \mathbf{u} as $(\mathbf{u}_1, \text{pth})$, where $\text{pth} = (z_0 = \text{rt}_i, z_1, \dots, z_d = \text{id} \parallel 0^{2\kappa - \log n} \parallel t_{\text{id}})$. Here rt_i is the root hash of TimeTree at the time when Tree_i , the tree containing id , was last updated.
 2. Parse ct as $(\text{pp}, \{\tilde{\mathbf{Q}}_{i,j}\}_{i,j}, \{\tilde{y}_{i,j}^{(0)}\}_{i,j}, \{\tilde{y}_{i,1}^{(1)}\}_i, \text{TEnc}_1(\text{pp}_T; \mathbf{r}))$.
 3. For $j = 1, \dots, d-1$, evaluate:
 - $\tilde{y}_{i,j+1}^{(1)} \leftarrow \tilde{\mathbf{Q}}_{i,j}(\tilde{y}_{i,j}, z_j)$
 4. Let $c_2 = \tilde{\mathbf{Q}}_{i,d}(\tilde{y}_{i,d}, z_d)$. If $c_2 = \perp$, set $c = \perp$, else set $c = \text{TEnc}_1(\text{pp}_T; \mathbf{r}) \parallel c_2$.⁵
 5. If $c = \perp$, output GetUpd , else output $\text{TDec}(\text{sk}, \mathbf{u}_1, c)$.

Theorem 5.1. *The RBE construction 5.1 satisfies the completeness, compactness, efficiency (Definition 2.2) and the stronger security notion of an anonymous RBE (Definition 2.5) properties.*

In the following subsections, we prove Theorem 5.1.

⁵Alternately, we could have performed these operations for each i , which would be the number of trees in \mathcal{T} . Here, we would have obtained a value $\neq \perp$ only for one i

5.2 Completeness, Compactness and Efficiency of the RBE Construction

Completeness. By the correctness of the hash garbling scheme and the completeness of the T-RBE scheme, completeness property follows.

Compactness of public parameters and update. Consider the public parameter $\text{pp} = (\text{hk}, \text{pp}_T, (\text{rt}_1, d_1), \dots, (\text{rt}_\eta, d_\eta))$. We observe that:

- The public parameter of T-RBE, pp_T is of size $O(\text{poly}(\kappa, \log n))$.
- The hash key hk and the depth d_i are each of size κ .
- The root of the time tree rt_i is of size 3κ .
- η is at any time is at most $\log(n)$, as the number of trees in \mathcal{T} at any time is at most $\log(n)$.

Hence, the size of pp is $O(\text{poly}(\kappa, \log n))$.

Consider the update $\mathbf{u} = (\mathbf{u}_1, \text{pth})$. By efficiency of T-RBE, size of \mathbf{u}_1 is $O(\text{poly}(\kappa, \log n))$. pth is the path of nodes from the leaf $\text{id} || t_{\text{id}} || 0^{2\kappa - \log n}$ to the root rt and hence of size at most $\kappa \cdot (3\kappa) = O(\kappa)$. Hence, the size of \mathbf{u} is $O(\text{poly}(\kappa, \log n))$.

Efficiency of runtime of registration and update. The registration process involves running the registration of underlying T-RBE, which takes time at most $\text{poly}(\kappa, \log n)$ and inserting the new identity into the Red-black Merkle tree `TimeTree`, which takes time at most $\log n$. Hence, each invocation of the registration process takes time $O(\text{poly}(\kappa, \log n))$.

Consider a single invocation of the update process. It involves a single invocation of the update algorithm of T-RBE, which takes time at most $\text{poly}(\kappa, \log n)$ and reading a single path from `TimeTree` in `aux`, which takes time $O((3\kappa) \cdot \kappa)$ and an additional $\log(n)$ time to figure out the correct version of `TimeTree` to read the path from. Hence a single invocation of the update algorithm takes time $O(\text{poly}(\kappa, \log n))$.

Efficiency of the number of updates. Each identity would require to invoke `Upd`, whenever the Merkle tree in \mathcal{T} containing id gets modified (by a merge). Unless a merge operation occurs, the identities can use the opening in the `TimeTree` corresponding to the time its Merkle tree (in \mathcal{T}) was last updated (even though `TimeTree` gets updated at each registration). Since the number of merges is at most $O(\log(n))$, the number of invocations of `Upd` by each identity is at most $O(\log n)$.

5.3 Proof of Anonymity and Security of the RBE construction

We now prove that the RBE scheme satisfies the stronger notion of security of an anonymous RBE (Definition 2.5).

Proof. We prove the security assuming that there is only one Merkle tree in \mathcal{T} at the time of encryption for simplicity. The proof for the case of multiple trees will be the same. Let `TimeTree` have root rt and let the challenge identity id be at depth d in `TimeTree`, at the time of encryption. For each $j \in [d]$, as in the construction, we have:

$$Q_j \equiv Q_{1,j}[\text{hk}, \text{state}_{1,j+1}, \text{id}, \text{m}, r_{1,j+1}, r, \text{pp}_T]$$

where all the variables are as in the encryption algorithm of the construction.

As in the construction, let the path from the leaf node $\text{id} || 0^{2\kappa - \log n} || t_{\text{id}}$ to the root rt be denoted by:

$$\text{pth} = ((\text{id}, 0^{2\kappa - \log n}, t_{\text{id}}), (a_1, \text{id}_1, b_1), \dots, (a_{d-1}, \text{id}_{d-1}, b_{d-1}), \text{rt})$$

As in the decryption algorithm of the construction, we denote the hash-garbling of the inputs of the circuits by $\tilde{y}_j = \tilde{y}_{1,j}^{(0)} || \tilde{y}_{1,j}^{(1)}$ for each $j \in [d]$. Then, in the actual game, the output of the encryption algorithm is $\text{ct}_0 := (\text{subct}_{0,1}, \text{subct}_{0,2})$, where $\text{subct}_{0,1} = (\text{pp}, \{\tilde{y}_j^{(0)}\}_j, \text{TEnc}_1(\text{pp}_T; r))$ and $\text{subct}_{0,2} = (\{\tilde{Q}_j\}_j, \tilde{y}_1^{(1)})$.

We describe the following sequence of hybrids, *for uniformly drawn message* m , where we first replace the garbled versions of the programs Q_j and the corresponding garbled inputs \tilde{y}_j by their simulated variants, which do not use state_j , one by one. Then, in the subsequent hybrids, we replace the simulated garbled programs and inputs by uniform, one by one. This uses the blindness of the T-RBE scheme and the blind hash garbling scheme.

- **Hybrid₀ (encryption in real game):** The ciphertext in this hybrid will be $\text{ct}_0 := (\text{subct}_{0,1}, \text{subct}_{0,2})$, where $\text{subct}_{0,1} = (\text{pp}, \{\tilde{y}_j^{(0)}\}_j, \text{TEnc}_1(\text{pp}_T; r))$ and $\text{subct}_{0,2} = (\tilde{Q}_1, \dots, \tilde{Q}_d, \tilde{y}_1^{(1)})$, as described above.
- **Hybrid₁:** We replace the first obfuscated program \tilde{Q}_1 with its simulated form. $\tilde{Q}_2, \dots, \tilde{Q}_d$ are sampled as in the construction. Let $\tilde{Q}_{1,\text{sim}}$ and $\tilde{y}_{1,\text{sim}}^{(1)}$ be sampled as:

$$(\tilde{Q}_{1,\text{sim}}, \tilde{y}_{1,\text{sim}}^{(1)}) \leftarrow \text{Sim}(\text{hk}, (a_{d-1}, \text{id}_{d-1}, b_{d-1}), 1^{|\text{Q}_1|}, \tilde{y}_2^{(1)})$$

Then, the ciphertext in this hybrid is

$$\text{ct}_1 := (\text{subct}_{1,1}, \text{subct}_{1,2}), \text{ where } \text{subct}_{1,1} = (\text{pp}, \{\tilde{y}_j^{(0)}\}_j, \text{TEnc}_1(\text{pp}_T; r)) \text{ and } \text{subct}_{1,2} = (\tilde{Q}_{1,\text{sim}}, \dots, \tilde{Q}_d, \tilde{y}_{1,\text{sim}}^{(1)}).$$

- **Hybrid_i,** for each $i \in [d-1]$: The ciphertext is $\text{ct}_i := (\text{subct}_{i,1}, \text{subct}_{i,2})$, where $\text{subct}_{i,1} = (\text{pp}, \{\tilde{y}_j^{(0)}\}_j, \text{TEnc}_1(\text{pp}_T; r))$ and $\text{subct}_{i,2} = (\tilde{Q}_{1,\text{sim}}, \dots, \tilde{Q}_{i,\text{sim}}, \tilde{Q}_{i+1}, \dots, \tilde{Q}_d, \tilde{y}_{1,\text{sim}}^{(1)})$ where for each $j \in [i]$:

$$(\tilde{Q}_{j,\text{sim}}, \tilde{y}_{j,\text{sim}}^{(1)}) \leftarrow \text{Sim}(\text{hk}, (a_{d-j}, \text{id}_{d-j}, b_{d-j}), 1^{|\text{Q}_j|}, \tilde{y}_{j+1}^{(1)})$$

- **Hybrid_d:** The ciphertext is $\text{ct}_d := (\text{subct}_{d,1}, \text{subct}_{d,2})$, where $\text{subct}_{d,1} = (\text{pp}, \{\tilde{y}_j^{(0)}\}_j, \text{TEnc}_1(\text{pp}_T; r))$ and $\text{subct}_{d,2} = (\tilde{Q}_{1,\text{sim}}, \dots, \tilde{Q}_{d,\text{sim}}, \tilde{y}_{1,\text{sim}}^{(1)})$ where for each $j \in [d-1]$

$$(\tilde{Q}_{j,\text{sim}}, \tilde{y}_{j,\text{sim}}^{(1)}) \leftarrow \text{Sim}(\text{hk}, (a_{d-j}, \text{id}_{d-j}, b_{d-j}), 1^{|\text{Q}_j|}, \tilde{y}_{j+1}^{(1)})$$

and

$$(\tilde{Q}_{d,\text{sim}}, \tilde{y}_{d,\text{sim}}^{(1)}) \leftarrow \text{Sim}(\text{hk}, (\text{id}, 0^{2\kappa - \log n}, t_{\text{id}}), 1^{|\text{Q}_d|}, \text{TEnc}_2(\text{pp}_T, t_{\text{id}}, \text{m}; r))$$

- **Hybrid_i,** for each $i = d+1, \dots, 2d-1$: The ciphertext is $\text{ct}_i := (\text{subct}_{i,1}, \text{subct}_{i,2})$, where $\text{subct}_{i,1} = (\text{pp}, \{\tilde{y}_j^{(0)}\}_j, \text{TEnc}_1(\text{pp}_T; r))$ and $\text{subct}_{i,2} = (\tilde{Q}_{1,\text{sim}}, \dots, \tilde{Q}_{2d-i,\text{sim}}, U_{|\tilde{Q}_{2d-i+1}|}, \dots, U_{|\tilde{Q}_d|}, \tilde{y}_{1,\text{sim}}^{(1)})$

- **Hybrid_{2d}**: The ciphertext is

$$\text{ct}_{2d} := (\text{subct}_{2d,1}, \text{subct}_{2d,2}), \text{ where } \text{subct}_{2d,1} = (\text{pp}, \{\tilde{y}_j^{(0)}\}_j, \text{TEnc}_1(\text{pp}_T; r)) \text{ and } \text{subct}_{2d,2} = (U_{|\tilde{Q}_1|+|\tilde{y}_1^{(1)}|}, U_{|\tilde{Q}_2|}, \dots, U_{|\tilde{Q}_d|})$$

By the simulation security of the hash garbling scheme, we know that, for each $j \in [d-1]$,

$$(\text{hk}, (a_{d-j}, \text{id}_{d-j}, b_{d-j}), \tilde{y}_j^{(0)}, \tilde{Q}_j, \tilde{y}_j^{(1)}) \stackrel{\epsilon}{\approx} (\text{hk}, (a_{d-j}, \text{id}_{d-j}, b_{d-j}), \tilde{y}_j^{(0)}, \tilde{Q}_{j,\text{sim}}, \tilde{y}_{j,\text{sim}}^{(1)})$$

and

$$(\text{hk}, (\text{id}, 0^{2\kappa-\log n}, t_{\text{id}}), \tilde{y}_d^{(0)}, \tilde{Q}_d, \tilde{y}_d^{(1)}) \stackrel{\epsilon}{\approx} (\text{hk}, (\text{id}, 0^{2\kappa-\log n}, t_{\text{id}}), \tilde{y}_d^{(0)}, \tilde{Q}_{d,\text{sim}}, \tilde{y}_{d,\text{sim}}^{(1)})$$

Hence, it follows that for each $i = 0, \dots, d-1$, **Hybrid_i** $\stackrel{\epsilon}{\approx}$ **Hybrid_{i+1}**.

Now, as the sequence of hybrids were defined for uniformly drawn message \mathbf{m} , by the blindness of underlying T-RBE scheme, we know that:

$$(\text{TEnc}_1(\text{pp}_T; r), \text{TEnc}_2(\text{pp}_T, t_{\text{id}}, \mathbf{m}; r)) \stackrel{\epsilon}{\approx} (\text{TEnc}_1(\text{pp}_T; r), U)$$

Hence, by the blindness property of the hash garbling scheme, it follows that:

$$(\text{hk}, (\text{id}, 0^{2\kappa-\log n}, t_{\text{id}}), \tilde{y}_d^{(0)}, \tilde{Q}_{d,\text{sim}}, \tilde{y}_{d,\text{sim}}^{(1)}) \stackrel{\epsilon}{\approx} (\text{hk}, (\text{id}, 0^{2\kappa-\log n}, t_{\text{id}}), \tilde{y}_d^{(0)}, U_{|\tilde{Q}_d|+|\tilde{y}_1^{(d)}|})$$

Hence, **Hybrid_d** $\stackrel{\epsilon}{\approx}$ **Hybrid_{d+1}**.

Similarly, by consecutive application of the blindness property of the hash garbling scheme, it holds that for each $i = d+1, \dots, 2d-1$, **Hybrid_i** $\stackrel{\epsilon}{\approx}$ **Hybrid_{i+1}**.

Hence, we have that, for uniformly drawn message \mathbf{m} ,

$$\mathbf{Hybrid}_0 \stackrel{\epsilon}{\approx} \mathbf{Hybrid}_{2d} \tag{1}$$

We now prove that, in fact, indistinguishability in Equation 1 is stronger than the desired anonymity of the RBE scheme:

For $b = 0, 1$, let $\text{ct}_0^{m_b, \text{id}_b} = (\text{subct}_{0,1}^{m_b, \text{id}_b}, \text{subct}_{0,2}^{m_b, \text{id}_b})$ and $\text{ct}_{2d}^{m_b, \text{id}_b} = (\text{subct}_{2d,1}^{m_b, \text{id}_b}, \text{subct}_{2d,2}^{m_b, \text{id}_b})$ denote the ciphertexts in **Hybrid₀** and **Hybrid_{2d}** respectively for the b -th message, identity pair.

1. By security of T-RBE, it would follow that **Hybrid_d⁰** $\stackrel{\epsilon}{\approx}$ **Hybrid_d¹**, where **Hybrid_i^β** denote the hybrids described above with use of message \mathbf{m}_β . Hence, it follows that **Hybrid₀⁰** $\stackrel{\epsilon}{\approx}$ **Hybrid₀¹**.
2. Using the above indistinguishability, it follows that, for a random message \mathbf{m} , we have $\text{ct}_0^{m_0, \text{id}_0} \stackrel{\epsilon}{\approx} \text{ct}_0^{m, \text{id}_0}$
3. Since **Hybrid₀** $\stackrel{\epsilon}{\approx}$ **Hybrid_{2d}** for random message \mathbf{m} , it then follows that $\text{ct}_0^{m, \text{id}_0} = (\text{subct}_{0,1}^{m, \text{id}_0}, \text{subct}_{0,2}^{m, \text{id}_0}) \stackrel{\epsilon}{\approx} (\text{subct}_{0,1}^{m, \text{id}_0}, U) = \text{ct}_{2d}^{m, \text{id}_0}$
4. Since $\text{subct}_{0,1}^{m, \text{id}_0} = \text{TEnc}_1(\text{pp}_T; r)$, is independent of identity or the message \mathbf{m} , it then follows that $\text{ct}_{2d}^{m, \text{id}_0} \stackrel{\epsilon}{\approx} \text{ct}_{2d}^{m, \text{id}_1}$.

Combining the steps 2, 3 and 4 above, it follows that $\text{ct}_0^{m_0, \text{id}_0} \stackrel{\epsilon}{\approx} \text{ct}_0^{m_1, \text{id}_1}$. This exactly represents the stronger security game of the anonymous RBE. Hence, the anonymity and security of the RBE scheme is proved. \square

References

- [AKTZ17] Nikolaos Alexopoulos, Aggelos Kiayias, Riivo Talviste, and Thomas Zacharias. Mcmix: Anonymous messaging via secure multiparty computation. In *USENIX Security Symposium*. USENIX Association, Vancouver, BC, pages 1217–1234, 2017.
- [ARP03] Sattam S Al-Riyami and Kenneth G Paterson. Certificateless public key cryptography. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 452–473. Springer, 2003.
- [BBDP01] Mihir Bellare, Alexandra Boldyreva, Anand Desai, and David Pointcheval. Key-privacy in public-key encryption. In Colin Boyd, editor, *Advances in Cryptology – ASIACRYPT 2001*, volume 2248 of *Lecture Notes in Computer Science*, pages 566–582, Gold Coast, Australia, December 9–13, 2001. Springer, Heidelberg, Germany.
- [BDCOP04] Dan Boneh, Giovanni Di Crescenzo, Rafail Ostrovsky, and Giuseppe Persiano. Public key encryption with keyword search. In *International conference on the theory and applications of cryptographic techniques*, pages 506–522. Springer, 2004.
- [BF01] Dan Boneh and Matthew K. Franklin. Identity-based encryption from the Weil pairing. In Joe Kilian, editor, *Advances in Cryptology – CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 213–229, Santa Barbara, CA, USA, August 19–23, 2001. Springer, Heidelberg, Germany.
- [BGB04] Nikita Borisov, Ian Goldberg, and Eric Brewer. Off-the-record communication, or, why not to use pgp. In *Proceedings of the 2004 ACM workshop on Privacy in the electronic society*, pages 77–84. ACM, 2004.
- [BGI⁺01] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In Joe Kilian, editor, *Advances in Cryptology – CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 1–18, Santa Barbara, CA, USA, August 19–23, 2001. Springer, Heidelberg, Germany.
- [BLSV18] Zvika Brakerski, Alex Lombardi, Gil Segev, and Vinod Vaikuntanathan. Anonymous IBE, leakage resilience and circular security from new assumptions. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018, Part I*, volume 10820 of *Lecture Notes in Computer Science*, pages 535–564, Tel Aviv, Israel, April 29 – May 3, 2018. Springer, Heidelberg, Germany.
- [BSJ⁺17] Mihir Bellare, Asha Camper Singh, Joseph Jaeger, Maya Nyayapati, and Igors Stepanovs. Ratcheted encryption and key exchange: The security of messaging. In *Annual International Cryptology Conference*, pages 619–650. Springer, 2017.
- [BW06] Xavier Boyen and Brent Waters. Anonymous hierarchical identity-based encryption (without random oracles). In *Annual International Cryptology Conference*, pages 290–307. Springer, 2006.
- [CB95] David A Cooper and Kenneth P Birman. Preserving privacy in a network of mobile computers. Technical report, Cornell University, 1995.

- [CCV04] Zhaohui Cheng, Richard Comley, and Luminita Vasii. Remove key escrow from the identity-based encryption system. In *Exploring New Frontiers of Theoretical Informatics*, pages 37–50. Springer, 2004.
- [CDG⁺17] Chongwon Cho, Nico Döttling, Sanjam Garg, Divya Gupta, Peihan Miao, and Antigoni Polychroniadou. Laconic oblivious transfer and its applications. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology – CRYPTO 2017, Part II*, volume 10402 of *Lecture Notes in Computer Science*, pages 33–65, Santa Barbara, CA, USA, August 20–24, 2017. Springer, Heidelberg, Germany.
- [CGBM15] Henry Corrigan-Gibbs, Dan Boneh, and David Mazières. Riposte: An anonymous messaging system handling millions of users. *arXiv preprint arXiv:1503.06115*, 2015.
- [CGCD⁺17] Katriel Cohn-Gordon, Cas Cremers, Benjamin Dowling, Luke Garratt, and Douglas Stebila. A formal security analysis of the signal messaging protocol. In *Security and Privacy (EuroS&P), 2017 IEEE European Symposium on*, pages 451–466. IEEE, 2017.
- [CGF10] Henry Corrigan-Gibbs and Bryan Ford. Dissent: accountable anonymous group messaging. In *Proceedings of the 17th ACM conference on Computer and communications security*, pages 340–350. ACM, 2010.
- [Cho09] Sherman SM Chow. Removing escrow from identity-based encryption. In *International Workshop on Public Key Cryptography*, pages 256–276. Springer, 2009.
- [DG17] Nico Döttling and Sanjam Garg. Identity-based encryption from the Diffie-Hellman assumption. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology – CRYPTO 2017, Part I*, volume 10401 of *Lecture Notes in Computer Science*, pages 537–569, Santa Barbara, CA, USA, August 20–24, 2017. Springer, Heidelberg, Germany.
- [DGHM18] Nico Döttling, Sanjam Garg, Mohammad Hajiabadi, and Daniel Masny. New constructions of identity-based and key-dependent message secure encryption schemes. In Michel Abdalla and Ricardo Dahab, editors, *PKC 2018: 21st International Conference on Theory and Practice of Public Key Cryptography, Part I*, volume 10769 of *Lecture Notes in Computer Science*, pages 3–31, Rio de Janeiro, Brazil, March 25–29, 2018. Springer, Heidelberg, Germany.
- [GGH⁺13] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th Annual Symposium on Foundations of Computer Science*, pages 40–49, Berkeley, CA, USA, October 26–29, 2013. IEEE Computer Society Press.
- [GHMR18] Sanjam Garg, Mohammad Hajiabadi, Mohammad Mahmoody, and Ahmadreza Rahimi. Registration-based encryption: Removing private-key generator from IBE. *Theory of Cryptography Conference (TCC)*, 2018. <https://eprint.iacr.org/2018/919>.
- [GLSW08] Vipul Goyal, Steve Lu, Amit Sahai, and Brent Waters. Black-box accountable authority identity-based encryption. In *Proceedings of the 15th ACM conference on Computer and communications security*, pages 427–436. ACM, 2008.

- [Goy07] Vipul Goyal. Reducing trust in the PKG in identity based cryptosystems. In Alfred Menezes, editor, *Advances in Cryptology – CRYPTO 2007*, volume 4622 of *Lecture Notes in Computer Science*, pages 430–447, Santa Barbara, CA, USA, August 19–23, 2007. Springer, Heidelberg, Germany.
- [HW15] Pavel Hubacek and Daniel Wichs. On the communication complexity of secure function evaluation with long output. In Tim Roughgarden, editor, *ITCS 2015: 6th Conference on Innovations in Theoretical Computer Science*, pages 163–172, Rehovot, Israel, January 11–13, 2015. Association for Computing Machinery.
- [JS18] Joseph Jaeger and Igors Stepanovs. Optimal channel security against fine-grained state compromise: The safety of messaging. In *Annual International Cryptology Conference*, pages 33–62. Springer, 2018.
- [Moh10] Payman Mohassel. A closer look at anonymity and robustness in encryption schemes. In Masayuki Abe, editor, *Advances in Cryptology – ASIACRYPT 2010*, volume 6477 of *Lecture Notes in Computer Science*, pages 501–518, Singapore, December 5–9, 2010. Springer, Heidelberg, Germany.
- [PR18] Bertram Poettering and Paul Rösler. Towards bidirectional ratcheted key exchange. In *Annual International Cryptology Conference*, pages 3–32. Springer, 2018.
- [RMS18] Paul Rösler, Christian Mainka, and Jörg Schwenk. More is less: On the end-to-end security of group chats in signal, whatsapp, and threema. 2018.
- [Sha84] Adi Shamir. Identity-based cryptosystems and signature schemes. In G. R. Blakley and David Chaum, editors, *Advances in Cryptology – CRYPTO’84*, volume 196 of *Lecture Notes in Computer Science*, pages 47–53, Santa Barbara, CA, USA, August 19–23, 1984. Springer, Heidelberg, Germany.
- [UDB⁺15] Nik Unger, Sergej Dechand, Joseph Bonneau, Sascha Fahl, Henning Perl, Ian Goldberg, and Matthew Smith. Sok: secure messaging. In *Security and Privacy (SP), 2015 IEEE Symposium on*, pages 232–249. IEEE, 2015.
- [WQT18] Quanyun Wei, Fang Qi, and Zhe Tang. Remove key escrow from the BF and Gentry identity-based encryption with non-interactive key generation. *Telecommunication Systems*, pages 1–10, 2018.