

Trace and Revoke with Optimal Parameters from Polynomial Hardness

Shweta Agrawal* Simran Kumari† Anshu Yadav‡ Shota Yamada§

Abstract

A *trace and revoke* (TR) scheme is an N user traitor tracing scheme which additionally enables the encryptor to specify a list $L \subseteq$ of revoked users so that these users can no longer decrypt ciphertexts. The “holy grail” of this line of work is a construction which resists unbounded collusions, achieves ciphertext, public and secret key sizes independent¹ of $|L|$ and $|N|$, and is based on polynomial hardness assumptions. In this work we make the following contributions:

1. *Public Trace Setting*: We provide a construction which (i) achieves optimal parameters, (ii) supports embedding identities (from an exponential space) in user secret keys, (iii) relies on polynomial hardness assumptions, namely compact functional encryption (FE) and a key-policy attribute based encryption (ABE) with special efficiency properties constructed by Boneh et al. (Eurocrypt 2014) from Learning With Errors (LWE), and (iv) enjoys adaptive security with respect to the revocation list. The previous best known construction by Nishimaki, Wichs and Zhandry (Eurocrypt 2016) which achieved optimal parameters and embedded identities, relied on indistinguishability obfuscation, which is considered an inherently subexponential assumption and achieved only selective security with respect to the revocation list.
2. *Secret Trace Setting*: We provide the first construction with optimal ciphertext, public and secret key sizes and embedded identities from any assumption outside Obfuscopia. In detail, our construction relies on Lockable Obfuscation which can be constructed using LWE (Goyal, Koppula, Waters and Wichs, Zirdelis, Focs 2017) and two ABE schemes: (i) the key-policy scheme with special efficiency properties by Boneh et al. (Eurocrypt 2014) and (ii) a ciphertext-policy ABE for P which was recently constructed by Wee (Eurocrypt 2022) using a new assumption called *evasive and tensor* LWE. This assumption, introduced to build an ABE, is believed to be much weaker than lattice based assumptions underlying FE or iO – in particular it is required even for lattice based broadcast, without trace.

Moreover, by relying on subexponential security of LWE, both our constructions can also support a *super-polynomial* sized revocation list, so long as it allows efficient representation and membership testing. Ours is the first work to achieve this, to the best of our knowledge.

*IIT Madras, shweta@cse.iitm.ac.in

†IIT Madras, sim78608@gmail.com

‡IIT Madras, anshu.yadav06@gmail.com

§National Institute of Advanced Industrial Science and Technology (AIST), Tokyo, yamada-shota@aist.go.jp

¹Ignoring logarithmic dependencies.

Contents

1	Introduction	3
1.1	Prior Work: Embedded Identity Trace and Revoke	5
1.2	Our Results	5
1.3	Technical Overview	6
2	Preliminaries	16
2.1	Functional Encryption	18
2.2	Attribute Based Encryption	21
2.3	Key-Policy ABE by Boneh et al. [BGG ⁺ 14]	23
2.4	Lockable Obfuscation	26
3	Revocable Predicate Encryption	26
4	Public-key RPE from FE and LWE	28
5	Revocable Mixed Functional Encryption	44
5.1	Definition	44
5.2	Construction	45
5.3	Security	49
6	Secret Key RPE from Evasive and Tensor LWE	57
6.1	Construction	57
6.2	Security	60
7	Embedded Identity Trace and Revoke	63
7.1	Indexed Trace and Revoke with Embedded Identity	65
7.2	Bounded Trace and Revoke with Embedded Identity	66
7.3	Unbounded Trace and Revoke with Embedded Identity	67
8	Indexed Trace and Revoke with Embedded Identity	68
8.1	Construction	68
8.2	Security	71
9	Bounded Trace and Revoke with Embedded identity	78
9.1	Construction	78
9.2	Security	81
10	Unbounded Trace and Revoke with Embedded Identities	86
10.1	Construction	86
10.2	Security	89
11	Extension to Super-Polynomial Size Revocation List	92

1 Introduction

Traitor Tracing. Traitor tracing (TT) schemes were first proposed by Chor, Fiat, and Naor [CFN94] to enable content providers to trace malicious users who exploit their secret keys to construct illegal decryption boxes. More formally, a TT system is a public key encryption system comprising N users for some large polynomial N . Each user $i \in [N]$ is provided with a unique secret key sk_i for decryption, and there is a common public key pk which is used by the content distributor to encrypt content. If any collection of users attempts to create and sell a new decoding box that can be used to decrypt the content, then the tracing algorithm, given black-box access to any such pirate decoder, is guaranteed to output an index $i \in [N]$ of one of the corrupt users, which in turn allows to hold them accountable. The literature has considered both public and secret tracing, where the former requires knowledge of a secret key to run the trace procedure and the latter does not suffer from this restriction.

Broadcast Encryption. Broadcast Encryption [FN93] (BE) introduced by Fiat and Naor, is also an N user system which supports an encrypted broadcast functionality. In BE, a content provider can transmit a single ciphertext over a broadcast channel so that only an authorized subset $S \subseteq N$ of users can decrypt and recover the message. More formally, each user $i \in [N]$ is provided with a unique decryption key sk_i and a ciphertext ct_m for a message m also encodes an authorized list S so that sk_i decrypts ct_m if and only if $i \in S$. Evidently, public key encryption provides a trivial construction of BE with ciphertext of size $O(N)$ – thus, the focus in such schemes is to obtain short ciphertext, ideally logarithmic in N .

Trace and Revoke. Naor and Pinkas [NP10] suggested a meaningful interleaving of these two functionalities so that traitors that are identified by the TT scheme can be removed from the set of authorized users in a BE scheme. To capture this, they defined the notion of “Trace and Revoke” (TR) where the content provider in a broadcast encryption scheme includes a list L of revoked users in the ciphertext, and sk_i works to decrypt ct_L if $i \notin L$. Moreover, it is required that revocation remain compatible with tracing, so that if an adversary builds a pirate decoder that can decrypt ciphertexts encrypted with respect to L , then the tracing algorithm should be able to output a corrupt non-revoked user who participated in building the illegal decoder. Trace and revoke systems provide a functionality which is richer than a union of BE and TT, since the traitor traced by the latter must belong to the set of non-revoked users for the guarantee to be meaningful. As such, TT schemes have been challenging to construct even given TT and BE schemes.

The Quest for Optimal Parameters. All the above primitives have been researched extensively over decades, resulting in a long sequence of beautiful constructions, non-exhaustively [BGW05, BW06, BWZ14, BZ17, NWZ16, KW20, GKW18, GKW19, AY20, Wee22]. A central theme in this line of work is to achieve optimal parameters, namely optimal sizes for the ciphertext, public key and secret key (and understanding tradeoffs thereof), while still supporting unbounded collusion resistance. Towards this, the powerful hammer of indistinguishability obfuscation (iO) [BGI⁺01] yielded the first feasibility results for traitor tracing [BZ17] as well as trace and revoke [NWZ16] while multilinear maps [GGH13, CLT13] led to the first construction for broadcast encryption [BWZ14]. Though there has been remarkable progress in the construction of iO from standard assumptions, with the breakthrough work of Jain, Lin and Sahai [JLS21, JLS22] finally reaching this goal, iO is an inherently subexponential assumption [GPSZ17] because the challenger is required to check whether two circuits are functionally equivalent, which can take exponential time in general. Indeed, all known constructions of iO assume subexponential hardness of the underlying algebraic assumptions. To address this limitation, a sequence of works [GPSZ17, AM18, BKS16, GPS16, KS20] has sought to replace iO by polynomially hard

assumptions such as functional encryption in different applications.

Optimal TT, BE and TR from Polynomial Assumptions: For traitor tracing, the first construction from standard assumptions was finally achieved by the seminal work of Goyal, Koppula and Waters [GKW18] in the secret trace setting, from the Learning With Errors (LWE) assumption. For broadcast encryption, this goal was achieved by Agrawal and Yamada [AY20] from LWE and the bilinear GGM. In the standard model, Agrawal, Wichs and Yamada [AWY20] provided a construction from a non-standard knowledge assumption on pairings, while Wee [Wee22] provided a construction from a new assumption on lattices, called Evasive and Tensor LWE. For trace and revoke, the only construction without iO that achieves collusion resistance and optimal parameters is by Goyal, Vusirikala and Waters (GVW) [GVW19] from *positional witness encryption* (PWE) which is a polynomial hardness assumption. However, their construction incurs an exponential loss in the security proof, requiring the underlying PWE to satisfy subexponential security. Moreover, although PWE is not an inherently subexponential assumption as are iO and witness encryption (WE), we do not currently know of any *constructions* of PWE that rely on standard polynomial hardness assumptions. In particular, [JLS21, JLS22] do *not* imply PWE from polynomial hardness.

Pathway via Secret Tracing. Both the iO and PWE based constructions of TR [NWZ16, GVW19] achieve public tracing. Taking a lesson from TT, where optimal parameters were achieved from standard assumptions only in the secret trace setting [GKW18], a natural approach towards optimal TR from better assumptions is to weaken the tracing algorithm to be secret key. This approach has been explored in a number of works – the current best parameters are achieved by Zhandry [Zha20] who obtains the best known tradeoff in ciphertext, public key and secret key size. In particular, Zhandry [Zha20] showed that all parameters can be of size $O(N^{1/3})$ by relying on the bilinear generic group model (GGM). Note that the generic group model is a strong assumption, and indeed a construction secure in this model cannot be considered as relying on standard assumptions, since several non-standard assumptions on pairings are secure in the GGM. Prior to [Zha20], Goyal et al. [GQWW19] provided a construction from LWE and Pairings, but their overall parameters are significantly worse – while their ciphertext can be arbitrarily small, $O(N^\epsilon)$, their public key is $O(N)$ and secret key is $O(N^c)$ for some large constant c^2 .

Thus, a central open question in TR is:

Can we construct collusion resistant Trace and Revoke with optimal parameters from concrete polynomial assumptions?

Embedding Identities. Traditionally, it was assumed that tracing the index $i \in [N]$ of a corrupt user is enough, and there is an external mapping, maintained by the content distributor or some other party which associates the number i to the identity of the user, i.e. name, national identity number and such, which is then used to ensure accountability. The work of Nishimaki, Wichs and Zhandry (henceforth NWZ) [NWZ16] argued that this assumption is problematic since it implies that a user must trust the content provider with her confidential information. Storing such a map is particularly worrisome in the setting of public tracing since the user either cannot map the recovered index to an actual person, or the index-identity map must be stored publicly.

NWZ provided an appealing solution to the above conundrum – they suggest that identifying information be embedded in the key of the user, so that if a coalition of traitors constructs a pirate decoder, the tracing algorithm can directly retrieve the identifying information from

²Zhandry [Zha20] states that the secret size in [GQWW19] is $O(N^2)$ but in fact the exponent is much larger due to the usage of arithmetic computations in NC^1 , which blows up the circuit size associated with the ABE secret keys.

one of the keys that was used to construct the decoder and no one needs to keep any records associating users to indices. Notably, the identities can live in an *exponential* sized space, which introduces significant challenges in the tracing procedure. Indeed, handling an exponential space in the tracing procedure is the key contribution of NWZ. They also provided constructions of traitor tracing as well as trace and revoke with embedded identities, denoted by EITT and EITR respectively, from various assumptions.

1.1 Prior Work: Embedded Identity Trace and Revoke.

In the *public* trace setting, the only work that achieves embedded identity trace and revoke (EITR) with full collusion resistance is that of NWZ. However, while it takes an important first step, the construction by NWZ suffers from the following drawbacks:

1. *Reliance on Subexponential Hardness Assumption.* The construction relies on indistinguishability obfuscation [BGI⁺01], which appears to be an inherently subexponential assumption as discussed above.
2. *Selective Security in Revocation List:* Despite relying on adaptive security of functional encryption, the notion of security achieved by their construction is selective – the adversary must announce the revocation list before making any key requests or seeing the challenge ciphertext.

In the *secret* trace setting, the work of Kim and Wu [KW20] achieves EITR from the subexponential Learning With Errors (LWE) assumption. However, their construction incurs a ciphertext size that grows with the size of the revocation list. Additionally, while they can achieve adaptive security with respect to the revocation list, this is either by incurring an exponential loss in the security proof, or by assuming sub-exponential security for an ingredient scheme.

1.2 Our Results

In this work, we provide the first constructions with optimal parameters from polynomial assumptions, which additionally support embedded identities from an exponential space. We detail our contributions below.

Public Trace Setting. We provide a construction of Trace and Revoke with public tracing which overcomes the limitations of NWZ – (i) it relies on polynomial hardness assumptions, namely functional encryption and “special” attribute based encryption, both of which can be constructed using standard polynomial hardness assumptions [BGG⁺14, JLS21, JLS22] (ii) it enjoys adaptive security in the revocation list.

A detailed comparison with prior work is provided in Table 1.

Our Assumptions. Functional Encryption (FE) and Attribute Based Encryption (ABE) are generalizations of Public Key Encryption. In FE, a secret key corresponds to a circuit C and a ciphertext corresponds to an input x from the domain of C . Given a function key sk_C and a ciphertext ct_x , the decryptor can learn $C(x)$ and nothing else. It has been shown that FE implies iO [AJ15, BV15] albeit with exponential loss. The aforementioned work of Jain, Lin and Sahai [JLS21, JLS22] provides a construction of compact FE from polynomial hardness assumptions, namely LPN, PRG in NC_0 and pairings. ABE is a special case of FE in which the input can be divided into a public and private part (x, m) and the circuit C in the secret key sk_C is only

Work	CT	SK	PK	Trace Space	Sel/Adp	Asspn	Identities
[NWZ16]	1	1	1	Exp	Selective	Subexp (iO)	Yes
[GVW19]	1	1	1	Poly	Adaptive	Subexp (subexp PWE)	No
This	1	1	1	Exp	Adaptive	Poly (FE and Special ABE)	Yes

Table 1: State of the art with Public Traceability.

evaluated on the public part x in the ciphertext $ct_{x,m}$. The private message m is revealed by decryption if and only if $C(x) = 1$. While FE implies ABE in general, we require our underlying ABE to satisfy special efficiency properties, which is not generically implied by FE. However, the desired ABE can be instantiated using the construction of Boneh et al. [BGG⁺14] which is based on LWE.

Secret Trace Setting. In the secret trace setting, we achieve the optimal size of $O(\log N)$ for ciphertext, public and secret key by relying on Lockable Obfuscation (LO) [GKW17, WZ17] and two special ABE schemes – one, the key-policy scheme with special efficiency properties by Boneh et al. [BGG⁺14] which is based on LWE, and two, a ciphertext-policy ABE for P which was recently constructed by Wee [Wee22] using the new evasive and tensor LWE assumption. The new assumptions by Wee, also independently discovered by Tsabary [Tsa22], are formulated for designing ciphertext-policy ABE which is much weaker than FE since ABE is an all or nothing primitive in contrast to FE. As such, these are believed to be much weaker than lattice based assumptions that have been introduced in the context of FE or iO.

Thus, we can replace the use of FE in the public trace setting by another ABE (though of type ciphertext-policy) and an LO in the private trace setting, albeit via a completely different construction. Along the way, we show that a small modification to the TR construction by Goyal et al. (Crypto 2019) yields a ciphertext of size $O(\log N)$ as against their original $O(N^\epsilon)$, from LWE and pairings. However, this construction retains the large public and secret keys of their construction, which depend at least linearly on N . Our results are summarized in Table 2.

Super-polynomial Revoke List. Lastly, by relying on subexponential security of LWE, both our constructions can support a *super-polynomial* sized revocation list, so long as it allows efficient representation and membership testing. Ours is the first work to achieve this, to the best of our knowledge.

1.3 Technical Overview

We proceed to give an overview of our techniques. We begin by defining the notion of revocable predicate encryption (RPE) in both the public and secret setting, then describe the ideas used to instantiate this primitive. Finally we outline how to upgrade public/secret RPE to build trace and revoke with embedded identities with public/secret tracing.

Work	CT	SK	PK	Trace Space	Asspn	Identities
[GQWW19]	N^ϵ	N^{poly}	N	Poly	LWE and Pairings	No
[Zha20]	N^a	N^{1-a}	N^{1-a}	Poly	GGM Pairings	No
[KW20]	L	1	1	Exp	Subexp LWE	Yes
This	1	1	1	Exp	Evasive Tensor LWE	Yes
Modified [GQWW19]	1	N^c	N	Poly	LWE and Pairings	No

Table 2: State of the Art with Secret Traceability. The column |CT| captures the dependence of ciphertext size on N and L where N denotes the number of users and L denotes the length of the revocation list. Parameters that are logarithmic in N , L or polynomial in the security parameter are represented as 1. Here, $0 < a < 1$ and $\epsilon > 0$ can be chosen arbitrarily. c is a large constant.

Revocable Predicate Encryption. NWZ introduced the notion of revocable functional encryption (RFE) and used it to construct EITR with public tracing. Subsequently, Kim and Wu [KW20] adapted this notion to the secret key setting, under the name of revocable predicate encryption (RPE) and used it to construct EITR with secret tracing. In this work, we extend Kim and Wu’s notion of RPE to the public key setting and use it to construct EITR with public tracing. Our notion of RPE in the public setting is similar to but weaker than RFE³ – it only supports “all or nothing” decryption in contrast to RFE. This weaker notion nevertheless suffices to construct EITR and moreover admits constructions from weaker assumptions.

In RPE, the key generation algorithm takes as input the master secret key msk , a label $\text{lb} \in \mathcal{L}$ and an attribute $x \in \mathcal{X}$. It outputs a secret key $\text{sk}_{\text{lb},x}$. The encryption algorithm takes as input the encryption key ek , a function f , a message $m \in \mathcal{M}$, and a revocation list $L \subseteq \mathcal{L}$. It outputs a ciphertext ct . Decryption recovers m if $f(x) = 1$ and $\text{lb} \notin L$. In the public variant of RPE, ek is a public key, while in the secret variant, ek is a secret key. In the secret variant, the scheme is also required to support a public “broadcast” functionality, i.e. there exists a public encryption algorithm that allows anyone to encrypt a message with respect to the “always-accept” policy, i.e. a policy that evaluates to true for all inputs. This is analogous to the primitive of “mixed FE” introduced by [GKW18].

In terms of security, we require RPE to satisfy message hiding and function hiding. At a high level, message hiding stipulates that an adversary cannot distinguish between encryptions of (f, m_0) and (f, m_1) as long as every key query for (lb, x) satisfies $f(x) = 0$ or $\text{lb} \in L$. Function hiding stipulates that an adversary cannot distinguish between encryptions of (f_0, m) and (f_1, m) as long as every key query for (lb, x) satisfies $f_0(x) = f_1(x)$ or $\text{lb} \in L$.

Before we describe our constructions, we highlight the chief difficulties that are inherent to designing RPE:

1. *Independence of parameter sizes from $|L|$.* A key requirement in TR schemes is that the ciphertext size should be independent of the length of the revocation list L – this constraint

³Syntactically, RPE is “ciphertext-policy” while RFE is “key-policy”, i.e. the function is embedded in the ciphertext in RPE as against the key in RFE.

must also be satisfied by the underlying RPE, in both the secret and public setting. In our work, we insist that even the public and secret keys satisfy $|L|$ independence. This constraint is inherited from broadcast encryption, and is challenging to satisfy. Further, note that L must be unbounded – its length cannot be fixed during setup, which introduces additional difficulties.

2. *Encrypted Computation.* While the revocation list L need not be hidden by the ciphertext, the function f^4 in the ciphertext is required to be hidden, as formalized by our function hiding requirement. Yet, this hidden function must participate in computing $f(x)$ where x is provided in the key. This requirement makes TR schemes worryingly close to collusion resistant functional encryption, an “obfustopia” primitive which we want to avoid in the secret trace setting.

Constructing Public Revocable Predicate Encryption. We proceed to describe the main ideas in constructing public RPE.

Overview of NWZ. The work of NWZ addresses the challenge of making the ciphertext size independent of $|L|$ by using a somewhere statistically binding (SSB) hash and hides the function f by using a functional encryption scheme, where f is encrypted in the ciphertext. However, they must additionally rely on iO – at a high level, this is because they require the *decryptor* to compute the SSB opening π and then run SSB verification on it (details of how SSB algorithms work are not relevant for this overview). In turn, the reason they need the decryptor to compute the opening π is because this needs both the set L and the label lb , which are available only to the decryptor – note that the encryptor has only L and the key generator has only lb . Now, since the decryptor has to compute π and run SSB verification, and since the program that computes SSB verification has some secrets, the decryptor is allowed to obtain obfuscation of this program. To implement this idea, they nest iO inside a compact FE scheme so that FE decryption outputs an iO which is then run by the decryptor on openings that it computes.

Trading iO for ABE. Above, note that the usage of iO is caused by the usage of SSB, which in turn is used to compress L . However, compression of a list has been achieved by much weaker primitives than iO in the literature of broadcast encryption – in particular, the construction of optimal broadcast encryption by Agrawal and Yamada uses the much weaker primitive of ABE (with special efficiency properties) to achieve this. However, ABE does not permit hiding anything other than a message, in particular, an ABE ciphertext cannot encrypt our function f since we desire f to participate in computation. ABE only permits computation on public values, and using ABE to encode f would force f to be public which we cannot allow.

In order to get around this difficulty, we leverage the power of functional encryption (FE), which permits encrypted computation and exactly fills the gap over ABE that we require. A natural candidate for RPE would be to simply use FE to encrypt f , L and m , and encode x and lb in the secret key for a functionality which tests that $lb \notin L$, that $f(x) = 1$ and outputs m if so. Indeed, this approach using FE is folklore, and was explicitly discussed by NWZ. Yet, they end up with a construction that additionally uses SSB, iO, a puncturable PRF and secret key encryption scheme because of the requirement of size independence from $|L|$ – we do not have candidates for FE with ciphertext size independent of the public attributes. In short, ABE gives us L compactness (in some cases by encoding L in the secret key [BGG⁺14] and in some cases by encoding L in the ciphertext [ALP11]) but does not hide f , whereas FE gives the opposite.

⁴For the informed reader, this function encodes the “index” and function hiding corresponds to “index hiding” in the literature.

Algebraic Synthesis of ABE and FE. We address this conundrum by combining the two primitives in a way that lets us get the best of both. In particular, we use ABE to check that $lb \notin L$ and use FE to compute $f(x)$. Evidently, the two steps cannot be performed independently in order to resist mix and match attacks so we use nesting, i.e. we use FE to generate ABE ciphertexts. Here, care is required, because ABE encryption takes L as input and done naively, this strategy will again induce a size dependence on L . We address this challenge by using the special ABE by Boneh et al. [BGG⁺14] which enjoys succinct secret keys and encoding L in the ABE secret key. In more detail, we let the RPE encryption generate $\text{ABE.sk}(C_L)$ for a circuit C_L which takes as input lb and checks that $lb \notin L$. Additionally, it generates an FE ciphertext for the function f and message m . The RPE key generator computes an FE key for a function which has (lb, x) hardwired and takes as input a function f , checks whether $f(x) = 1$ and if so, generates a fresh ABE ciphertext with attribute lb and message m . Thus the decryptor can first compute FE decryption to recover the ABE ciphertext $\text{ABE.ct}(lb, m)$ and then use ABE decryption with $\text{ABE.sk}(C_L)$ to output m if and only if $lb \notin L$. It is easy to verify that this construction achieves optimal parameters – this is because ABE has optimal parameters and we used FE only for a simple functionality that does not involve L .

Putting it all Together. The above description is over-simplified and ignores technical challenges such as how to leverage indistinguishability based security of FE, how to generate the randomness used for ABE encryption and such others – we refer the reader to Section 4 for details. However, even having filled in these details, we get only a selectively secure RPE. Substantial work and several new ideas are required for adaptive security, as we discuss next.

Adaptive Security. Next, we outline our ideas to achieve adaptive security, namely where the revocation list L is chosen adaptively by the adversary. Note that to avoid complexity leveraging, we are required to rely only on the selective security of the underlying ABE – this creates multiple technical difficulties which are resolved by very carefully using specific algebraic properties of our ingredients.

Leveraging Late Generation of ABE. Our first observation is that full adaptive security of ABE may be unnecessary, since in our construction of RPE, the generation of the ABE instance is deferred until the generation of the challenge ciphertext, at which time the set of revoked users is known. This intuition turns out to be true, but via a complicated security proof as we outline next. Below, we consider the case of function hiding in the RPE ciphertext, the case of message hiding is similar.

Recall that function hiding says that two ciphertexts encoding (f_b, m, L) , where $b \in \{0, 1\}$ should be indistinguishable so long as for any requested key $\text{sk}_{lb, x}$ it holds that $f_0(x) = f_1(x)$ or $lb \in L$. Note that the adversary is permitted to query for keys that allow decryption of the ciphertexts, i.e. $f_0(x) = f_1(x) = 1$.

Embedding ABE CTs in FE keys. In order to use ABE security to prove RPE security, a first (by now standard) step is to use the “trapdoor technique” [CIJ⁺13, ABSV15, BS18], which allows us to hardwire ABE ciphertexts into FE secret keys. In the security game with the ABE challenger, the reduction submits the label lb associated with each RPE secret key as its challenge attribute and embeds the returned ABE ciphertext into the FE key. Here we immediately run into a difficulty, since in the RPE setting some ABE ciphertexts are decryptable by the adversary and we cannot leverage ABE security. Moreover, we cannot even hope to guess which keys will correspond to decryptable ABE ciphertexts since there are an unbounded polynomial number of key queries in the RPE security game. The same difficulty is faced by NWZ and is the main reason why their construction does not achieve adaptive security in the revocation list.

Polynomial Function Space Suffices for TR. To overcome this hurdle, we leverage the serendipitous fact that for the purpose of constructing TR, it suffices to construct RPE whose function space (recall that functions are encoded in the ciphertext) is only of polynomial size. This observation, which was implicitly present in [GKW19], is abstracted and used explicitly in our proof. In particular, we can assume that the reduction algorithm knows the challenge functions (f_0, f_1) at the beginning of the game, since it can simply guess them. Now, given the secret key query (lb, x) , the reduction checks whether $f_0(x) = f_1(x)$. If yes, then there is no need to use ABE security, for the ABE ciphertexts in this case will encode the same message, and will hence be independent of the challenge bit. On the other hand, if $f_0(x) \neq f_1(x)$, then we have by the admissibility condition that $\text{lb} \in L$, even when L is not known. In this case, the reduction can use the security of the ABE without any difficulty.

Additional Hurdles Stemming from ABE Selective Security. We now highlight another challenge in the proof. For concreteness, let us consider the second key query $(\text{lb}^{(2)}, x^{(2)})$, which we assume is a pre-challenge query, and assume that $f_0(x^{(2)}) \neq f_1(x^{(2)})$. Hence, by the above discussion, we are required to use ABE security for the ciphertexts with attribute $\text{lb}^{(2)}$. However, according to the selective definition, the reduction is required to choose the challenge attribute at the very start of the game, without even seeing the public parameters. At the same time, the reduction is required to simulate the ABE ciphertext for the first key query, before receiving the second key query from the adversary, that is, without seeing the ABE parameters, leading to an apparent impasse.

We address this issue by considering the following two cases separately: for the first query $(\text{lb}^{(1)}, x^{(1)})$, we have (1) $f_0(x^{(1)}) \neq f_1(x^{(1)})$ or (2) $f_0(x^{(1)}) = f_1(x^{(1)})$. In first case, it is tempting to think that one can simply use a hybrid argument to change the ABE ciphertext associated with each key query satisfying $f_0(x^{(i)}) \neq f_1(x^{(i)})$ for $i \in [2]$. However, this does not work as is, since the ABE ciphertext may leak information about the ABE public key. To address this, we rely on the pseudorandomness of ciphertexts in our ABE [BGG⁺14] due to which we are guaranteed that the ciphertext does not reveal any information about the public parameters, enabling the hybrid strategy above. To handle the second case, we change the way in which the ABE ciphertext for the first key is generated. In more detail, we stop hardwiring the the ABE ciphertext into the first key and instead generate it directly using ABE parameters. This removes the aforementioned problem since we no longer need to embed the ABE ciphertext or public key into the first FE key. To enable this idea, we introduce additional branch of trapdoor mode for the construction to separate the paths of computation for the cases $f_0(x) = f_1(x)$ and $f_0(x) \neq f_1(x)$. To handle post-challenge queries, we need to address additional challenges, which we do not describe here. We refer the reader to Section 4 for details.

Handling Super-polynomial Revocation List. Our construction (also the secret version, described next) organically supports super-polynomially large revocation list, something that was not known before, to the best of our knowledge. In more detail, let L be a list of super-polynomial size, such that L can be represented as a string of polynomial length and there exists a circuit C_L of polynomial size which takes as input some string lb and checks whether $\text{lb} \in L$ or not. Note that any super-polynomially large list must have efficient representation in order to even allow various algorithms to read it. Then, the key generation of [BGG⁺14] can naturally encode the circuit C_L as before and the construction works as before. A subtlety that arises with super-polynomial L is that when we deal with post challenge key queries in the proof, we have to deal with the ABE queries in the order of key first and ciphertext later. With polynomial size L , this does not pose a problem because when the adversary chooses L , all the labels for which we use ABE security are in L and we can perform a hybrid argument over these labels. However, this is not possible for super polynomial L , which requires to rely on subexponentially secure

LWE. Please see Section 11 for details.

Instantiating Public RPE. Overall, armed with the above ideas, we get a public RPE from compact FE and efficient ABE supporting exponential sized identity space and adaptive security in the revocation list L . Currently, we only know how to instantiate our desired ABE from LWE [BGG⁺14], whereas FE can be instantiated in multiple different ways. A natural candidate would be the FE from standard assumptions [JLS21, JLS22] which relies on pairings, LPN and low depth PRG – in this case, our RPE will require the extra assumption of LWE. Another option is to instantiate FE with a post-quantum candidate [GP21, LPST16, WW21, Agr19, DQV⁺21] from non-standard strengthenings of LWE – this has the advantage that the ABE does not incur any extra assumption in the final construction. For super-polynomial L , we need subexponential hardness of LWE in either pathway to instantiation, as discussed above.

Revocable Predicate Encryption in Private Setting. For private revocable predicate encryption, our starting point is the work of Goyal et al. [GQWW19], who show how to combine “broadcast mixed FE” (called BMFE) together with ABE to achieve RPE (via a different abstraction which they call AugBE). They construct BMFE by adding the broadcast functionality to the primitive of mixed FE defined by [GKW18]. They embed BMFE ciphertext into an ABE ciphertext to achieve RPE, where BMFE is constructed from LWE and ABE is instantiated using pairings.

Supporting Exponential Identity Space. To begin, we upgrade their notion of BMFE to support an exponential space of identities (which we refer to as labels) towards the goal of embedded identity trace and revoke. We refer to our notion as Revocable Mixed FE (denoted by RMFE) and construct it from LWE. Both [GQWW19] and our work start with a mixed FE scheme and add broadcast to it, but their construction builds upon the scheme based on constrained PRFs [CVW⁺18] while ours begins with the scheme based on Lockable Obfuscation (LO), also from [CVW⁺18]. Our construction of RMFE deviates significantly from theirs, and achieves significantly better secret key size – $O(\log N)$ as against $O(N)$ – in addition to supporting exponential instead of polynomial space. We describe this construction next.

Mixed FE. The notion of mixed FE was introduced by Goyal, Koppula and Waters in the context of traitor tracing [GKW18]. Identifying and constructing this clever primitive is the key insight that enables [GKW18] to construct traitor tracing with optimal parameters from LWE. Mixed FE is, as the name suggests, a mix of public and secret key FE. Thus, it has a secret as well as a public encryption procedure. The secret encryption procedure takes as input a function f and computes ct_f . This is decryptable by a key sk_x to recover $f(x)$. The adversary can make one query to the encryption oracle in addition to getting the challenge ciphertext for challenge (f_0, f_1) . It can also make an unbounded number of key requests so long as $f_0(x) = f_1(x)$. The public encryption algorithm computes a ciphertext for the “always accept” function, i.e. a function which evaluates to 1 for any input x . It is required that the public ciphertext be indistinguishable from the secret ciphertext.

One of the constructions of mixed FE suggested by [CVW⁺18] uses a secret key FE scheme (SKFE) to construct the secret encryption algorithm and leverages the power of lockable obfuscation (LO) to construct the public encryption procedure. Recall that in a lockable obfuscation scheme [GKW17, WZ17] there exists an obfuscation algorithm Obf that takes as input a program C , a message m and a (random) “lock value” α and outputs an obfuscated program P . One can evaluate the obfuscated program on any input x to obtain as output m if $P(x) = \alpha$ and \perp otherwise. Intuitively, the idea of [CVW⁺18] is to wrap the FE ciphertext using LO and to define the public key encryption algorithm as outputting a simulated version of the LO obfuscated circuit, which is publicly sampleable.

In more detail, the construction works as follows. The secret key for a user with input x is an SKFE secret key $\text{SKFE.sk}(x)$. The secret ciphertext of MFE for function f is constructed as follows.

1. First, SKFE ciphertext $\text{SKFE.ct}(H_{f,\alpha})$ is generated, where α is a freshly chosen random value and $H_{f,\alpha}$ is a circuit that takes as input x and outputs α if $f(x) = 0$ and 0 otherwise.
2. Then, LO with lock value α and any message $m \neq \perp$ is used to obfuscate the circuit $\text{SKFE.Dec}(\text{SKFE.ct}(H_{f,\alpha}), \cdot)$, namely the circuit that takes as input an SKFE secret key and decrypts the hardwired ciphertext using this.

The decryption result of MFE is defined as 1 if the evaluation result of the LO circuit on the given input SKFE secret key is \perp and 0 otherwise. Correctness follows from correctness of SKFE and LO. In particular, if $f(x) = 0$, then SKFE decryption outputs α , which unlocks the LO to give m , otherwise \perp . By definition, MFE decryption will output 1 if LO outputs \perp which happens when $f(x) = 1$, and 0 otherwise.

Revocable Mixed FE. RMFE augments MFE so that the encryption algorithms (both secret and public) now include a revocation list L and the secret key additionally includes a label lb . A secret key $\text{sk}_{\text{lb},x}$ decrypts a secret ciphertext $\text{ct}_{f,L}$ to recover $f(x)$ if $\text{lb} \notin L$ and 1 otherwise. For a public ciphertext ct_L , the output of decryption is always 1 regardless of which secret key is being used. For security, we need two properties: function hiding and mode hiding. For function hiding, we require that a secret ciphertext $\text{ct}_{f_0,L}$ is indistinguishable from $\text{ct}_{f_1,L}$ if for all queries, either $f_0(x) = f_1(x)$ or $\text{lb} \in L$. For mode hiding, we require that a secret ciphertext $\text{ct}_{f,L}$ is indistinguishable from a public ciphertext ct_L . Recall that L is not required to be hidden, but we require that the parameters do not depend on $|L|$.

To extend MFE to RMFE, we retain the idea of letting the secret ciphertext be an LO obfuscated circuit and public ciphertext be the simulated LO. To incorporate the list L , we must ensure that the LO lock value α is recovered only when $f(x) = 0$ and $\text{lb} \notin L$. To do so, we consider two subsystems such that one system outputs partial decryption result α_1 only when $f(x) = 0$ and the second system outputs partial decryption result α_2 only when $\text{lb} \notin L$ such that $\alpha = \alpha_1 + \alpha_2$. We must ensure that α_1 and α_2 are user specific decryption results to avoid collusion attacks.

Note that the second subsystem, which entails L , should be constructed so that the hardwired values inside the circuit do not depend on $|L|$, but still control access to the value α_2 depending on L . To satisfy these apparently conflicting requirements, we make use of the unique algebraic properties of the ABE construction by Boneh et al [BGG⁺14], as described below. For the first subsystem, we use SKFE.

In more detail, our candidate scheme is as follows.

1. **Secret Key:** The RMFE secret key consists of $\text{ABE.ct}(\text{lb}, K)$ and $\text{SKFE.ct}((x, K, R))$ where K and R are user specific random strings, lb is used as an attribute and K is the plaintext for ABE encryption.
2. **Ciphertext:** To generate RMFE ciphertext, the secret key encryption procedure is as follows:
 - It first generates $\text{ABE.sk}(C_L)$, where C_L is a circuit that takes as input a label lb and outputs 1 only when $\text{lb} \notin L$.
 - It also generates $\text{SKFE.sk}(H_{f,\alpha})$, where $H_{f,\alpha}$ takes as input (x, K, R) and outputs $K \oplus \alpha$ if $f(x) = 0$ and R if $f(x) = 1$.

- Now, consider the circuit $CC[\text{ABE.sk}(C_L), \text{SKFE.sk}(H_{f,\alpha})]$, which takes as an input the pair $(\text{ABE.ct}, \text{SKFE.ct})$, decrypts both ABE and SKFE ciphertexts using their respective keys, and then outputs the XOR between the decryption results.
- The final ciphertext is an LO of $CC[\text{ABE.sk}(C_L), \text{SKFE.Enc}(H_{f,K,\alpha})]$ with lock value α and any arbitrary message $m \neq \perp$.

By key compactness of [BGG⁺14], the size of $\text{ABE.sk}(C_L)$ is independent of $|L|$. A subtle point here is that ABE decryption is happening inside the LO and this depends on L . If the LO must process L , then the size of the LO and hence ciphertext blows up with L ! Fortunately, the algebraic structure of the ABE scheme we use [BGG⁺14] again comes to our rescue. At a high level, ABE decryption can be divided into an “ L -dependent” step which results in a short processed ciphertext, followed by an “ L -independent” step. Importantly, the L -dependent step does not depend on the ABE secret key which is hardwired in the LO and hence inaccessible, and can hence be performed *outside* the LO by the decryptor! The resultant short processed ciphertext can then be provided as input to the LO preventing the problematic size blowup.

RMFE Proof Overview. Next we outline some of the ideas developed for the security proof. For ease of understanding, we limit ourselves to the simpler setting where the adversary does not have access to the encryption oracle. This restriction can be removed using combinatorial tricks, similar to [CVW⁺18]. For security, we must argue two properties – mode indistinguishability and function hiding. The former can be established by relying on security of SKFE and LO analogously to the MFE proof in [CVW⁺18]. Hence, we focus on function hiding for the rest of the overview, which is subtle and requires several new ideas.

For function hiding, we must make use of the security of ABE and SKFE. Intuitively, security of SKFE guarantees that the values encoded in SKFE ciphertexts and secret keys are hidden, beyond what is revealed by decryption.⁵ First note that given a key for (lb, x) such that $f_0(x) = f_1(x)$, no information about the challenge bit is revealed by decryption, since the decryption results of SKFE are the same for both cases. The case with $f_0(x) \neq f_1(x)$ is more challenging. Let us assume $f_0(x) = 0$ and $f_1(x) = 1$. In this case, the decryption result of the challenge ciphertext is R or $K \oplus \alpha$ depending on the value of the challenge bit. Since both are random strings, it is tempting to conclude that they do not reveal any information of the challenge bit.

However, in reality, information about K is encoded in the ciphertext $\text{ABE.ct}(\text{lb}, K)$ and creates a correlation which must be handled. Indeed, a computationally unbounded attacker can learn the challenge bit by breaking open the ABE ciphertext, recovering K and then correlating it with the decryption result of SKFE. Hence, security of ABE must play a role and fortunately, we show that security of ABE suffices to overcome this difficulty. Recall that our security definition of RMFE requires that if $f_0(x) \neq f_1(x)$, then it should hold that $\text{lb} \in L$. This means that the ciphertext $\text{ABE.ct}(\text{lb}, K)$ is computationally indistinguishable from $\text{ABE.ct}(\text{lb}, 0)$, since the only ABE secret key available to the adversary is $\text{ABE.sk}(C_L)$. Now, in the adversary’s view, both $K \oplus \alpha$ and R are random strings that are independent from other parameters. Therefore, the adversary cannot obtain any information of the challenge bit from the decryption result in this case as well. For more details, please see Section 5.

Comparison with the BMFE by Goyal et al. [GQWW19]. We observe that both our RMFE as well as the BMFE by [GQWW19] rely solely on LWE. However, our secret key is $\text{ABE.ct}(\text{lb}, K)$ and $\text{SKFE.ct}((x, K, R))$, which has optimal size, being clearly independent of N and L . In

⁵We note that we need message *and* function hiding security for the underlying SKFE, while [CVW⁺18] only needs message hiding security.

contrast their secret key depends linearly on N . We also observe that our RMFE can support an exponentially large space of identities, while their BMFE does not.

Combining RMFE and ABE to get RPE. Finally, we nest our RMFE inside an outer ABE scheme to obtain RPE. This step is very similar to [GQWW19], but we need to use a different ABE scheme. In particular, in the construction of RPE in [GQWW19], a key policy ABE (kpABE) is used to encrypt the message m with attributes as the RMFE ciphertext along with the list L . The RPE secret key for (x, lb) is a kpABE secret key for a the RMFE decryption circuit $\text{RMFE.Dec}(\text{RMFE.sk}, \cdot, \cdot)$.

An obvious difficulty here is that encoding the attribute $(L, \text{RMFE.ct})$ in the ABE ciphertext can cause the ciphertext size to depend on the size of L . To avoid this blowup, [GQWW19] use a special kpABE which has the property that the ciphertext size is independent of the size of the attribute. They instantiate this kpABE with the scheme [ALP11] which uses pairings⁶. However, we cannot use [ALP11, Tak14] because of the following two reasons:

1. First, the ABE scheme by [ALP11] only supports NC_1 . However, our circuit $\text{RMFE.Dec}(\text{RMFE.sk}, \cdot, \cdot)$ does not fit into NC_1 ⁷.
2. Furthermore, even if the above problem could be resolved, using [ALP11] is problematic since their ABE has secret and public keys at least as large as $O(|L|)$. While the scheme of [GQWW19] also suffers from this blow-up, our goal is to obtain short keys, independent of $|L|$.

The first problem cannot be resolved even if we use the ABE schemes for circuits [GVW13, BGG⁺14], since their ciphertext size also depends on $|L|$. To instantiate our ABE, we use recent construction of compact cpABE from evasive and tensor LWE [Wee22], whose parameter sizes depend only on the input length of the circuit and are independent of its size. Armed with the above ideas, we suggest the following RPE:

1. The encryption algorithm of RPE, given m, f, L computes RMFE ciphertext encoding (f, L) and then computes $\text{cpABE.Enc}(\text{RMFE.Dec}(\text{RMFE.ct}, \cdot, L), m)$.
2. The key generation algorithm RPE given (lb, x) , computes RMFE secret key for (lb, x) and outputs $\text{cpABE.sk}(\text{RMFE.sk})$.

Correctness of RPE follows from correctness of cpABE and RMFE while optimality of parameters follows from the efficiency of the underlying schemes. In particular, observe that all parameters are independent of $|L|$. Also note that evasive and tensor LWE are required only to instantiate cpABE with the desired efficiency. If future work standardizes the assumptions underlying the cpABE, our construction will inherit these assumptions. For more details, we refer the reader to Section 6.

Instantiating Secret RPE. Currently, the only two suitable ABE schemes that we know to instantiate our compiler are the LWE based kpABE by Boneh et al. [BGG⁺14] and the evasive and tensor LWE based cpABE by Wee [Wee22]. These two ABEs give us a secret RPE scheme supporting exponential identities and with optimal parameters, from evasive and tensor LWE. Note that this construction does *not* achieve adaptive security in the revoke list. Nevertheless, it is the first construction of optimal RPE, even without embedded identities, from any assumption

⁶In fact, one could instead use the kpABE constructed by [Tak14]. This enjoys the same efficiency properties and is based on the standard DLIN assumption as against the q -type assumption of [ALP11].

⁷The informed reader may wonder whether we can solve this issue by using preprocessing as in [GQWW19] but this does not work due to technical reasons.

outside Obfustopia. Note that the usage of a non-standard assumption outside of obfustopia (in particular, only from lattice techniques) is somewhat inherent given that even broadcast encryption *without* tracing requires non-standard assumption if we instantiate it only from lattices. We are hopeful that future improvements in cpABE will yield a construction from completely standard assumptions.

Trace and Revoke with Optimal Ciphertext from LWE and Pairings. Along the way, we observe that the broadcast and trace construction provided by Goyal et al. [GQWW19], without embedded identities, can be easily modified to achieve at least optimal ciphertext size, from the same assumptions. At a high level, they construct a broadcast mixed FE from LWE with optimal ciphertext size and then nest this inside the kpABE by [ALP11], which enjoys ciphertext size independent of the attribute length, and can support computation in NC_1 . Since their BMFE decryption does not fit into NC_1 , they preprocess the ciphertext so that part of the decryption is performed “outside”, namely, they group $\log N$ matrix tuples into c groups of $(\log N)/c$ tuples each. Then they precompute all possible $2^{(\log N)/c} = N^{1/c}$ subset-products within each group. Due to this, BMFE decryption only needs to multiply together c of the preprocessed matrices, which can be done in NC_1 so long as c is constant. Unfortunately, this step increases their ciphertext size to $O(N^\epsilon)$ for any $\epsilon > 0$ though the BMFE ciphertext size was optimal.

We observe that they are “under-using” the ciphertext size independence of [ALP11] – in particular, while the attribute length has indeed been blown up to $O(N^\epsilon)$, this does not affect the ciphertext size of [ALP11]. Moreover, while the attribute must also be provided outside in the clear, this part can be compressed, i.e. the preprocessing which expands the attribute to size N^ϵ can be performed by the decryptor directly by grouping and multiplying matrices as described above, and there is no need for the encryptor to provide this expanded form. Thus, their scheme tweaked with this simple modification already achieves ciphertext of optimal size, though with large secret key $O(N^\epsilon)$ for some large constant c .

Trace and Revoke from Revocable Predicate Encryption. It remains to show how to construct the final goal of trace and revoke with embedded identities. As discussed earlier, we follow [NWZ16, KW20] and use the abstraction of RPE to build trace and revoke. However, to embed identities in our trace and revoke schemes, we deviate from these works and instead build upon ideas developed by [GKW19] (henceforth GKW) in the context of traitor tracing.

Embedded Identity Traitor Tracing (EITT) by GKW. The work of Goyal, Koppula and Waters [GKW19] provided an alternative approach for embedding identities in traitor tracing schemes. A well known approach for constructing Traitor Tracing systems suggested by Boneh, Sahai and Waters [BSW07] is via the intermediate primitive of *Private Linear Broadcast Encryption* (PLBE), which allows to construct a tracing algorithm that performs a linear search over the space of users to recover the traitor. Since the number of users was polynomial, this algorithm could be efficient. However, if we allow arbitrary identities to correspond to user indices then the space over which this search must be performed becomes exponential even if the number of users is polynomial, and the trace algorithm is no longer efficient. The main new idea in NWZ that enables them to handle exponentially large identity spaces is to replace a linear search over indices by a clever generalization of binary search, which efficiently solves an “oracle jump problem” which in turn suffices for tracing.

Goyal, Koppula and Waters (GKW) provided an alternate route to the problem of embedding identities. Instead of using PLBE and generalizing the search procedure, they instead extend the definition of PLBE to support embedded identities, denoted by EI-PLBE, and then used this to get a full fledged EITT scheme. This approach has the notable advantage that even

if the space of identities is exponential, it can use the fact that the number of users is only *polynomial* and hence rely on only *selective* security of the underlying primitives. In particular, they demonstrate a “nested” tracing approach, where the tracing algorithm works in two steps: first, it outputs a set of indices that correspond to the users that are traitors, and then it uses each index within this set to recover the corresponding identity. Additionally, GKW provide a sequence of (increasingly stronger) TT primitives with embedded identities, namely, indexed EITT, bounded EITT and finally unbounded EITT where unbounded EITT satisfies the most general notion of embedded identity traitor tracing. They also provide generic transformations between these notions, which allows to focus on the weakest notion for any new instantiation.

Embedded Identity Trace and Revoke (EITR). We adapt the approach of GKW and show how to use their nested approach to trace embedded identities even in the more challenging setting of trace and revoke. As in their case, this lets us use polynomial hardness assumptions in obtaining EITR, in contrast to NWZ. We also define indexed, bounded and unbounded EITR and provide transformations between them. Our definitions as well as transformations are analogous to GKW albeit care is required to incorporate the revoke list L in each step and adapt the definitions and proofs of security accordingly. We then construct indexed EITR using RPE, and obtain unbounded EITR via our generic conversions.

We note that our framework unifies the approaches of Kim and Wu [KW20] who used the framework of RPE in the context of TR and that of GKW who used the framework of EI-PLBE in the context of TT, to obtain EITR. This unification yields a clean abstraction which can be used for both public and secret key settings. We believe this framework is of independent interest. We refer the reader to Sections 7, 8, 9 and 10 for details.

Organization of the paper We provide notation and preliminaries in Section 2. We define RPE in Section 3 and provide a construction of public-key RPE in Section 4. We give our construction of RMFE in Section 5. Then we construct secret-key RPE using RMFE in Section 6. We define different versions of trace and revoke with embedded identities in Section 7 and construct indexed-EITR in Section 8, bounded-EITR in Section 9 and unbounded-EITR in Section 10. Our goal is unbounded-EITR and other variants are introduced as intermediate goals. Secret and public tracing unbounded-EITR can be obtained by applying the conversions in Section 8, 9, and 10 to the secret-key and public key RPE, respectively. We show how to extend our constructions for super-polynomial sized revocation list in Section 11.

2 Preliminaries

In this section we define the notation and preliminaries used in our work.

Notation. We use bold letters to denote vectors and the notation $[a, b]$ to denote the set of integers $\{k \in \mathbb{N} \mid a \leq k \leq b\}$. We use $[n]$ to denote the set $[1, n]$. Concatenation is denoted by the symbol \parallel . We say a function $f(n)$ is *negligible* if it is $O(n^{-c})$ for all $c > 0$, and we use $\text{negl}(n)$ to denote a negligible function of n . We say $f(n)$ is *polynomial* if it is $O(n^c)$ for some constant $c > 0$, and we use $\text{poly}(n)$ to denote a polynomial function of n . We use the abbreviation PPT for probabilistic polynomial-time. We say an event occurs with *overwhelming probability* if its probability is $1 - \text{negl}(n)$. For two distributions X_λ and Y_λ , $X_\lambda \approx_c Y_\lambda$ denotes that they are computationally indistinguishable for any PPT algorithm. For a vector \mathbf{x} , we let x_i denote its i -th entry. For a set S , we let $|S|$ denote the number of elements in S . For a binary string x , we let $|x|$ denote the length of x .

Definition 2.1 (Pseudorandom Functions). A function $F : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$, with key space $\mathcal{K} = \{\mathcal{K}_\lambda\}_{\lambda \in [\mathbb{N}]}$, domain $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda \in [\mathbb{N}]}$, and range $\mathcal{Y} = \{\mathcal{Y}_\lambda\}_{\lambda \in [\mathbb{N}]}$ is a pseudorandom function if it satisfies the following properties:

- **Efficiency:** For all $k \in \mathcal{K}_\lambda$ and $x \in \mathcal{X}_\lambda$, $F(k, x)$ is efficiently computable.
- **Security:** There exists a negligible function negl such that for all PPT adversary \mathcal{A} , for all λ , the advantage of \mathcal{A} in the following security experiment is negl .
 $\text{Expt}_{\mathcal{A}}^{\text{PRF}}(\lambda)$

- The challenger samples a key $k \leftarrow \mathcal{K}_\lambda$ and a bit $b \in \{0, 1\}$.
- The adversary issues polynomially many queries of the following two types in any order:
 - Evaluation Queries:** \mathcal{A} outputs $x \in \mathcal{X}_\lambda$. The challenger returns $y = F(k, x)$.
 - Challenge Queries:** \mathcal{A} outputs $x \in \mathcal{X}_\lambda$. The challenger returns y_b , where $y_0 = F(k, x)$ and $y_1 \leftarrow \mathcal{Y}_\lambda$.
- In the end, \mathcal{A} outputs its guess bit b' .

\mathcal{A} wins the experiment if $b' = b$.

Remark 2.2. The above security notion is tailored to our purpose and may look stronger than more standard security notion, where the adversary is not allowed to make evaluation queries. However, we can easily show that more standard security notion implies the above by considering the following hybrid games. The first game is the same as above game with $b = 0$. Then, we consider the game where we change all the answers to both evaluation and challenge queries to be random. The game is indistinguishable from the previous one assuming the standard security notion. Finally, we consider a game where answers to the evaluation queries are changed to be $F(k, x)$, while answers to the challenge queries remain random. Again, using the standard security notion, this game is indistinguishable from the previous game. Furthermore, notice that this game is the same as the above game with $b = 1$.

Definition 2.3 (Symmetric Key Encryption). A symmetric key encryption scheme for message space $\mathcal{M} = \{\mathcal{M}_\lambda\}_{\lambda \in [\mathbb{N}]}$ and key space $\mathcal{K} = \{\mathcal{K}_\lambda\}_{\lambda \in [\mathbb{N}]}$ and ciphertext space $\mathcal{CT}_{\text{SKE}}$ has the following syntax:

$\text{Setup}(1^\lambda) \rightarrow \text{sk}$. The setup algorithm takes as input the security parameter λ and outputs a secret key sk .

$\text{Enc}(\text{sk}, m) \rightarrow \text{ct}$. The encryption algorithm takes as input the secret key sk and a message $m \in \mathcal{M}_\lambda$ and outputs a ciphertext ct .

$\text{Dec}(\text{sk}, \text{ct}) \rightarrow m'$. The decryption algorithm takes as input a secret key sk and a ciphertext ct and outputs a message $m' \in \mathcal{M}_\lambda$.

Correctness: A SKE scheme is said to be correct if there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, for every message $m \in \mathcal{M}_\lambda$, we have

$$\Pr \left[\begin{array}{l} m' = m : \\ \text{sk} \leftarrow \text{Setup}(1^\lambda); \\ \text{ct} \leftarrow \text{Enc}(\text{sk}, m); \\ m' = \text{Dec}(\text{sk}, \text{ct}). \end{array} \right] \geq 1 - \text{negl}(\lambda),$$

Security: A SKE scheme is said to have pseudorandom ciphertext if there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, for every message $m \in \mathcal{M}_\lambda$, we have

$$\Pr \left[\beta' = \beta : \begin{array}{l} \text{sk} \leftarrow \text{Setup}(1^\lambda); \\ \beta' \leftarrow \mathcal{A}^{\text{Enc}(\text{sk}, \cdot), \text{Enc}^\beta(\text{sk}, \cdot)} \end{array} \right] \leq 1/2 + \text{negl}(\lambda),$$

where the $\text{Enc}(\text{sk}, \cdot)$ oracle, on input a message m , returns $\text{Enc}(\text{sk}, m)$ and $\text{Enc}(\text{sk}, \cdot)$ oracle, on input a message m , returns ct_β , where $\text{ct}_0 \leftarrow \text{Enc}(\text{sk}, m)$ and $\text{ct}_1 \leftarrow \mathcal{CT}_{\text{SKE}}$.

Remark 2.4. We note that similarly to the case of PRF, the above security notion looks a bit different from more standard security notion where the adversary does not have access to $\text{Enc}(\text{sk}, \cdot)$ oracle. However, by a similar reduction to the case of PRF explained in Remark 2.2, it can be seen that they are equivalent.

2.1 Functional Encryption

Here, we recall the definition of public-key and secret-key functional encryption.

2.1.1 Public-Key Functional Encryption

Consider a function family $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$, with input space $\mathcal{M} = \{\mathcal{M}_\lambda\}_{\lambda \in \mathbb{N}}$ and output space $\mathcal{Y} = \{\mathcal{Y}_\lambda\}_{\lambda \in \mathbb{N}}$, i.e, $\mathcal{F}_\lambda = \{f : \mathcal{M}_\lambda \rightarrow \mathcal{Y}_\lambda\}$. A public-key functional encryption scheme FE for \mathcal{F} consists of four polynomial time algorithms (Setup, KeyGen, Enc, Dec):

$\text{Setup}(1^\lambda) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm takes as input the security parameter λ and outputs a public key mpk and a master secret key msk .

$\text{KeyGen}(\text{msk}, f) \rightarrow \text{sk}_f$. The key generation algorithm takes as input the master secret key msk and a function $f \in \mathcal{F}_\lambda$ and it outputs a functional secret key sk_f .

$\text{Enc}(\text{mpk}, \mathbf{m}) \rightarrow \text{ct}$. The encryption algorithm takes as input the public key mpk and a message $\mathbf{m} \in \mathcal{M}_\lambda$ and outputs a ciphertext ct .

$\text{Dec}(\text{sk}_f, \text{ct}) \rightarrow y$. The decryption algorithm takes as input a functional secret key sk_f and a ciphertext ct and outputs $y \in \mathcal{Y}_\lambda$.

Definition 2.5 (Correctness). A functional encryption scheme is said to be correct if there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, for every message $\mathbf{m} \in \mathcal{M}_\lambda$, and for every function $f \in \mathcal{F}_\lambda$, we have

$$\Pr \left[f(\mathbf{m}) \leftarrow \text{Dec}(\text{KeyGen}(\text{msk}, f), \text{Enc}(\text{mpk}, \mathbf{m})) \right] \geq 1 - \text{negl}(\lambda)$$

where $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$ and the probability is taken over the random coins of all algorithms.

Definition 2.6 (Selective Message Privacy). A functional encryption scheme over a function space $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$ is said to have selective message privacy (or simply is selectively secure) if for any PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, we have

$$\Pr \left[\begin{array}{l} \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{mpk}, \text{ct}) = b : \\ \begin{array}{l} (\mathbf{m}_0, \mathbf{m}_1) \leftarrow \mathcal{A}; \\ (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda); \\ b \leftarrow \{0, 1\}; \text{ct} \leftarrow \text{Enc}(\text{mpk}, \mathbf{m}_b) \end{array} \end{array} \right] \leq 1/2 + \text{negl}(\lambda),$$

where each key query for a function $f \in \mathcal{F}_\lambda$, queried by \mathcal{A} to the KeyGen oracle must satisfy the condition that $f(\mathbf{m}_0) = f(\mathbf{m}_1)$.

Compactness : We say that a FE scheme is compact if the running time of the encryption algorithm only depends on the security parameter and the input message length. In particular, it is independent of the complexity of the function class supported by the scheme.

Definition 2.7 (Fully Compact FE). A functional encryption scheme, $\text{FE} = (\text{FE.Setup}, \text{FE.KeyGen}, \text{FE.Enc}, \text{FE.Dec})$ for input space $\mathcal{M} = \{\mathcal{M}_\lambda\}_{\lambda \in \mathbb{N}}$ and function class $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$, where each \mathcal{F}_λ is a finite collection of functions, is said to be fully compact if the running time of the encryption algorithm FE.Enc , on input FE.mpk and a message $\mathbf{m} \in \mathcal{M}_\lambda$, is $\text{poly}(\lambda, |\mathbf{m}|)$.

Jain, Lin and Sahai [JLS21, JLS22] provided the first construction of FE with sublinear compactness, from standard assumptions. In more detail:

Lemma 2.8 ([JLS22, GS16, AJS15]). *There exists a public-key functional encryption scheme for polynomial sized circuits having selective security (as per Definition 2.6) and full compactness (as per Definition 2.7) with encryption time $\text{poly}(\lambda, |\mathbf{m}|)$, where $\lambda \in \mathbb{N}$ is the security parameter, \mathbf{m} is the input message, assuming LPN, DLIN and existence of boolean PRGs in NC^0 .*

2.1.2 Secret Key Functional Encryption

A secret key functional encryption scheme is the same as the public key functional encryption scheme, except that the setup algorithm only outputs msk and the encryption algorithm takes the master secret key msk as input, instead of the master public key mpk .

Definition 2.9 (Ada-IND Function-Message Privacy ([BS18], adapted)). A SKFE $= (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ scheme over an input space $\mathcal{M} = \{\mathcal{M}_\lambda\}_{\lambda \in \mathbb{N}}$ and a function space $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$ is said to be Ada-IND function and message private if for any PPT algorithm \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that

$$|\Pr[\text{Expt}_{\mathcal{A}}(\lambda, 0) = 1] - \Pr[\text{Expt}_{\mathcal{A}}(\lambda, 1) = 1]| \leq \text{negl}(\lambda)$$

where for each $b \in \{0, 1\}$ and $\lambda \in \mathbb{N}$, the $\text{Expt}_{\mathcal{A}}(\lambda, b)$ is defined as follows

1. $\text{msk} \leftarrow \text{Setup}(1^\lambda)$.
2. $b' \leftarrow \mathcal{A}^{\text{KeyGen}_b(\text{msk}, \cdot, \cdot), \text{Enc}_b(\text{msk}, \cdot, \cdot)}$, where the oracle $\text{KeyGen}_b(\text{msk}, \cdot, \cdot)$ on input (f_0, f_1) outputs $\text{KeyGen}(\text{msk}, f_b)$ and $\text{Enc}_b(\text{msk}, \cdot, \cdot)$ on input $(\mathbf{m}_0, \mathbf{m}_1)$ outputs $\text{Enc}(\text{msk}, \mathbf{m}_b)$.
3. Output b' .

and \mathcal{A} is admissible only if for all the key queries $(f_0, f_1) \in \mathcal{F} \times \mathcal{F}$ and encryption queries $(\mathbf{m}_0, \mathbf{m}_1) \in \mathcal{M} \times \mathcal{M}$, it must hold that $f_0(\mathbf{m}_0) = f_1(\mathbf{m}_1)$.

Note: We refer to the t -bounded Ada-IND function and message private SKFE scheme where Def. 2.9 holds only if at most t queries are made to the $\text{KeyGen}_b(\text{msk}, \cdot, \cdot)$ oracle.

Remark 2.10. We can construct a SKFE scheme satisfying t -bounded Ada-IND message privacy (without function privacy, i.e., $f_0 = f_1$ should hold in the security game) from one-way functions [AV19]. This implies a SKFE scheme satisfying t -bounded Ada-IND function-message privacy, using the conversion results from [BS18].

We also need a 2-bounded semi-adaptive simulation based function-message private SKFE scheme, defined next.

Definition 2.11 (Simulation Based Function-Message Privacy). A secret-key functional encryption scheme is said to satisfy t -bounded semi-adaptive simulation based function and message privacy, if for all PPT stateful algorithm \mathcal{A} , there exists PPT stateful algorithms $\text{Sim}^{\text{SK}}, \text{Sim}^{\text{CT}}$ such that:

$$\{\text{Exp}_{\mathcal{A}}^{\text{real}}(1^\lambda)\}_{\lambda \in \mathbb{N}} \approx_c \{\text{Exp}_{\mathcal{A}, \text{Sim}^{\text{SK}}, \text{Sim}^{\text{CT}}}^{\text{ideal}}(1^\lambda)\}_{\lambda \in \mathbb{N}}$$

where the real and ideal experiments of stateful algorithms $\mathcal{A}, \text{Sim}^{\text{SK}}, \text{Sim}^{\text{CT}}$ are as follows:

$\text{Exp}_{\mathcal{A}}^{\text{real}}(1^\lambda)$ $\text{msk} \leftarrow \text{Setup}(1^\lambda)$ For $i \in [t]$: $\mathcal{A} \rightarrow f_i \in \mathcal{F};$ $\mathcal{A} \leftarrow \text{sk}_{f_i} = \text{KeyGen}(\text{msk}, f_i)$ Repeat polynomially many times: $\mathcal{A} \rightarrow \mathbf{m} \in \mathcal{M};$ $\mathcal{A} \leftarrow \text{Enc}(\text{msk}, \mathbf{m})$ $\mathcal{A} \rightarrow b; \text{ Output } b$	$\text{Exp}_{\mathcal{A}, \text{Sim}}^{\text{ideal}}(1^\lambda)$ $\text{msk} \leftarrow \text{Setup}(1^\lambda)$ For $i \in [t]$: $\mathcal{A} \rightarrow f_i \in \mathcal{F};$ $\mathcal{A} \leftarrow \text{sk}_{f_i} = \text{Sim}^{\text{SK}}(\text{msk}, 1^{ f_i })$ Repeat polynomially many times: $\mathcal{A} \rightarrow \mathbf{m} \in \mathcal{M};$ $\mathcal{A} \leftarrow \text{Sim}^{\text{CT}}(\text{msk}, \{f_i(\mathbf{m})\}_{i \in [t]})$ $\mathcal{A} \rightarrow b; \text{ Output } b$
---	--

Lemma 2.12. *There exists a 2-bounded semi-adaptive simulation based function and message private SKFE scheme, assuming the existence of one-way functions.*

Proof. Let $\Pi = (\Pi.\text{Setup}, \Pi.\text{KeyGen}, \Pi.\text{Enc}, \Pi.\text{Dec})$ be a 2-bounded Ada-IND function-message private SKFE scheme. We construct a 2-bounded semi-adaptive simulation based function-message private SKFE = (Setup, KeyGen, Enc, Dec) scheme as follows:

1. Setup(1^λ): Sample $\text{msk} \leftarrow \Pi.\text{Setup}(1^\lambda)$.
2. KeyGen(msk, f): $\text{sk} \leftarrow \Pi.\text{KeyGen}(\text{msk}, E[f, 0])$. Here, $E[f, \text{mode}]$ is defined as $E[f, \text{mode}](\mathbf{m}, y_1, y_2) = y_{\text{mode}}$ for $\text{mode} \in \{0, 1, 2\}$, where $y_0 := f(\mathbf{m})$.
3. Enc(msk, \mathbf{m}): $\text{ct} \leftarrow \Pi.\text{Enc}(\text{msk}, (\mathbf{m}, \perp, \perp))$.
4. Dec(sk, ct): $\mathbf{m}' \leftarrow \Pi.\text{Dec}(\text{sk}, \text{ct})$.

Simulators for encryption and keygen are described as follows:

$\text{Sim}^{\text{CT}}(\text{msk}, \{f_1(\mathbf{m}), f_2(\mathbf{m})\})$ outputs $\Pi.\text{Enc}(\text{msk}, (\perp, f_1(\mathbf{m}), f_2(\mathbf{m})))$.

$\text{Sim}^{\text{SK}}(\text{msk}, 1^{|f|})$ outputs $\Pi.\text{KeyGen}(\text{msk}, E[1^{|f|}, b])$ for the b -th key query, for $b \in \{1, 2\}$.

Security:

Firstly, observe that the simulator thus constructed is valid since its output does not depend on the function f or input \mathbf{m} . Then proof of security follows directly from Ada-IND security of Π , because for any set of queries consisting of two functions f_1 and f_2 and polynomial many inputs $\{\mathbf{m}_i\}$, $f_b(\mathbf{m}_i) = g_b(\mathbf{z}_i)$ for $b \in \{1, 2\}$, where g_b is the circuit defined as $(E[1^{|f_b|}, b])$ and $\mathbf{z}_i = (\perp, f_1(\mathbf{m}_i), f_2(\mathbf{m}_i))$ used by the simulators to generate the simulated keys and ciphertexts, respectively. \square

2.2 Attribute Based Encryption

We define both ciphertext policy attribute-based encryption (cpABE) and key policy attribute-based encryption (kpABE) in a unified form below.

Let $R = \{R_\lambda : A_\lambda \times B_\lambda \rightarrow \{0, 1\}\}_{\lambda \in \mathbb{N}}$ be a relation where A_λ and B_λ denote “ciphertext attribute” and “key attribute” spaces. An attribute-based encryption (ABE) scheme for R and a message space $\mathcal{M} = \{\mathcal{M}_\lambda\}_{\lambda \in \mathbb{N}}$ is defined by the following PPT algorithms:

$\text{Setup}(1^\lambda) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm takes as input the unary representation of the security parameter λ and outputs a master public key mpk and a master secret key msk .

$\text{Enc}(\text{mpk}, X, \mu) \rightarrow \text{ct}$. The encryption algorithm takes as input a master public key mpk , a ciphertext attribute $X \in A_\lambda$, and a message $\mu \in \mathcal{M}_\lambda$. It outputs a ciphertext ct .

$\text{KeyGen}(\text{msk}, Y) \rightarrow \text{sk}_Y$. The key generation algorithm takes as input the master secret key msk and a key attribute $Y \in B_\lambda$. It outputs a private key sk_Y .

$\text{Dec}(\text{mpk}, \text{sk}_Y, Y, \text{ct}, X) \rightarrow \mu$ or \perp . The decryption algorithm takes as input the master public key mpk , a private key sk_Y , private key attribute $Y \in B_\lambda$, a ciphertext ct and ciphertext attribute $X \in A_\lambda$. It outputs the message μ or \perp which represents that the ciphertext is not in a valid form.

Definition 2.13 (Correctness). An ABE scheme for relation family R is correct if for all $\lambda \in \mathbb{N}$, $X \in A_\lambda, Y \in B_\lambda$ such that $R(X, Y) = 1$, and for all messages $\mu \in \mathcal{M}_\lambda$,

$$\Pr \left[\begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda), \\ \text{sk}_Y \leftarrow \text{KeyGen}(\text{msk}, Y), \\ \text{ct} \leftarrow \text{Enc}(\text{mpk}, X, \mu) : \\ \text{Dec}(\text{mpk}, \text{sk}_Y, Y, \text{ct}, X) \neq \mu \end{array} \right] = \text{negl}(\lambda)$$

where the probability is taken over the coins of Setup , KeyGen , and Enc .

Definition 2.14 (Sel-IND security for ABE). For an ABE scheme $\text{ABE} = \{\text{Setup}, \text{Enc}, \text{KeyGen}, \text{Dec}\}$ for a relation family $R = \{R_\lambda : A_\lambda \times B_\lambda \rightarrow \{0, 1\}\}_{\lambda \in \mathbb{N}}$ and a message space $\{\mathcal{M}_\lambda\}_{\lambda \in \mathbb{N}}$ and an adversary \mathcal{A} , let us define Sel-IND security game as follows.

1. \mathcal{A} outputs the challenge ciphertext attribute $X^* \in A_\lambda$.
2. **Setup phase:** On input 1^λ , the challenger samples $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$ and gives mpk to \mathcal{A} .
3. **Query phase:** During the game, \mathcal{A} adaptively makes the following queries, in an arbitrary order. \mathcal{A} can make unbounded many key queries, but can make only single challenge query.
 - (a) **Key Queries:** \mathcal{A} chooses an input $Y \in B_\lambda$. For each such query, the challenger replies with $\text{sk}_Y \leftarrow \text{KeyGen}(\text{msk}, Y)$.
 - (b) **Challenge Query:** At some point, \mathcal{A} submits a pair of equal length messages $(\mu_0, \mu_1) \in \mathcal{M}^2$ to the challenger. The challenger samples a random bit $b \leftarrow \{0, 1\}$ and replies to \mathcal{A} with $\text{ct} \leftarrow \text{Enc}(\text{mpk}, X^*, \mu_b)$.

We require that $R(X^*, Y) = 0$ holds for any Y such that \mathcal{A} makes a key query for Y in order to avoid trivial attacks.

4. **Output phase:** \mathcal{A} outputs a guess bit b' as the output of the experiment.

We define the advantage $\text{Adv}_{\text{ABE}, \mathcal{A}}^{\text{Sel-IND}}(1^\lambda)$ of \mathcal{A} in the above game as

$$\text{Adv}_{\text{ABE}, \mathcal{A}}^{\text{Sel-IND}}(1^\lambda) := \left| \Pr[\text{Exp}_{\text{ABE}, \mathcal{A}}(1^\lambda) = 1 | b = 0] - \Pr[\text{Exp}_{\text{ABE}, \mathcal{A}}(1^\lambda) = 1 | b = 1] \right|.$$

The ABE scheme ABE is said to satisfy Sel-IND security (or simply *selective security*) if for any stateful PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that $\text{Adv}_{\text{ABE}, \mathcal{A}}^{\text{Sel-IND}}(1^\lambda) = \text{negl}(\lambda)$.

We can consider the following stronger version of the security where we require the ciphertext to be pseudorandom.

Definition 2.15 (Sel-INDr security for ABE). We define Sel-INDr security game similarly to Sel-IND security game except that the adversary \mathcal{A} chooses single message μ instead of (μ_0, μ_1) at the challenge phase and the challenger returns $\text{ct} \leftarrow \text{Enc}(\text{mpk}, X^*, \mu)$ if $b = 0$ and a random ciphertext $\text{ct} \leftarrow \mathcal{CT}$ from a ciphertext space \mathcal{CT} if $b = 1$. Here, we assume that uniform sampling from the ciphertext space \mathcal{CT} is possible without any parameter other than the security parameter λ . We define the advantage $\text{Adv}_{\text{ABE}, \mathcal{A}}^{\text{Sel-INDr}}(1^\lambda)$ of the adversary \mathcal{A} accordingly and say that the scheme satisfies Sel-INDr security if the quantity is negligible.

We also consider adaptive version of the security.

Definition 2.16 (Ada-IND security for ABE). We define Ada-IND security game similarly to Sel-IND security game except that the adversary \mathcal{A} can choose the challenge ciphertext attribute X^* adaptively. We define the advantage $\text{Adv}_{\text{ABE}, \mathcal{A}}^{\text{Ada-IND}}(1^\lambda)$ of the adversary \mathcal{A} accordingly and say that the scheme satisfies Ada-IND security if the quantity is negligible.

In the following, we recall definitions of various ABEs by specifying the relation.

Ciphertext-policy Attribute Based encryption (cpABE). We define cpABE for circuit class $\{\mathcal{C}_{\ell(\lambda), d(\lambda)}\}_\lambda$ by specifying the relation. Here, $\mathcal{C}_{\ell(\lambda), d(\lambda)}$ is a set of circuits with binary output whose input length is $\ell(\lambda)$ and the depth is at most $d(\lambda)$. Note that we do not pose any restriction on the size of the circuits. We define $A_\lambda^{\text{cpABE}} = \mathcal{C}_{\ell(\lambda), d(\lambda)}$ and $B_\lambda^{\text{cpABE}} = \{0, 1\}^\ell$. Furthermore, we define the relation R_λ^{cpABE} as

$$R_\lambda^{\text{cpABE}}(C, \mathbf{x}) = C(\mathbf{x}).$$

Key-policy Attribute Based encryption (kpABE). To define kpABE for circuits, we simply swap key and ciphertext attributes in cpABE for circuits. More formally, to define kpABE for circuits, we define $A_\lambda^{\text{kpABE}} = \{0, 1\}^\ell$ and $B_\lambda^{\text{kpABE}} = \mathcal{C}_{\ell(\lambda), d(\lambda)}$. We also define $R_\lambda^{\text{kpABE}} : A_\lambda^{\text{kpABE}} \times B_\lambda^{\text{kpABE}} \rightarrow \{0, 1\}$ as

$$R_\lambda^{\text{kpABE}}(\mathbf{x}, C) = C(\mathbf{x}).$$

Boneh et al. [BGG⁺14] provided a construction of kpABE which satisfies key compactness and ciphertext succinctness. The following theorem summarizes the efficiency properties of their construction.

Theorem 2.17 (Properties of [BGG⁺14]). *There exists a key-policy ABE scheme kpABE = (kpABE.Setup, kpABE.KeyGen, kpABE.Enc, kpABE.Dec) for function class $\mathcal{C}_{\ell, d}$ which is selectively secure under the LWE assumption and has the following properties. In particular:*

Key Compactness. We have $|\text{ABE.sk}_C| \leq \text{poly}(\lambda, d)$ for any $C \in \mathcal{C}_{\ell, d}$, where $(\text{ABE.mpk}, \text{ABE.msk}) \leftarrow \text{ABE.Setup}(1^\lambda)$ and $\text{ABE.sk}_C \leftarrow \text{ABE.KeyGen}(\text{ABE.msk}, C)$. In particular, the length of the secret key is independent of the attribute length ℓ and the size of the circuit C .

Parameters Succinctness. We have $|\text{ABE.mpk}|, |\text{ABE.msk}| \leq \text{poly}(\lambda, d, \ell)$ and $|\text{ABE.ct}| \leq \text{poly}(\lambda, d, \ell) + |\mu|$ for any $\mathbf{x} \in \mathcal{X}_\lambda$ and $\mu \in \mathcal{M}_\lambda$, where $(\text{ABE.mpk}, \text{ABE.msk}) \leftarrow \text{ABE.Setup}(1^\lambda)$ and $\text{ABE.ct} \leftarrow \text{ABE.Enc}(\text{ABE.mpk}, \mathbf{x}, \mu)$.

Online-Offline Decryption. The decryption algorithm $\text{Dec}(\text{mpk}, \text{sk}_C, C, \text{ct}_\mathbf{x}, \mathbf{x})$ can be divided into two parts, which we call

- $\text{Dec}^{\text{off}}(\text{mpk}, C, \mathbf{x}) \rightarrow \text{off}$, which performs the heavier computation that involves the circuit C and attribute \mathbf{x} offline without knowing the ciphertext $\text{ct}_\mathbf{x}$ or the secret key sk_C to get a "help" off . We have that the size of off is $\text{poly}(\lambda, \ell, d)$ and the depth of the circuit $\text{Dec}^{\text{off}}(\cdot, \cdot, \cdot)$ is bounded by $\text{poly}(\lambda, \ell, \text{depth}(C))$, where $\text{depth}(C)$ is the depth of the circuit C .
- $\text{Dec}^{\text{on}}(\text{sk}_C, \text{ct}_\mathbf{x}, \text{off})$ takes the help off generated offline along with the secret key sk_C and the ciphertext $\text{ct}_\mathbf{x}$ and outputs the underlying message μ . We note that this part does not take C as input and in particular, the size of the circuit $\text{Dec}^{\text{on}}(\cdot, \cdot, \cdot)$ is $\text{poly}(\lambda, \ell, d)$, which is independent from the size of the circuit C .

We will provide the detail of the kpABE scheme given by [BGG⁺14] in Sec. 2.3. There, we will show that the scheme satisfies the online-offline decryption property defined above.

We will also use the cpABE scheme given by [Wee22]. The following theorem summarizes the properties of the scheme.

Theorem 2.18 (Properties of [Wee22]). *There exists a ciphertext policy ABE scheme $\text{cpABE} = (\text{cpABE.Setup}, \text{cpABE.KeyGen}, \text{cpABE.Enc}, \text{cpABE.Dec})$ for function class $\mathcal{C}_{\ell, d}$, which is selectively secure under the evasive and tensor LWE assumption and has the following properties. In particular:*

Ciphertext Compactness. We have $|\text{cpABE.ct}| \leq \text{poly}(\lambda, d) + |\mu|$ for any $C \in \mathcal{C}_{\ell, d}$ and $\mu \in \mathcal{M}_\lambda$, where $(\text{cpABE.mpk}, \text{cpABE.msk}) \leftarrow \text{cpABE.Setup}(1^\lambda)$ and $\text{cpABE.ct} \leftarrow \text{cpABE.Enc}(\text{cpABE.mpk}, C, \mu)$. In particular, the size of the ciphertext is independent from the size of the circuit C and its input length.

Parameters Succinctness. We have $|\text{cpABE.mpk}|, |\text{cpABE.msk}|, |\text{cpABE.sk}_\mathbf{x}| \leq \text{poly}(\lambda, \ell, d)$ for any $\mathbf{x} \in \{0, 1\}^\ell$, where $(\text{cpABE.mpk}, \text{cpABE.msk}) \leftarrow \text{cpABE.Setup}(1^\lambda)$ and $\text{cpABE.sk}_\mathbf{x} \leftarrow \text{cpABE.KeyGen}(\text{cpABE.msk}, \mathbf{x})$.

2.3 Key-Policy ABE by Boneh et al. [BGG⁺14]

We will use the kpABE scheme proposed by Boneh et al. [BGG⁺14]. We provide the description of the scheme in the following while showing that the decryption algorithm can be divided into two phases as stated in Theorem 2.17. We note that the presentation here is largely based on [AY20].

2.3.1 Lattice Preliminaries

Here, we introduce necessary backgrounds for presenting the scheme.

Trapdoors. Let $\text{SampZ}(\gamma)$ be an output of discrete Gaussian distribution with parameter γ over \mathbb{Z} . Let us consider a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$. For all $\mathbf{V} \in \mathbb{Z}_q^{n \times m'}$, we let $\mathbf{A}_\gamma^{-1}(\mathbf{V})$ be an output of $\text{SampZ}(\gamma)^{m \times m'}$ conditioned on $\mathbf{A} \cdot \mathbf{A}_\gamma^{-1}(\mathbf{V}) = \mathbf{V}$. A γ -trapdoor for \mathbf{A} is a trapdoor that enables one to sample from the distribution $\mathbf{A}_\gamma^{-1}(\mathbf{V})$ in time $\text{poly}(n, m, m', \log q)$ for any \mathbf{V} . We slightly overload notation and denote a γ -trapdoor for \mathbf{A} by \mathbf{A}_γ^{-1} . We also define the special gadget matrix $\mathbf{G} \in \mathbb{Z}_q^{n \times m}$ as the matrix obtained by padding $\mathbf{I}_n \otimes (1, 2, 4, 8, \dots, 2^{\lceil \log q \rceil})$ with

zero-columns. The following properties had been established in a long sequence of works. We refer to [AY20] and references therein for the details.

Lemma 2.19 (Properties of Trapdoors). *Lattice trapdoors exhibit the following properties.*

1. Given \mathbf{A}_{τ}^{-1} , one can obtain $\mathbf{A}_{\tau'}^{-1}$ for any $\tau' \geq \tau$.
2. Given \mathbf{A}_{τ}^{-1} , one can obtain $[\mathbf{A} \parallel \mathbf{B}]_{\tau}^{-1}$ and $[\mathbf{B} \parallel \mathbf{A}]_{\tau}^{-1}$ for any \mathbf{B} .
3. There exists an efficient procedure $\text{TrapGen}(1^n, 1^m, q)$ that outputs $(\mathbf{A}, \mathbf{A}_{\tau_0}^{-1})$ where $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ for some $m = O(n \log q)$ and is 2^{-n} -close to uniform, where $\tau_0 = \omega(\sqrt{n} \log q \log m)$.

Lattice Evaluation. The following is an abstraction of the evaluation procedure in previous LWE based FHE and ABE schemes. We follow the presentation by Tsabary [Tsa19].

Lemma 2.20 (Fully Homomorphic Computation [BGG⁺14]). *There exists a pair of deterministic algorithms $(\text{EvalF}, \text{EvalFX})$ with the following properties.*

- $\text{EvalF}(\mathbf{B}, F) \rightarrow \mathbf{H}_F$. Here, $\mathbf{B} \in \mathbb{Z}_q^{n \times m \ell}$ and $F : \{0, 1\}^{\ell} \rightarrow \{0, 1\}$ is a circuit.
- $\text{EvalFX}(F, \mathbf{x}, \mathbf{B}) \rightarrow \widehat{\mathbf{H}}_{F, \mathbf{x}}$. Here, $\mathbf{x} \in \{0, 1\}^{\ell}$ and $F : \{0, 1\}^{\ell} \rightarrow \{0, 1\}$ is a circuit with depth d . We have $[\mathbf{B} - \mathbf{x} \otimes \mathbf{G}] \widehat{\mathbf{H}}_{F, \mathbf{x}} = \mathbf{B} \mathbf{H}_F - F(\mathbf{x}) \mathbf{G} \pmod{q}$, where we denote $[x_1 \mathbf{G} \parallel \dots \parallel x_k \mathbf{G}]$ by $\mathbf{x} \otimes \mathbf{G}$. Furthermore, we have $\|\mathbf{H}_F\|_{\infty} \leq m^{O(d)}$ and $\|\widehat{\mathbf{H}}_{F, \mathbf{x}}\|_{\infty} \leq m^{O(d)}$.
- The running time of $(\text{EvalF}, \text{EvalFX})$ is bounded by $\text{poly}(n, m, \log q, d)$.

Note that the last item implies that the circuits computing EvalF and EvalFX can be implemented with depth $\text{poly}(n, m, \log q, d)$.

2.3.2 Key-Policy ABE by Boneh et al. [BGG⁺14]

The scheme supports the circuit class $\mathcal{C}_{\ell(\lambda), d(\lambda)}$, which is the set of all circuits with input length $\ell(\lambda)$ and depth at most $d(\lambda)$ with arbitrary $\ell(\lambda) = \text{poly}(\lambda)$ and $d(\lambda) = \text{poly}(\lambda)$. In our case, we will set $d(\lambda) = \omega(\log \lambda)$. In the description below, we focus on the case where the message space is $\{0, 1\}$ for simplicity. To encrypt a long message, we run the construction in parallel to encrypt an SKE key $K \in \{0, 1\}^{\lambda}$ and then use it to encrypt the message.

Setup(1^{λ}): On input 1^{λ} , the setup algorithm defines the parameters $n = n(\lambda)$, $m = m(\lambda)$, noise distributions χ over \mathbb{Z} , $\tau_0 = \tau_0(\lambda)$, $\tau = \tau(\lambda)$, and $B = B(\lambda)$ as specified later. It then proceeds as follows.

1. Sample $(\mathbf{A}, \mathbf{A}_{\tau_0}^{-1}) \leftarrow \text{TrapGen}(1^n, 1^m, q)$ such that $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$.
2. Sample random matrix $\mathbf{B} = (\mathbf{B}_1, \dots, \mathbf{B}_{\ell}) \leftarrow (\mathbb{Z}_q^{n \times m})^{\ell}$ and a random vector $\mathbf{u} \leftarrow \mathbb{Z}_q^n$.
3. Output the master public key $\text{mpk} = (\mathbf{A}, \mathbf{B}, \mathbf{u})$ and the master secret key $\text{msk} = \mathbf{A}_{\tau_0}^{-1}$.

KeyGen(msk, C): The key generation algorithm takes as input the master public key mpk , the master secret key msk , and a circuit $C \in \mathcal{C}_{\ell, d}$ and proceeds as follows.

1. Set $F := -C$ to be the same circuit as C except that the output bit is flipped.⁸
2. Compute $\mathbf{H}_F = \text{EvalF}(\mathbf{B}, F)$ and $\mathbf{B}_F = \mathbf{B} \mathbf{H}_F$.

⁸While we follow the standard definition of ABE and require the decryption to be possible when $C(x) = 1$, it is when $C(\mathbf{x}) = 0$ in [BGG⁺14]. To fill the gap, we flip the output bit.

3. Compute $[\mathbf{A} \parallel \mathbf{B}_F]_{\tau}^{-1}$ from $\mathbf{A}_{\tau_0}^{-1}$ and sample $\mathbf{r} \in \mathbb{Z}^{2m}$ as $\mathbf{r}^{\top} \leftarrow [\mathbf{A} \parallel \mathbf{B}_F]_{\tau}^{-1}(\mathbf{u}^{\top})$.
4. Output the secret key $\text{sk}_C := \mathbf{r}$.

$\text{Enc}(\text{mpk}, \mathbf{x}, \mu)$: The encryption algorithm takes as input the master public key mpk , an attribute $\mathbf{x} \in \{0, 1\}^{\ell}$, and a message $\mu \in \{0, 1\}$ and proceeds as follows.

1. Sample $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $e_0 \leftarrow \chi$, $\mathbf{e} \leftarrow \chi^m$, and $\mathbf{e}_{i,b} \leftarrow \widetilde{\chi}^m$ for $i \in [\ell]$ and $b \in \{0, 1\}$, where $\widetilde{\chi}^m$ is the distribution obtained by first sampling $\mathbf{x} \leftarrow \chi^m$ and $\mathbf{S} \leftarrow \{-1, 1\}^{m \times m}$ and then outputting $\mathbf{S}\mathbf{x}$.
2. Compute

$$\begin{aligned} \text{For all } i \in [\ell], b \in \{0, 1\}, \psi_{i,b} &:= \mathbf{s}(\mathbf{B}_i - b\mathbf{G}) + \mathbf{e}_{i,b} \in \mathbb{Z}_q^m \\ \psi_{2\ell+1} &:= \mathbf{s}\mathbf{A} + \mathbf{e} \in \mathbb{Z}_q^m, \psi_{2\ell+2} := \mathbf{s}\mathbf{u}^{\top} + e_0 + \mu \lceil q/2 \rceil \in \mathbb{Z}_q, \end{aligned}$$

3. Output the ciphertext $\text{ct}_{\mathbf{x}} := (\{\psi_{i,x_i}\}_{i \in [\ell]}, \psi_{2\ell+1}, \psi_{2\ell+2})$, where x_i is the i -th bit of \mathbf{x} .

$\text{Dec}(\text{mpk}, \text{sk}_C, C, \text{ct}_{\mathbf{x}}, \mathbf{x})$: The decryption algorithm takes as input the master public key mpk , a secret key sk_C for a circuit C , and a ciphertext $\text{ct}_{\mathbf{x}}$ for an attribute \mathbf{x} and proceeds as follows. The decryption algorithm can be divided into offline and online phase Dec^{off} and Dec^{on} , respectively, as defined below.

$\text{Dec}^{\text{off}}(\text{mpk}, C, \mathbf{x})$:

1. Compute $C(\mathbf{x})$ and $\widehat{\mathbf{H}}_{F,\mathbf{x}} = \text{EvalF}(F, \mathbf{x}, \mathbf{B})$, where $F = -C$.
2. Output $\text{off} = (C(\mathbf{x}), \widehat{\mathbf{H}}_{F,\mathbf{x}})$.

$\text{Dec}^{\text{on}}(\text{sk}_C, \text{ct}_{\mathbf{x}}, \text{off})$:

1. Parse off as $(C(\mathbf{x}), \widehat{\mathbf{H}}_{F,\mathbf{x}})$ and $\text{ct}_{\mathbf{x}} \rightarrow (\{\psi_{i,x_i} \in \mathbb{Z}_q^m\}_{i \in [\ell]}, \psi_{2\ell+1} \in \mathbb{Z}_q^m, \psi_{2\ell+2} \in \mathbb{Z}_q)$, and $\text{sk}_C = \mathbf{r} \in \mathbb{Z}^{2m}$. If any of the component is not in the corresponding domain or $C(\mathbf{x}) = 0$, output \perp .
2. Concatenate $\{\psi_{i,x_i}\}_{i \in [\ell]}$ to form $\psi_{\mathbf{x}} = (\psi_{1,x_1}, \dots, \psi_{\ell,x_{\ell}})$.
3. Compute

$$\psi' := \psi_{2\ell+2} - [\psi_{2\ell+1} \parallel \psi_{\mathbf{x}} \widehat{\mathbf{H}}_{F,\mathbf{x}}] \mathbf{r}^{\top}.$$

4. Output 0 if $\psi' \in [-B, B]$ and 1 if $[-B + \lceil q/2 \rceil, B + \lceil q/2 \rceil]$.

Parameters and Security. We choose the parameters for the scheme as follows for concreteness:

$$\begin{aligned} d &= \log n \log \log n, & n &= \tilde{\Theta}(\lambda^c), & m &= n^{1.1} \log q, & q &= n^{O(\log^3 n)}, \\ \chi &= \text{SampZ}(3\sqrt{n}), & \tau_0 &= n \log q \log m, & \tau &= n^{O(\log^2 n)}, & B &= n^{O(\log^2 n)}, \end{aligned}$$

where c is some constant (e.g., $c = 1$).

It is easy to see that the efficiency requirement stated in Theorem 2.17 directly follows from the above parameter settings and Lemma 2.20.

Theorem 2.21 (Adapted from [BGG⁺14]). *The above scheme satisfies Sel-INDr security (Definition 2.15) if we assume the LWE assumption with $n^{O(\log^3 n)}$ approximation factor.*

2.4 Lockable Obfuscation

We define lockable obfuscation [GKW17, WZ17] below. Consider a function family $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$, with input space $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda \in \mathbb{N}}$ and output space $\mathcal{Y} = \{\mathcal{Y}_\lambda\}_{\lambda \in \mathbb{N}}$, i.e, $\mathcal{F}_\lambda = \{f : \mathcal{X}_\lambda \rightarrow \mathcal{Y}_\lambda\}$. A lockable obfuscation scheme for \mathcal{F} consists of algorithms Obf and Eval with the following syntax:

$\text{Obf}(1^\lambda, f, \alpha) \rightarrow \tilde{f}$. The obfuscation algorithm takes as input the security parameter λ , a function $f \in \mathcal{F}_\lambda$ and a lock value $\alpha \in \mathcal{Y}_\lambda$. It outputs an obfuscated program \tilde{f} .

$\text{Eval}(\tilde{f}, x) \rightarrow 1 \cup \{\perp\}$. The evaluation algorithm takes as input the obfuscated program \tilde{f} and an input $x \in \mathcal{X}_\lambda$. It outputs 1 or \perp .

Definition 2.22 (Correctness). A lockable obfuscation scheme is said to be correct if it satisfies the following properties:

1. For all $\lambda \in \mathbb{N}$, $f \in \mathcal{F}_\lambda$, $x \in \mathcal{X}_\lambda$ and $\alpha \in \mathcal{Y}_\lambda$ such that $f(x) = \alpha$, we have

$$\text{Eval}(\text{Obf}(1^\lambda, f, \alpha), x) = 1$$

2. There exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $f \in \mathcal{F}_\lambda$, $x \in \mathcal{X}_\lambda$ and $\alpha \in \mathcal{Y}_\lambda$ such that $f(x) \neq \alpha$, we have

$$\Pr[\text{Eval}(\text{Obf}(1^\lambda, f, \alpha), x) = \perp] \geq 1 - \text{negl}(\lambda)$$

where the probability is taken over the random coins used during obfuscation.

Definition 2.23 (Security). A lockable obfuscation scheme is said to be secure if there is a PPT simulator Sim such that for all $f \in \mathcal{F}_\lambda$, we have

$$\text{Obf}(1^\lambda, f, \alpha) \approx_c \text{Sim}(1^\lambda, 1^{|f|})$$

where $\alpha \leftarrow \mathcal{Y}_\lambda$ and the probability is taken over the randomness of the obfuscator and simulator Sim.

Theorem 2.24 ([GKW17, WZ17]). *There exists lockable obfuscation for all circuits with lock space $\{0, 1\}^\lambda$ from the LWE assumption.*

3 Revocable Predicate Encryption

In this section we define revocable predicate encryption (RPE), in both public and secret key setting. Since the two notions differ only in the encryption algorithm, we present them here in a unified way.

Definition 3.1. A RPE scheme for an attribute space $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda \in [\mathbb{N}]}$, a function family $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in [\mathbb{N}]}$ where $\mathcal{F}_\lambda = \{f : \mathcal{X}_\lambda \rightarrow \{0, 1\}\}$, a label space $\mathcal{L} = \{\mathcal{L}_\lambda\}_{\lambda \in [\mathbb{N}]}$ and a message space $\mathcal{M} = \{\mathcal{M}_\lambda\}_{\lambda \in [\mathbb{N}]}$ has the following probabilistic polynomial time algorithms:

$\text{Setup}(1^\lambda) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm takes the security parameter λ as input and it outputs a master public key mpk and a master secret key msk.

$\text{KeyGen}(\text{msk}, \text{lb}, x) \rightarrow \text{sk}_{\text{lb},x}$ ⁹. The key generation algorithm takes as input the master secret key msk , a label $\text{lb} \in \mathcal{L}_\lambda$ and an attribute $x \in \mathcal{X}_\lambda$. It outputs a secret key $\text{sk}_{\text{lb},x}$.

$\text{Enc}(\text{ek}, f, m, L) \rightarrow \text{ct}$. The encryption algorithm takes as input the encryption key ek , a function f , a message $m \in \mathcal{M}_\lambda$, and a revocation list $L \subseteq \mathcal{L}_\lambda$. It outputs a ciphertext ct .

$\text{Dec}(\text{sk}_{\text{lb},x}, \text{ct}, L) \rightarrow m'$. The decryption algorithm takes the secret key $\text{sk}_{\text{lb},x}$, a ciphertext ct , and a revocation list L and it outputs $m' \in \mathcal{M}_\lambda \cup \{\perp\}$.

In *public-key RPE*, we take $\text{ek} = \text{mpk}$ in the Enc algorithm, and in *secret-key RPE*, we take $\text{ek} = \text{msk}$.

Furthermore, there is an additional algorithm in the secret key setting defined as follows:

$\text{Broadcast}(\text{mpk}, m, L) \rightarrow \text{ct}$. On input the master public key, a message m , and a revocation list $L \subseteq \mathcal{L}_\lambda$, the broadcast algorithm outputs a ciphertext ct .

Definition 3.2 (Correctness). A revocable predicate encryption scheme is said to be correct if there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, label $\text{lb} \in \mathcal{L}_\lambda$, attributes $x \in \mathcal{X}_\lambda$, predicates $f \in \mathcal{F}_\lambda$ such that $f(x) = 1$, all messages $m \in \mathcal{M}_\lambda$ and any set of revoked users $L \subseteq \mathcal{L}_\lambda$ such that $\text{lb} \notin L$, if we set $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$ and $\text{sk}_{\text{lb},x} \leftarrow \text{KeyGen}(\text{msk}, \text{lb}, x)$, then the following holds

$$\Pr [\text{Dec}(\text{sk}_{\text{lb},x}, \text{ct}, L) = m] \geq 1 - \text{negl}(\lambda),$$

for $\text{ct} \leftarrow \text{Enc}(\text{ek}, f, m, L)$ (*Encryption correctness*) and $\text{ct} \leftarrow \text{Broadcast}(\text{mpk}, m, L)$ (*Broadcast correctness*).

Security. We now define the security properties of a RPE scheme.

In the following security definitions, we assume for simplicity that the adversary does not make key queries for same input (lb, x) more than once.

Definition 3.3 (q -query Message Hiding). Let $q(\cdot)$ be any fixed polynomial. A RPE scheme satisfies q -query message hiding property if for every PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for every $\lambda \in \mathbb{N}$, all messages $m \in \mathcal{M}_\lambda$ and any subset of revoked users $L \subseteq \mathcal{L}_\lambda$, the following holds

$$\Pr \left[\begin{array}{l} \beta' = \beta : \\ (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda); \\ (f, m_0, m_1, L) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot), \text{Enc}(\text{ek}, \cdot, \cdot, \cdot)}(\text{mpk}); \\ \beta \leftarrow \{0, 1\}; \text{ct}_\beta \leftarrow \text{Enc}(\text{ek}, f, m_\beta, L); \\ \beta' \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot), \text{Enc}(\text{ek}, \cdot, \cdot, \cdot)}(\text{ct}_\beta) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda)$$

where \mathcal{A} can make at most $q(\lambda)$ queries to the encryption oracle $\text{Enc}(\text{ek}, \cdot, \cdot, \cdot)$, and \mathcal{A} is admissible if and only if for all the key queries (lb, x) to the $\text{KeyGen}(\text{msk}, \cdot, \cdot)$ oracle, either $f(x) = 0$ or $\text{lb} \in L$.

Definition 3.4 (q -query Selective Message Hiding). This is the same as the Def 3.3 except that \mathcal{A} outputs the revocation list L in the beginning of the game, before the Setup algorithm is run.

Definition 3.5 (q -query Function Hiding). Let $q(\cdot)$ be any fixed polynomial. A RPE scheme satisfies q -query function hiding property if for every PPT adversary \mathcal{A} , there exists a negligible

⁹We want to point out that the secret key $\text{sk}_{\text{lb},x}$ does not hide the corresponding label lb and attribute x and we assume these to be included in the secret key.

function $\text{negl}(\cdot)$ such that for every $\lambda \in \mathbb{N}$, all messages $m \in \mathcal{M}_\lambda$ and any subset of revoked users $L \subseteq \mathcal{L}_\lambda$, the following holds

$$\Pr \left[\beta' = \beta : \begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda); \\ (f_0, f_1, m, L) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot), \text{Enc}(\text{ek}, \cdot, \cdot, \cdot)}(\text{mpk}); \\ \beta \leftarrow \{0, 1\}; \text{ct}_\beta \leftarrow \text{Enc}(\text{ek}, f_\beta, m, L); \\ \beta' \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot), \text{Enc}(\text{ek}, \cdot, \cdot, \cdot)}(\text{ct}_\beta) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda)$$

where \mathcal{A} can make at most $q(\lambda)$ queries to the encryption oracle $\text{Enc}(\text{ek}, \cdot, \cdot, \cdot)$, and \mathcal{A} is admissible if and only if for all the key queries (lb, x) to the $\text{KeyGen}(\text{msk}, \cdot, \cdot)$ oracle, either $f_0(x) = f_1(x)$ or $\text{lb} \in L$.

Definition 3.6 (q -query Selective Function Hiding). This is the same as the Def 3.5 except that \mathcal{A} outputs the revocation list L in the beginning of the game, before the Setup algorithm is run.

The following security notion is defined only for secret-key RPE scheme.

Definition 3.7 (q -query Selective Broadcast Security). Let $q(\cdot)$ be any fixed polynomial. A RPE scheme satisfies q -query selective broadcast security if there exists a negligible function $\text{negl}(\cdot)$ such that for every PPT adversary \mathcal{A} , for every $\lambda \in \mathbb{N}$, all messages $m \in \mathcal{M}_\lambda$ and any subset of revoked users $L \subseteq \mathcal{L}_\lambda$, the following holds

$$\Pr \left[\beta' = \beta : \begin{array}{l} L \leftarrow \mathcal{A}(1^\lambda); \\ (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda); \\ f, m \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot), \text{Enc}(\text{msk}, \cdot, \cdot, \cdot)}(\text{mpk}); \\ \beta \leftarrow \{0, 1\}; \text{ct}_0 \leftarrow \text{Enc}(\text{msk}, f, m, L); \\ \text{ct}_1 \leftarrow \text{Broadcast}(\text{mpk}, m, L); \\ \beta' \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot), \text{Enc}(\text{msk}, \cdot, \cdot, \cdot)}(\text{ct}_\beta) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda)$$

where \mathcal{A} can make at most $q(\lambda)$ queries to the encryption $\text{Enc}(\text{msk}, \cdot, \cdot, \cdot)$ oracle and \mathcal{A} is admissible if and only if $f(x) = 1, \forall x \in \mathcal{X}_\lambda$.

Remark 3.8. In the *public-key* RPE scheme, the adversary \mathcal{A} can itself simulate the encryption oracle $\text{Enc}(\text{ek}, \cdot, \cdot, \cdot)$, as $\text{ek} = \text{mpk}$ in this setting. Therefore, in public-key setting, we refer to the security definitions without imposing the q -query bound on the encryption oracle.

Remark 3.9. We note that when the message space is binary, function space \mathcal{F}_λ is polynomially small and q is a constant, the weaker security definitions where adversary outputs the challenge function f , the challenge message m and the SK-Enc query functions $\{\tilde{f}_i\}_{i \in [q]}$ at the beginning of the game, before the $\text{Setup}(1^\lambda)$ algorithm is run, is equivalent to the definitions where the adversary outputs $f, m, \{\tilde{f}_i\}_{i \in [q]}$ adaptively. First, the functions can be guessed with polynomial loss. Furthermore, if we restrict the message space to be binary, we can guess the challenge message as well. To extend the message space, we can encrypt each bit by parallel systems.

4 Public-key RPE from FE and LWE

In this section we provide our construction of a public key RPE scheme $\text{RPE} = (\text{RPE.Setup}, \text{RPE.KeyGen}, \text{RPE.Enc}, \text{RPE.Dec})$ for an attribute space $\mathcal{X} = \{\mathcal{X}_\lambda\}_\lambda$, a function family $\mathcal{F} = \{\mathcal{F}_\lambda\}_\lambda$ where $\mathcal{F}_\lambda = \{f : \mathcal{X}_\lambda \rightarrow \{0, 1\}\}$, a label space $\mathcal{L} = \{\mathcal{L}_\lambda\}_\lambda$ and a message space $\mathcal{M} = \{\mathcal{M}_\lambda\}_\lambda$ from polynomial hardness assumptions. We assume that $|\mathcal{F}_\lambda|$ and $|\mathcal{M}_\lambda|$ are bounded by some polynomial in λ . The restriction on $|\mathcal{F}_\lambda|$ is sufficient for our purpose and the restriction on $|\mathcal{M}_\lambda|$ can be removed by running the scheme in parallel.

Our construction uses the following building blocks:

1. A Sel-INDr secure key-policy ABE scheme $\text{kpABE} = (\text{kpABE.Setup}, \text{kpABE.Enc}, \text{kpABE.KeyGen}, \text{kpABE.Dec})$ for circuit class $\mathcal{C}_{\ell(\lambda), d(\lambda)}$ with parameter succinctness and key compactness (Theorem 2.17). Here $\ell(\lambda)$ is the input length and is the length of labels in our setting and the depth of the circuit is $d(\lambda) \in \omega(\log \lambda)$ to support unbounded revocation list. The message space of the scheme kpABE is $\mathcal{M} = \{\mathcal{M}_\lambda\}_\lambda$ and $\mathcal{CT}_{\text{kpABE}}$ denotes the ciphertext space. We assume that uniform sampling from $\mathcal{CT}_{\text{kpABE}}$ is efficiently possible without any parameter.
2. A (fully) compact, selectively secure, public-key functional encryption scheme $\text{FE} = (\text{FE.Setup}, \text{FE.Enc}, \text{FE.KeyGen}, \text{FE.Dec})$ that supports polynomial sized circuits. We assume that the message space is sufficiently large so that it can encrypt an ABE master public key, a (description of) function $f \in \mathcal{F}_\lambda$, a PRF key, two secret keys of SKE, and a trit mode $\in \{0, 1, 2\}$.
3. A PRF $F : \{0, 1\}^\lambda \times \mathcal{X} \rightarrow \{0, 1\}^t$ where t is the length of the randomness used in kpABE encryption (Def. 2.1).
4. A symmetric key encryption schemes $\text{SKE} = (\text{SKE.KeyGen}, \text{SKE.Enc}, \text{SKE.Dec})$ with pseudorandom ciphertexts (Def. 2.3). We let $\mathcal{CT}_{\text{SKE}}$ denote the ciphertext space of SKE .¹⁰ We assume that uniform sampling from $\mathcal{CT}_{\text{SKE}}$ is efficiently possible without any parameter.

We describe our construction below.

$\text{RPE.Setup}(1^\lambda) \rightarrow (\text{RPE.mpk}, \text{RPE.msk})$. The setup algorithm does the following:

- Generate $(\text{FE.mpk}, \text{FE.msk}) \leftarrow \text{FE.Setup}(1^\lambda)$.
- Output $\text{RPE.mpk} = \text{FE.mpk}$ and $\text{RPE.msk} = \text{FE.msk}$.

$\text{RPE.KeyGen}(\text{RPE.msk}, \text{lb}, x) \rightarrow \text{RPE.sk}_{\text{lb}, x}$. The key generation algorithm does the following:

- Sample random values $\gamma_1, \gamma_2, \delta \leftarrow \mathcal{CT}_{\text{SKE}}$.
- Construct a circuit $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ which has the label lb , attribute x , γ_1, γ_2 and δ hardwired, as defined in Figure 1.
- Compute $\text{FE.sk}_{\text{lb}, x} \leftarrow \text{FE.KeyGen}(\text{FE.msk}, \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta])$.
- Output $\text{RPE.sk}_{\text{lb}, x} = \text{FE.sk}_{\text{lb}, x}$.

$\text{RPE.Enc}(\text{RPE.mpk}, f, m, L) \rightarrow \text{RPE.ct}$. The encryption algorithm does the following:

- Parse $\text{RPE.mpk} = \text{FE.mpk}$.
- Sample a PRF key $K \leftarrow \{0, 1\}^\lambda$.
- Generate $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$.
- Compute $\text{FE.ct} \leftarrow \text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f, m, K, 0, \perp, \perp))$.
- Construct a circuit C_L , with revocation list L hardwired defined as follows:
On input a label $\text{lb} \in \mathcal{L}_\lambda$,

$$C_L(\text{lb}) = 1 \text{ if and only if } \text{lb} \notin L. \quad (4.1)$$

Compute $\text{kpABE.sk}_L \leftarrow \text{kpABE.KeyGen}(\text{kpABE.msk}, C_L)$.

¹⁰ We note that we use the same ciphertext space for simplicity even though messages with different lengths are going to be encrypted. To have the same ciphertext space, we can, for example, pad short messages to be some fixed length, which is possible when the message length is bounded by some polynomial.

Function Re-Enc[lb, x , $\gamma_1, \gamma_2, \delta$]

Hardwired values: A label lb, an attribute x , and SKE ciphertexts γ_1, γ_2 , and δ .

Inputs: A kpABE master public key kpABE.mpk, a function $f \in \mathcal{F}_\lambda$, a message $m \in \mathcal{M}_\lambda$, a PRF key K , a trapdoor mode $\text{mode} \in \{0, 1, 2\}$ and SKE keys SKE.key₁ and SKE.key₂.

Output : A kpABE ciphertext.

1. Parse the input as (ABE.mpk, f , m , K , mode, SKE.key₁, SKE.key₂).
2. Set $\tilde{m} = \begin{cases} m & \text{if } f(x) = 1 \\ 0 & \text{if } f(x) = 0. \end{cases}$
3. Compute kpABE.ct_{lb} = kpABE.Enc(kpABE.mpk, lb, \tilde{m} ; $F(K, (\text{lb}, x))$).
4. Compute flag = SKE.Dec(SKE.key₂, δ).
5. Compute out _{i} = SKE.Dec(SKE.key _{i} , γ_i) for $i \in \{1, 2\}$.
6. If mode = 0, output kpABE.ct_{lb}.
7. If mode = 1, output out₁.
8. If mode = 2, output $\begin{cases} \text{out}_2 & \text{if flag} = 1 \\ \text{kpABE.ct}_{\text{lb}} & \text{if flag} = 0. \end{cases}$

Figure 1: Function to compute kpABE ciphertexts depending on various conditions.

- Output RPE.ct = (kpABE.mpk, kpABE.sk _{L} , FE.ct).

RPE.Dec(RPE.sk_{lb, x} , RPE.ct, L) $\rightarrow m'$. The decryption algorithm does the following:

- Parse RPE.ct = (kpABE.mpk, kpABE.sk _{L} , FE.ct) and RPE.sk_{lb, x} = FE.sk_{lb, x} .
- Compute ct' = FE.Dec(FE.sk_{lb, x} , FE.ct).
- Construct circuit C_L from L and compute $m' = \text{kpABE.Dec}(\text{kpABE.mpk}, \text{kpABE.sk}_L, C_L, \text{ct}', \text{lb})$.
- Output m' .

Correctness. We now show that the above construction is correct via the following theorem.

Theorem 4.1. *Suppose FE and kpABE schemes are correct. Then the above construction satisfies the encryption correctness (Def. 3.2).*

Proof. Firstly, for any label lb, attribute x , and function f such that $f(x) = 1$, we have FE.Dec(FE.sk_{lb, x} , FE.ct) = kpABE.ct_{lb}, where kpABE.ct_{lb} = kpABE.Enc(kpABE.mpk, lb, m ; $F(K, (\text{lb}, x))$), by the correctness of FE and the definition of Re-Enc. We then observe that $C_L(\cdot)$ can be implemented with depth $O(\log(|L| \cdot |\text{lb}|)) = O(\log \text{poly}(\lambda)) = O(\log \lambda) \leq d$ and thus $C_L(\cdot) \in \mathcal{C}_{\ell, d}$. Then if $\text{lb} \notin L$ we have $C_L(\text{lb}) = 1$ and hence from the correctness of the kpABE scheme it follows that $\text{kpABE.Dec}(\text{kpABE.mpk}, \text{kpABE.sk}_L, C_L, \text{kpABE.ct}_{\text{lb}}, \text{lb}) = m$. So the decryption correctly recovers the message when $f(x) = 1$ and $\text{lb} \notin L$. \square

Efficiency. Here we argue that our construction achieves optimal parameters. Namely, we show that the size of each parameter is independent from $|L|$. We note that $|f|$ refers to the description size of the function, not the circuit size that implements the function. When $|x|$ is very long and f has succinct description, the former can be much shorter than the latter.

1. Public key size $|RPE.mpk|$: We have $|RPE.mpk| = |FE.mpk|$. Since we assumed that FE is fully compact (Def. 2.7), the length of FE.mpk only depends on the input length of Re-Enc. We have that the input length is $|ABE.mpk| + |f| + |m| + |K| + |\text{mode}| + 2|SKE.key| = |ABE.mpk| + |f| + O(\lambda)$. We have $|ABE.mpk| \leq \text{poly}(\lambda, |\text{lb}|, d) = \text{poly}(\lambda, |\text{lb}|)$ by Theorem 2.17. The total length is therefore $\text{poly}(\lambda, |f|, |\text{lb}|)$.
2. Secret key size $|RPE.sk_{|\text{lb},x}|$: We have $|RPE.sk_{|\text{lb},x}| = |FE.sk_{|\text{lb},x}|$. Since the size of the latter is polynomially dependent on the size of Re-Enc, we evaluate its size. We can see that the size of Re-Enc is polynomial in the total length of the input and the hardwired values. The length of the input is bounded by $\text{poly}(\lambda, |f|, |\text{lb}|)$ as analyzed in the above item. The length of the hardwired values are $|\text{lb}| + |x| + |\gamma_1| + |\gamma_2| + |\delta|$. We have $|\gamma_1| + |\gamma_2| + |\delta| = 3|\gamma_2|$ ¹¹ and $|\gamma_2| = \text{poly}(\lambda, \text{kpABE.ct}_{|\text{lb}}) = \text{poly}(\lambda, \text{lb}, d) = \text{poly}(\lambda, \text{lb})$. Therefore, the size of Re-Enc is $\text{poly}(\lambda, |x|, |f|, |\text{lb}|)$ and so is the size of the secret key.
3. Ciphertext size $|RPE.ct|$: We have $|RPE.ct| = |\text{kpABE.mpk}| + |\text{kpABE.sk}_L| + |FE.ct|$. We have $|ABE.mpk| = \text{poly}(\lambda, |\text{lb}|)$ as we showed in the first item. We also have $|\text{kpABE.sk}_L| \leq \text{poly}(\lambda, d) \leq \text{poly}(\lambda)$ by the key compactness of kpABE (2.17). By similar analysis to the first item, full compactness of FE implies $|FE.ct| \leq \text{poly}(\lambda, |f|, |\text{lb}|)$. Therefore, the overall length of the ciphertext is $\text{poly}(\lambda, |f|, |\text{lb}|)$.

Security. Now we prove that the above construction of RPE satisfies both function hiding and message hiding security.

Function Hiding

Theorem 4.2. *Assume that F is a secure PRF, SKE is correct and secure, FE and kpABE are secure as per definitions 2.6 and 2.15, respectively. Furthermore, assume $|\mathcal{F}_\lambda| \leq \text{poly}(\lambda)$ and $|\mathcal{M}_\lambda| \leq \text{poly}(\lambda)$. Then the RPE constructed above is function hiding (Def. 3.5).*

Proof. Recall that for function hiding we need $RPE.\text{Enc}(RPE.mpk, f_0, m, L) \approx_c RPE.\text{Enc}(RPE.mpk, f_1, m, L)$, where for all the key queries (lb, x) , either $f_0(x) = f_1(x)$ or $\text{lb} \in L$.

The proof proceeds via a sequence of hybrid games between the challenger and a PPT adversary \mathcal{A} .

Hybrid₀. This is the real world with $\beta = 0$, i.e. the challenge ciphertext is computed using the function f_0 . We write the complete game here to set up the notations and easy reference in later hybrids.

1. The adversary outputs the challenge functions f_0 and f_1 and the challenge message m ¹².

¹¹We assumed that $|\gamma_1| = |\gamma_2| = |\delta|$. See Footnote 10.

¹²To keep the proofs and notations simple, we let f_0, f_1, m to be given selectively. This is sufficient to achieve security as in Def 3.5, as mentioned in Remark 3.9.

2. The challenger generates $(\text{FE.mpk}, \text{FE.msk}) \leftarrow \text{FE.Setup}(1^\lambda)$, sets $\text{RPE.mpk} = \text{FE.mpk}$ and sends it to the adversary. The challenger then responds to different queries from \mathcal{A} as follows:
3. **Key Queries :** For each key query (lb, x) , the challenger does the following:
 - Samples random values $\gamma_1, \gamma_2, \delta \leftarrow \mathcal{CT}_{\text{SKE}}$.
 - Defines the circuit $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ as in Figure 1 and computes $\text{FE.sk}_{\text{lb},x} \leftarrow \text{FE.KeyGen}(\text{FE.msk}, \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta])$.
 - Returns $\text{RPE.sk}_{\text{lb},x} = \text{FE.sk}_{\text{lb},x}$ to the adversary.
4. **Challenge Query:** When the adversary outputs the revocation list L for the challenge query, the challenger does the following:
 - Samples a PRF key K .
 - Generates $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$.
 - Computes FE.ct as
$$\text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f_0, m, K, 0, \perp, \perp)).$$
 - Defines C_L as in Eq. (4.1) and computes $\text{kpABE.sk}_L \leftarrow \text{kpABE.KeyGen}(\text{kpABE.msk}, C_L)$.
 - Returns $\text{RPE.ct} = (\text{kpABE.mpk}, \text{kpABE.sk}_L, \text{FE.ct})$ to the adversary.
5. In the end, the adversary outputs a bit β' .

Hybrid₁. This hybrid is same as the previous hybrid except the following changes:

- The challenger generates $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$ in the beginning of the game after the adversary outputs the challenge functions f_0 and f_1 and the challenge message m . It also samples a SKE secret key SKE.key_1 and a PRF key K .
- For each key query (lb, x) , γ_1 is computed differently. In particular, the challenger does the following
 - Sets $\tilde{m} = \begin{cases} m & \text{if } f_0(x) = 1 \\ 0 & \text{if } f_0(x) = 0 \end{cases}$ and computes $\text{kpABE.ct}'_{\text{lb}} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; F(K, (\text{lb}, x)))$.
 - Sets γ_1 as $\text{SKE.Enc}(\text{SKE.key}_1, \text{kpABE.ct}'_{\text{lb}})$.

Hybrid₂. This hybrid is same as the previous hybrid except the following changes:

- The challenger samples two SKE secret keys SKE.key_1 and SKE.key_2 in the beginning of the game (after receiving f_0, f_1, m from \mathcal{A}).
- For each key query (lb, x) , γ_2 and δ are computed differently from the previous hybrid. In particular, the challenger does the following:
 - Sets $\text{flag} = \begin{cases} 1 & \text{if } f_0(x) \neq f_1(x) \\ 0 & \text{otherwise.} \end{cases}$
 - Sets $\delta \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{flag})$ and $\gamma_2 \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct}'_{\text{lb}})$ if $\text{flag} = 1$; else $\gamma_2 \leftarrow \mathcal{CT}_{\text{SKE}}$. We note that $\text{kpABE.ct}'_{\text{lb}}$ and γ_1 are computed as in the previous hybrid.

Hybrid₃. This hybrid is same as the previous hybrid except that FE.ct in the challenge ciphertext is computed as

$$\text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f_0, m, \perp, 1, \text{SKE.key}_1, \text{SKE.key}_2)).$$

Hybrid₄. This hybrid is same as the previous hybrid except that for each key query (lb, x) , $kpABE.ct'_{lb}$ is computed as follows:

- If $flag = 1$, $kpABE.ct'_{lb} = kpABE.Enc(kpABE.mpk, lb, \tilde{m}; r)$, where $r \leftarrow \{0, 1\}^t$.
- Else, if $flag = 0$, $kpABE.ct'_{lb} = kpABE.Enc(kpABE.mpk, lb, \tilde{m}; F(K, (lb, x)))$.

Hybrid₅. This hybrid is same as the previous hybrid except that $FE.ct$ in the challenge ciphertext is computed as

$$FE.Enc(FE.mpk, (kpABE.mpk, f_0, m, K, 2, \perp, SKE.key_2)).$$

Hybrid₆. This hybrid is same as the previous hybrid except that for each key query (lb, x) , γ_1 is set differently as $\gamma_1 \leftarrow \mathcal{CT}_{SKE}$.

Hybrid₇. This hybrid is same as the previous hybrid except that for each such key query (lb, x) where $flag = 1$, $kpABE.ct'_{lb}$ is sampled uniformly from \mathcal{CT}_{kpABE} , i.e., $kpABE.ct'_{lb} \leftarrow \mathcal{CT}_{kpABE}$.

Hybrid₈. This hybrid is same as the previous hybrid except that $FE.ct$ in the challenge ciphertext is computed as

$$FE.Enc(FE.mpk, (kpABE.mpk, f_1, m, K, 2, \perp, SKE.key_2)).$$

We note that in the hybrids hereafter, we rewind the changes made in the preceding hybrids.

Hybrid₉. This hybrid is same as the previous hybrid except that for each such key query (lb, x) where $flag = 1$, $kpABE.ct'_{lb}$ is changed back to $kpABE.Enc(kpABE.mpk, lb, \tilde{m}; r)$ for $r \leftarrow \{0, 1\}^t$, where \tilde{m} is now defined as m if $f_1(x) = 1$ and 0 otherwise.

Hybrid₁₀. This hybrid is same as the previous hybrid except that for all the key queries (lb, x) , the challenger sets γ_1 as $SKE.Enc(SKE.key_1, kpABE.ct'_{lb})$.

Hybrid₁₁. This hybrid is same as the previous hybrid except that $FE.ct$ in the challenge ciphertext is computed as

$$FE.Enc(FE.mpk, (kpABE.mpk, f_1, m, \perp, 1, SKE.key_1, SKE.key_2)).$$

Hybrid₁₂. This hybrid is same as the previous hybrid except that for each key query (lb, x) , $kpABE.ct'_{lb}$ is computed as $kpABE.Enc(kpABE.mpk, lb, \tilde{m}; F(K, (lb, x)))$.

Hybrid₁₃. This hybrid is same as the previous hybrid except that $FE.ct$ in the challenge ciphertext is computed as

$$FE.Enc(FE.mpk, (kpABE.mpk, f_1, m, K, 0, \perp, \perp)).$$

Hybrid₁₄. This hybrid is same as the previous hybrid except that γ_2 and δ_2 are set as $\gamma_2, \delta \leftarrow \mathcal{CT}_{SKE}$.

Hybrid₁₅. This hybrid is same as the previous hybrid except that γ_1 is set as $\gamma_1 \leftarrow \mathcal{CT}_{SKE}$. Note that this is the real world with $\beta = 1$, i.e., f_1 is encrypted in the challenge ciphertext.

Indistinguishability of hybrids We now show that the consecutive hybrids are indistinguishable.

Claim 4.2.1. *Assume that SKE is secure, then $\text{Hybrid}_0 \approx_c \text{Hybrid}_1$.*

Proof. We show that if \mathcal{A} can distinguish between Hybrid_0 and Hybrid_1 with non-negligible advantage ϵ , then there exists a PPT adversary \mathcal{B} against the security of SKE scheme with advantage ϵ . The reduction is as follows.

1. The SKE challenger samples $\text{SKE.key}_1 \leftarrow \text{SKE.Setup}(1^\lambda)$ and a bit $\hat{\beta} \leftarrow \{0, 1\}$ and starts the SKE security game with \mathcal{B} .
2. \mathcal{B} invokes \mathcal{A} , which then outputs the challenge functions f_0 and f_1 and the challenge message m .
3. \mathcal{B} generates $(\text{FE.mpk}, \text{FE.msk}) \leftarrow \text{FE.Setup}(1^\lambda)$, sets $\text{RPE.mpk} = \text{FE.mpk}$ and sends RPE.mpk to \mathcal{A} .
 \mathcal{B} also generates $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$ and samples a PRF key K .
4. **Key Queries :** Whenever \mathcal{A} issues a key query (lb, x) , \mathcal{B} does the following:
 - It sets $\tilde{m} = \begin{cases} m & \text{if } f_0(x) = 1 \\ 0 & \text{if } f_0(x) = 0 \end{cases}$
and computes $\text{kpABE.ct}'_{\text{lb}} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; F(K, (\text{lb}, x)))$.
 - It sends $\text{kpABE.ct}'_{\text{lb}}$ as the challenge message to the SKE challenger. The SKE challenger returns $\text{ct}_{\hat{\beta}}$ to \mathcal{B} , where $\text{ct}_0 \leftarrow \mathcal{CT}_{\text{SKE}}$ and $\text{ct}_1 \leftarrow \text{SKE.Enc}(\text{SKE.key}_1, \text{kpABE.ct}'_{\text{lb}})$.
 - It sets $\gamma_1 = \text{ct}_{\hat{\beta}}$, samples random values $\gamma_2, \delta \leftarrow \mathcal{CT}_{\text{SKE}}$ and computes the FE key for the circuit $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ and returns it as the secret key $\text{RPE.sk}_{\text{lb}, x}$ to \mathcal{A} .
5. **Challenge Query :** When \mathcal{A} outputs the revocation list L for the challenge ciphertext, \mathcal{B} does the following:
 - Computes $\text{FE.ct} \leftarrow \text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f_0, m, K, 0, \perp, \perp))$.
 - Defines C_L as in Eq. 4.1 and computes $\text{kpABE.sk}_L \leftarrow \text{kpABE.KeyGen}(\text{kpABE.msk}, C_L)$.
 - Returns $\text{RPE.ct} = (\text{kpABE.mpk}, \text{kpABE.sk}_L, \text{FE.ct})$ to \mathcal{A} .
6. In the end, \mathcal{A} outputs a bit β' . \mathcal{B} sends β' to the SKE challenger.

We observe that if the SKE challenger samples $\hat{\beta} = 0$, then \mathcal{B} simulated Hybrid_0 , else Hybrid_1 with \mathcal{A} . Hence, advantage of $\mathcal{B} = |\Pr(\beta' = 1 | \hat{\beta} = 0) - \Pr(\beta' = 1 | \hat{\beta} = 1)| = |\Pr(\beta' = 1 | \text{Hybrid}_0) - \Pr(\beta' = 1 | \text{Hybrid}_1)| = \epsilon$ (by assumption). □

Claim 4.2.2. *Assume that SKE is secure. Then $\text{Hybrid}_1 \approx_c \text{Hybrid}_2$.*

Proof. The proof follows the same steps as that for the claim 4.2.1 and hence omitted. □

Claim 4.2.3. *Assume that FE satisfies selective security (Def. 2.6) and SKE is correct. Then $\text{Hybrid}_2 \approx_c$ and Hybrid_3 .*

Proof. We show that if \mathcal{A} can distinguish between the two hybrids with non-negligible advantage ϵ , then there exists a PPT adversary \mathcal{B} against the security of FE with the same advantage ϵ . \mathcal{B} is defined as follows:

1. Upon being invoked by the FE challenger, \mathcal{B} invokes \mathcal{A} . \mathcal{A} outputs the challenge functions f_0 and f_1 and the challenge message m .
2. \mathcal{B} samples two SKE secret keys $\text{SKE.key}_1, \text{SKE.key}_2$, a PRF key K and generates $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$.
3. It sets $\mu_0 = (\text{kpABE.mpk}, f_0, m, K, 0, \perp, \perp)$ and $\mu_1 = (\text{kpABE.mpk}, f_0, m, \perp, 1, \text{SKE.key}_1, \text{SKE.key}_2)$ and sends (μ_0, μ_1) to the FE challenger as challenge messages.
4. The FE challenger generates $(\text{FE.mpk}, \text{FE.msk}) \leftarrow \text{FE.Setup}(1^\lambda)$, samples $\hat{\beta} \leftarrow \{0, 1\}$. It then computes $\text{FE.ct}_{\hat{\beta}} \leftarrow \text{FE.Enc}(\text{FE.mpk}, \mu_{\hat{\beta}})$ and sends $(\text{FE.mpk}, \text{FE.ct}_{\hat{\beta}})$ to \mathcal{B} .
5. \mathcal{B} sets $\text{RPE.mpk} = \text{FE.mpk}$ and sends RPE.mpk to \mathcal{A} .
6. **Key Queries :** When \mathcal{A} issues a key query (lb, x) , \mathcal{B} does the following:

- It sets $\tilde{m} = \begin{cases} m & \text{if } f_0(x) = 1 \\ 0 & \text{if } f_0(x) = 0 \end{cases}$ and computes $\text{kpABE.ct}'_{\text{lb}} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; F(K, (\text{lb}, x)))$.
- It sets $\text{flag} = \begin{cases} 1 & \text{if } f_0(x) \neq f_1(x) \\ 0 & \text{otherwise} \end{cases}$.
- It computes $\gamma_1 \leftarrow \text{SKE.Enc}(\text{SKE.key}_1, \text{kpABE.ct}'_{\text{lb}})$, $\delta \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{flag})$ and $\gamma_2 \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct}'_{\text{lb}})$ if $\text{flag} = 1$; else samples $\gamma_2 \leftarrow \mathcal{CT}_{\text{SKE}}$.
- It defines $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ as in Figure 1 and sends a key query $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ to the FE challenger. The FE challenger returns the secret key $\text{FE.sk}_{\text{lb}, x}$ to \mathcal{B} .
- \mathcal{B} returns $\text{RPE.sk}_{\text{lb}, x} = \text{FE.sk}_{\text{lb}, x}$ to \mathcal{A} .

7. **Challenge Query :** When \mathcal{A} outputs the revocation list L for the challenge ciphertext, \mathcal{B} does the following:

- It defines C_L as in Eq. (4.1) and computes $\text{kpABE.sk}_L \leftarrow \text{kpABE.KeyGen}(\text{kpABE.msk}, C_L)$.
- Returns $\text{RPE.ct} = (\text{kpABE.mpk}, \text{kpABE.sk}_L, \text{FE.ct}_{\hat{\beta}})$ to \mathcal{A} .

8. In the end, \mathcal{A} outputs a bit β' . \mathcal{B} sends β' to the FE challenger.

We observe that if FE challenger chose $\hat{\beta} = 0$, then \mathcal{B} simulated Hybrid_2 , else Hybrid_3 with \mathcal{A} . Hence, advantage of $\mathcal{B} = |\Pr(\beta' = 1 | \hat{\beta} = 0) - \Pr(\beta' = 1 | \hat{\beta} = 1)| = |\Pr(\beta' = 1 | \text{Hybrid}_2) - \Pr(\beta' = 1 | \text{Hybrid}_3)| = \epsilon$ (by assumption).

Admissibility of \mathcal{B}

Firstly, we observe that the only key queries that \mathcal{B} issues to the FE challenger are for the Re-Enc functions defined for each key query (lb, x) by \mathcal{A} . Next, we observe that for any function $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$, $\mu_0 = (\text{kpABE.mpk}, f_0, m, K, 0, \perp, \perp)$, and $\mu_1 = (\text{kpABE.mpk}, f_0, m, \perp, 1, \text{SKE.key}_1, \text{SKE.key}_2)$ the following holds true from the definition of Re-Enc and correctness of SKE decryption,

$$\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_0) = \text{kpABE.ct}_{\text{lb}}$$

$$= \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; F(K, (\text{lb}, x)))$$

$$\begin{aligned} \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_1) &= \text{SKE.Dec}(\text{SKE.key}_1, \gamma_1) \\ &= \text{kpABE.ct}'_{\text{lb}} \\ &= \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; F(K, (\text{lb}, x))) \end{aligned}$$

Thus, for all the keys queried to the FE challenger, $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_0) = \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_1)$. This establishes the admissibility of \mathcal{B} . \square

Claim 4.2.4. *Assume that PRF is secure, then $\text{Hybrid}_3 \approx_c \text{Hybrid}_4$.*

Proof. We show that if \mathcal{A} can distinguish between the two hybrids with non-negligible advantage ϵ , then there exists a PPT adversary \mathcal{B} against PRF security with the same advantage ϵ . The reduction \mathcal{B} is defined as follows:

1. The PRF challenger samples a PRF key K and a bit $\hat{\beta} \leftarrow \{0, 1\}$ and starts the game with \mathcal{B} .
2. \mathcal{B} then invokes \mathcal{A} which outputs the challenge functions f_0 and f_1 and the challenge message m .
3. \mathcal{B} samples two SKE secret keys $\text{SKE.key}_1, \text{SKE.key}_2$, generates $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$, $(\text{FE.mpk}, \text{FE.msk}) \leftarrow \text{FE.Setup}(1^\lambda)$, sets $\text{RPE.mpk} = \text{FE.mpk}$ and sends RPE.mpk to \mathcal{A} .
4. **Key Queries :** When \mathcal{A} issues a key query (lb, x) , \mathcal{B} does the following:
 - If $\text{flag} = 0$, it sends an evaluation query for input (lb, x) to the PRF challenger and gets back $F(K, (\text{lb}, x))$. It then computes $\text{kpABE.ct}'_{\text{lb}} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; F(K, (\text{lb}, x)))$.
 - If $\text{flag} = 1$, it sends (lb, x) as a challenge query to the PRF challenger and gets back $r_{\hat{\beta}}$, where $r_0 = F(K, (\text{lb}, x))$ and $r_1 \leftarrow \{0, 1\}^t$. It then computes $\text{kpABE.ct}'_{\text{lb}} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; r_{\hat{\beta}})$.
 - Computes $\gamma_1 \leftarrow \text{SKE.Enc}(\text{SKE.key}_1, \text{kpABE.ct}'_{\text{lb}})$, $\delta \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{flag})$ and $\gamma_2 \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct}'_{\text{lb}})$ if $\text{flag} = 1$; else $\gamma_2 \leftarrow \mathcal{CT}_{\text{SKE}}$.
 - Defines $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ and computes the FE key for the circuit $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ and returns it as the secret key $\text{RPE.sk}_{\text{lb}, x}$ to \mathcal{A} .
5. **Challenge Query :** When \mathcal{A} outputs the revocation list L for the challenge ciphertext, \mathcal{B} does the following
 - Computes $\text{FE.ct} \leftarrow \text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f_0, m, \perp, 1, \text{SKE.key}_1, \text{SKE.key}_2))$.
 - Defines C_L as in Eq. 4.1 and computes $\text{kpABE.sk}_L \leftarrow \text{kpABE.KeyGen}(\text{kpABE.msk}, C_L)$.
 - Returns $\text{RPE.ct} = (\text{kpABE.mpk}, \text{kpABE.sk}_L, \text{FE.ct})$ to \mathcal{A} .
6. In the end, \mathcal{A} outputs a bit β' . \mathcal{B} sends β' to the SKE challenger.

We observe that if $\hat{\beta} = 0$, then \mathcal{B} simulated Hybrid_3 , else Hybrid_4 with \mathcal{A} . Hence, advantage of $\mathcal{B} = |\Pr(\beta' = 1 | \hat{\beta} = 0) - \Pr(\beta' = 1 | \hat{\beta} = 1)| = |\Pr(\beta' = 1 | \text{Hybrid}_3) - \Pr(\beta' = 1 | \text{Hybrid}_4)| = \epsilon$ (by assumption). \square

Claim 4.2.5. Assume that FE is selectively secure (as per Def 2.6) and SKE is correct. Then $\text{Hybrid}_4 \approx_c \text{Hybrid}_5$.

Proof. We show that if \mathcal{A} wins with non-negligible advantage ϵ in distinguishing the two hybrids, then there exists an adversary \mathcal{B} against FE security with the same advantage ϵ . The steps of the reduction are similar as in the proof of the Claim 4.2.3, with $\mu_0 = (\text{kpABE.mpk}, f_0, m, \perp, 1, \text{SKE.key}_1, \text{SKE.key}_2)$ and $\mu_1 = (\text{kpABE.mpk}, f_0, m, K, 2, \perp, \text{SKE.key}_2)$. Hence, here we only argue the admissibility of \mathcal{B} in the FE security game.

Admissibility of \mathcal{B}

Firstly, we observe that the only key queries that \mathcal{B} issues to the FE challenger are for the $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ functions, where each function is defined corresponding to a key query (lb, x) by \mathcal{A} . Here, $\gamma_1 \leftarrow \text{SKE.Enc}(\text{SKE.key}_1, \text{kpABE.ct}'_{\text{lb}})$, $\delta \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{flag})$ and $\gamma_2 \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct}'_{\text{lb}})$ if $\text{flag} = 1$; else $\gamma_2 \leftarrow \mathcal{CT}_{\text{SKE}}$. Also $\text{kpABE.ct}'_{\text{lb}} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; F(K, (\text{lb}, x)))$ if $\text{flag} = 0$; else $\text{kpABE.ct}'_{\text{lb}} \leftarrow \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; r)$ where $r \leftarrow \{0, 1\}^t$.

Next, we observe that for any function $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$, $\mu_0 = (\text{kpABE.mpk}, f_0, m, \perp, 1, \text{SKE.key}_1, \text{SKE.key}_2)$, and $\mu_1 = (\text{kpABE.mpk}, f_0, m, K, 2, \perp, \text{SKE.key}_2)$ the following holds true from the definition of Re-Enc and correctness of SKE decryption,

$$\begin{aligned} \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_0) &= \text{SKE.Dec}(\text{SKE.key}_1, \gamma_1) \\ &= \text{kpABE.ct}'_{\text{lb}} \\ &= \begin{cases} \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; F(K, (\text{lb}, x))) & \text{if flag} = 0 \\ \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; r) & \text{if flag} = 1 \end{cases} \end{aligned}$$

$$\begin{aligned} \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_1) &= \begin{cases} \text{kpABE.ct}_{\text{lb}} & \text{if flag} = 0 \\ \text{SKE.Dec}(\text{SKE.key}_2, \gamma_2) = \text{kpABE.ct}'_{\text{lb}} & \text{if flag} = 1 \end{cases} \\ &= \begin{cases} \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; F(K, (\text{lb}, x))) & \text{if flag} = 0 \\ \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; r) & \text{if flag} = 1 \end{cases} \end{aligned}$$

Thus, $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_0) = \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_1)$. This establishes the admissibility of \mathcal{B} . \square

Claim 4.2.6. Assume that SKE is secure, then $\text{Hybrid}_5 \approx_c \text{Hybrid}_6$.

Proof. The proof is similar to the proof for the claim 4.2.1 and hence omitted. \square

Claim 4.2.7. Assume that kpABE is Sel-INDr secure (Def. 2.15), then $\text{Hybrid}_6 \approx_c \text{Hybrid}_7$.

Proof. To argue indistinguishability between the two hybrids, we define an intermediate hybrid Hybrid_{6a} as follows - this hybrid is same as Hybrid_6 except that for each such pre-challenge key query (lb, x) where $\text{flag} = 1$, $\text{kpABE.ct}'_{\text{lb}} \leftarrow \mathcal{CT}_{\text{kpABE}}$. The proof then follows from the following two claims:

Claim 4.3 (1). Assuming kpABE is Sel-INDr secure (Def. 2.15), $\text{Hybrid}_6 \approx_c \text{Hybrid}_{6a}$.

Proof. Let Q_{fpre} be the number of pre-challenge key queries with $\text{flag} = 1$ ¹³. Then, we further define the following sub hybrids: for $i = 0$ to Q_{fpre} , define $\text{Hybrid}_{6.i}$ which is same as Hybrid_6 ,

¹³Note that we can upper bound Q_{fpre} as $Q_{\text{fpre}} \leq |L|$.

except that for the first i key queries with $\text{flag} = 1$, $\text{kpABE.ct}'_{\text{lb}} \leftarrow \mathcal{CT}_{\text{kpABE}}$. Thus, $\text{Hybrid}_{6.0} = \text{Hybrid}_6$ and $\text{Hybrid}_{6.Q_{\text{fpre}}} = \text{Hybrid}_{6a}$. Next, we show that for all $i \in [Q_{\text{fpre}}]$, $\text{Hybrid}_{6.i-1} \approx_c \text{Hybrid}_{6.i}$. In particular, we show that if \mathcal{A} distinguishes between the two hybrids with non-negligible advantage ϵ , then there exists a PPT algorithm \mathcal{B} against Sel-INDr security of kpABE with the same advantage ϵ .

Observe that the two hybrids differ only in the value of $\text{kpABE.ct}'_{\text{lb}}$ used in the computation of $\text{RPE.sk}_{\text{lb},x}$ for the i -th key query (lb, x) with $\text{flag} = 1$; in the former hybrid we have $\text{kpABE.ct}'_{\text{lb}} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; r)$, while in the latter hybrid, $\text{kpABE.ct}'_{\text{lb}} \leftarrow \mathcal{CT}_{\text{kpABE}}$. Now, we define the reduction \mathcal{B} .

1. \mathcal{B} firstly invokes \mathcal{A} and gets f_0, f_1 and m .
2. \mathcal{B} then samples two SKE secret keys $\text{SKE.key}_1, \text{SKE.key}_2$, a PRF key K and generates $(\text{FE.mpk}, \text{FE.msk}) \leftarrow \text{FE.Setup}(1^\lambda)$. It sets $\text{RPE.mpk} = \text{FE.mpk}$ and sends it to \mathcal{A} . It then answers different queries from \mathcal{A} as follows:
 3. **Key Queries:** For each key query (lb, x) , \mathcal{B} does the following:
 - Sets $\text{flag} = 0$, if $f_0(x) = f_1(x)$; else $\text{flag} = 1$. It computes $\delta = \text{SKE.Enc}(\text{SKE.key}_2, \text{flag})$ and samples $\gamma_1 \leftarrow \mathcal{CT}_{\text{SKE}}$.
 - Computes γ_2 as follows:
 - If $\text{flag} = 0$, $\gamma_2 \leftarrow \mathcal{CT}_{\text{SKE}}$.
 - Else, if $\text{flag} = 1$, then let this be the j -th key query with $\text{flag} = 1$. Then,
 - * For $j < i$, $\gamma_2 \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct}'_{\text{lb}})$, where $\text{kpABE.ct}'_{\text{lb}} \leftarrow \mathcal{CT}_{\text{kpABE}}$. (Note that this does not require kpABE.mpk).
 - * For $j = i$, \mathcal{B} does the following:
 - Sends lb and \tilde{m} as the challenge attribute and message, respectively, to the kpABE challenger.
 - The kpABE challenger samples $\hat{\beta} \leftarrow \{0, 1\}$ and computes $\text{kpABE.ct} \leftarrow \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; r)$, if $\hat{\beta} = 0$, else samples $\text{kpABE.ct} \leftarrow \mathcal{CT}_{\text{kpABE}}$. The kpABE challenger sends $\{\text{kpABE.mpk}, \text{kpABE.ct}\}$ to \mathcal{B} .
 - \mathcal{B} then computes $\gamma_2 \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct})$.
 - * For $j > i$, $\gamma_2 \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct}'_{\text{lb}})$, where $\text{kpABE.ct}'_{\text{lb}} \leftarrow \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; r)$.
 - Defines $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$, computes an FE key for this circuit and returns it as the secret key $\text{RPE.sk}_{\text{lb},x}$ to \mathcal{A} .
 4. **Challenge Query:** When \mathcal{A} outputs the revocation list L for the challenge query, \mathcal{B} does the following:
 - Computes $\text{FE.ct} \leftarrow \text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f_0, m, K, 2, \perp, \text{SKE.key}_2))$.
 - Defines C_L as in Eq. 4.1 and sends a key query for the circuit C_L to the kpABE challenger. The kpABE challenger returns kpABE.sk_L .
 - Returns $\text{RPE.ct} = (\text{kpABE.mpk}, \text{kpABE.sk}_L, \text{FE.ct})$ to \mathcal{A} .
 5. In the end, \mathcal{A} outputs its guess bit β' . \mathcal{B} sends β' to the kpABE challenger as its guess bit.

We observe that if the kpABE challenger chose $\hat{\beta} = 0$, then \mathcal{B} simulated $\text{Hybrid}_{6,i-1}$, else $\text{Hybrid}_{6,i}$ with \mathcal{A} . Hence, advantage of $\mathcal{B} = |\Pr(\beta' = 1 | \hat{\beta} = 0) - \Pr(\beta' = 1 | \hat{\beta} = 1)| =$

$|\Pr(\beta' = 1 | \text{Hybrid}_{6,i-1}) - \Pr(\beta' = 1 | \text{Hybrid}_{6,i})| = \epsilon$ (by assumption).

Admissibility of \mathcal{B} : Observe that \mathcal{B} issues a single kpABE key query which is for circuit C_L and the challenge attribute is the label lb corresponding to a key query (lb, x) by \mathcal{A} for which $\text{flag} = 1$. Hence, by the admissibility condition of \mathcal{A} , $\text{lb} \in L$ and hence $C_L(\text{lb}) = 0$. \square

Claim 4.4 (2). Assume kpABE is Sel-INDr secure (Def. 2.15). Then, $\text{Hybrid}_{6a} \approx_c \text{Hybrid}_7$.

Proof. To prove the claim, we again consider sub hybrids, $\text{Hybrid}_{6a,i}$ for $i = 0$ to $|L|$, defined as follows: let $L_{[1:j]}$ be the set of first j labels in the revocation list L . Then, $\text{Hybrid}_{6a,i}$ is same as Hybrid_{6a} except the following changes: for any post-challenge key query (lb, x) such that $\text{lb} \in L_{[1:i]}$ and $\text{flag} = 1$, $\text{kpABE.ct}'_{\text{lb}} \leftarrow \mathcal{CT}_{\text{kpABE}}$. Thus, $\text{Hybrid}_{6a,0} = \text{Hybrid}_{6a}$ and $\text{Hybrid}_{6a,|L|} = \text{Hybrid}_7$. Next we prove the following claim:

Claim 4.5. Assume kpABE is Sel-INDr secure (Def. 2.15). Then for $i \in [|L|]$, $\text{Hybrid}_{6a,i-1} \approx_c \text{Hybrid}_{6a,i}$.

Proof. Let lb_i be the i -th label in L . Then we observe that if there is no post-challenge key query (lb, x) such that $\text{flag} = 1$ and $\text{lb} = \text{lb}_i$ then the two hybrids are identical. Else, we show that if \mathcal{A} can distinguish between the two hybrids with non negligible advantage ϵ then there exists a PPT algorithm \mathcal{B} against Sel-INDr security of kpABE security with the same advantage ϵ . \mathcal{B} is defined as follows:

1. \mathcal{B} firstly invokes \mathcal{A} and gets f_0, f_1 and m .
2. It then samples a SKE secret key SKE.key_2 , a PRF key K and generates FE keys as $(\text{FE.mpk}, \text{FE.msk}) \leftarrow \text{FE.Setup}(1^\lambda)$. It sets $\text{RPE.mpk} = \text{FE.mpk}$ and sends it to \mathcal{A} . It then answers different queries from \mathcal{A} as follows:
 3. **Pre-challenge Key Queries:** For pre-challenge key query (lb, x) , \mathcal{B} does the following:
 - Sets $\text{flag} = 0$ if $f_0(x) = f_1(x)$; else $\text{flag} = 1$ and computes $\delta = \text{SKE.Enc}(\text{SKE.key}_2, \text{flag})$. It also samples $\gamma_1 \leftarrow \mathcal{CT}_{\text{SKE}}$.
 - Computes γ_2 as follows: if $\text{flag} = 0$, $\gamma_2 \leftarrow \mathcal{CT}_{\text{SKE}}$; else $\gamma_2 \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct}'_{\text{lb}})$, where $\text{kpABE.ct}'_{\text{lb}} \leftarrow \mathcal{CT}_{\text{kpABE}}$. (Note that this does not require kpABE.mpk).
 - Defines $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$, computes an FE key for this circuit and returns it as the secret key $\text{RPE.sk}_{\text{lb},x}$ to \mathcal{A} .
4. **Challenge Query:** When \mathcal{A} outputs the revocation list $L = \{\text{lb}_1, \dots, \text{lb}_{|L|}\}$ for the challenge ciphertext, \mathcal{B} does the following:
 - Sends lb_i to the kpABE challenger as the challenge attribute. The kpABE challenger samples $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$ and $\hat{\beta} \leftarrow \{0, 1\}$, and sends kpABE.mpk to \mathcal{B} .
 - Constructs circuit C_L as defined in the construction and sends a key query for C_L to the kpABE challenger and gets kpABE.sk_L in response.
 - Computes $\text{FE.ct} \leftarrow \text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f_0, m, K, 2, \perp, \text{SKE.key}_2))$
 - Returns $\text{RPE.ct} = (\text{kpABE.mpk}, \text{kpABE.sk}_L, \text{FE.ct})$ to \mathcal{A} .

5. **Post-challenge Key Queries:** For each post-challenge key query (lb, x) , \mathcal{B} computes flag , δ and γ_1 as defined for the hybrid and defines $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$, where γ_2 is computed as follows:

- if $\text{flag} = 0$, $\gamma_2 \leftarrow \mathcal{CT}_{\text{SKE}}$. Else,
 - if $\text{lb} \in L_{[1:i-1]}$, $\gamma_2 = \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct}'_{\text{lb}})$, where $\text{kpABE.ct}'_{\text{lb}} \leftarrow \mathcal{CT}_{\text{kpABE}}$.
 - if $\text{lb} = \text{lb}_i$, \mathcal{B} sends challenge query with message \tilde{m} to the kpABE challenger. The kpABE challenger returns a ciphertext $\text{kpABE.ct} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}_i, \tilde{m}; r)$, if $\hat{\beta} = 0$; else $\text{kpABE.ct} \leftarrow \mathcal{CT}_{\text{kpABE}}$. Then \mathcal{B} computes $\gamma_2 \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct})$.¹⁴
 - if $\text{lb} \notin L_{[1:i]}$, then $\gamma_2 \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct}'_{\text{lb}})$, where $\text{kpABE.ct}'_{\text{lb}} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; r)$.

\mathcal{B} computes an FE key for $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ and returns it as the secret key $\text{RPE.sk}_{\text{lb}, x}$ to \mathcal{A} .

6. In the end, \mathcal{A} outputs its guess bit β' . \mathcal{B} sends β' to the kpABE challenger as its guess bit.

We observe that if the kpABE challenger chose $\hat{\beta} = 0$, then \mathcal{B} simulated $\text{Hybrid}_{6a, i-1}$, else $\text{Hybrid}_{6a, i}$ with \mathcal{A} . Hence, advantage of $\mathcal{B} = |\Pr(\beta' = 1 | \hat{\beta} = 0) - \Pr(\beta' = 1 | \hat{\beta} = 1)| = |\Pr(\beta' = 1 | \text{Hybrid}_{6a, i-1}) - \Pr(\beta' = 1 | \text{Hybrid}_{6a, i})| = \epsilon$ (by assumption).

Admissibility of \mathcal{B} : Observe that \mathcal{B} issues a single ABE key query, which is for circuit C_L and the challenge attribute is $\text{lb}_i \in L$. Hence, by the design of C_L , $C_L(\text{lb}_i) = 0$ as desired. \square

\square

\square

Claim 4.5.1. Assume that FE is secure, then $\text{Hybrid}_7 \approx_c \text{Hybrid}_8$.

Proof. We show that if \mathcal{A} wins with non-negligible advantage ϵ in distinguishing the two hybrids, then there exists an adversary \mathcal{B} against FE security with the same advantage ϵ . The steps of the reduction are similar to the proof of Claim 4.2.3, with $\mu_0 = (\text{kpABE.mpk}, f_0, m, K, 2, \perp, \text{SKE.key}_2)$ and $\mu_1 = (\text{kpABE.mpk}, f_1, m, K, 2, \perp, \text{SKE.key}_2)$. Hence, here we only argue the admissibility of \mathcal{B} in the FE security game.

Admissibility of \mathcal{B} :

Firstly, we observe that the only key queries that \mathcal{B} issues to the FE challenger are for the $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ functions, where each function is defined corresponding to a key query (lb, x) by \mathcal{A} . Here, $\gamma_1 \leftarrow \mathcal{CT}_{\text{SKE}}$, $\delta \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{flag})$ and $\gamma_2 \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct}'_{\text{lb}})$ if $\text{flag} = 1$; else $\gamma_2 \leftarrow \mathcal{CT}_{\text{SKE}}$, where $\text{kpABE.ct}'_{\text{lb}} \leftarrow \mathcal{CT}_{\text{kpABE}}$ for $\text{flag} = 1$. Next, we observe that for any $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ function, $\mu_0 = (\text{kpABE.mpk}, f_0, m, K, 2, \perp, \text{SKE.key}_2)$, and $\mu_1 = (\text{kpABE.mpk}, f_1, m, K, 2, \perp, \text{SKE.key}_2)$ the following holds true from the definition of Re-Enc and correctness of SKE decryption,

¹⁴In case multiple queries with $\text{lb} = \text{lb}_i$ are made, we need to simulate the ciphertext multiple times. In that case, we rely on multi-challenge version of Sel-INDr, which is easily seen to be equivalent with the single challenge version.

- when $\text{flag} = 0$, i.e. $f_0(x) = f_1(x)$, \tilde{m} in $\text{kpABE.ct}_{\text{lb}} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; F(K, (\text{lb}, x)))$, computed inside the $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ function is same for both f_0 and f_1 , and hence same on both the inputs μ_0 and μ_1 . So,

$$\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_0) = \text{kpABE.ct}_{\text{lb}} \quad (\text{since mode} = 2 \text{ and flag} = 0).$$

$$\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_1) = \text{kpABE.ct}_{\text{lb}} \quad (\text{since mode} = 2 \text{ and flag} = 0).$$

- When $\text{flag} = 1$, i.e. $f_0(x) \neq f_1(x)$, by definition of $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ we have

$$\begin{aligned} \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_0) &= \text{SKE.Dec}(\text{SKE.key}_2, \gamma_2) \\ &= \text{kpABE.ct}'_{\text{lb}} \quad (\text{since mode} = 2 \text{ and flag} = 1). \end{aligned}$$

$$\begin{aligned} \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_1) &= \text{SKE.Dec}(\text{SKE.key}_2, \gamma_2) \\ &= \text{kpABE.ct}'_{\text{lb}} \quad (\text{since mode} = 2 \text{ and flag} = 1). \end{aligned}$$

This establishes the admissibility of \mathcal{B} . □

The rest of the hybrids, Hybrid_9 to Hybrid_{15} , are simply unwinding the previous hybrids and their proofs of indistinguishability are same as their corresponding counterparts in the first set of hybrids and hence, omitted. □

Message Hiding

Theorem 4.6. *Assume that F is a secure PRF, SKE is correct and secure, FE and kpABE are secure as per definitions 2.6 and 2.15, respectively. Furthermore, assume $|\mathcal{F}_\lambda| \leq \text{poly}(\lambda)$ and $|\mathcal{M}_\lambda| \leq \text{poly}(\lambda)$. Then the construction for RPE satisfies message hiding property as defined in Def. 3.3.*

Proof. Recall that for message hiding we want

$$\text{RPE.Enc}(\text{RPE.mpk}, f, m_0, L) \approx_c \text{RPE.Enc}(\text{RPE.mpk}, f, m_1, L),$$

where for all the key queries (lb, x) , either $f(x) = 0$ or $\text{lb} \in L$. The proof is given via a similar sequence of hybrid games between the challenger and a PPT adversary \mathcal{A} as in the proof of Theorem 4.2. The hybrids are defined as follows:

Hybrid_0 . This is the real world with $\beta = 0$, i.e., the challenge ciphertext is computed using the message m_0 . We write the complete game here to set up notations and easy reference in later hybrids.

1. The adversary outputs the challenge messages (m_0, m_1) and the challenge function f ¹⁵.
2. The challenger generates $(\text{FE.mpk}, \text{FE.msk}) \leftarrow \text{FE.Setup}(1^\lambda)$, sets $\text{RPE.mpk} = \text{FE.mpk}$ and sends RPE.mpk to the adversary.
3. **Key Queries :** When adversary issues a key query on (lb, x) , the challenger does the following:
 - Samples random values $\gamma_1, \gamma_2, \delta \leftarrow \mathcal{CT}_{\text{SKE}}$.

¹⁵To keep the proofs and notations simple, we let f, m_0, m_1 to be output selectively. This is sufficient to achieve security as in Def 3.3, as mentioned in Remark 3.9.

- Defines the circuit $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ as in the Figure 1.
 - Computes $\text{FE.sk}_{\text{lb},x} \leftarrow \text{FE.KeyGen}(\text{FE.msk}, \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta])$.
 - Returns $\text{RPE.sk}_{\text{lb},x} = \text{FE.sk}_{\text{lb},x}$ to the adversary.
4. **Challenge Query** : When the adversary outputs the revocation list L for the challenge ciphertext, the challenger does the following:
- Samples a PRF key K .
 - Generates $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$.
 - Computes FE.ct as
- $$\text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f, m_0, K, 0, \perp, \perp)).$$
- Defines C_L as in Eq. 4.1 and computes $\text{kpABE.sk}_L \leftarrow \text{kpABE.KeyGen}(\text{kpABE.msk}, C_L)$.
 - Returns $\text{RPE.ct} = (\text{kpABE.mpk}, \text{kpABE.sk}_L, \text{FE.ct})$ to the adversary.
5. In the end, the adversary outputs a bit β' .

Hybrid₁. This hybrid is same as the previous hybrid except the following changes:

- The challenger generates $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$ in the beginning of the game after the adversary outputs (m_0, m_1) and f . It also samples a SKE secret key SKE.key_1 and a PRF key K .
- For each key query (lb, x) , γ_1 is computed differently. In particular, the challenger does the following:
 - Sets $\tilde{m}_0 = \begin{cases} m_0 & \text{if } f(x) = 1, \\ 0 & \text{if } f(x) = 0. \end{cases}$
 - and computes $\text{kpABE.ct}'_{\text{lb}} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}_0; F(K, (\text{lb}, x)))$.
 - Sets γ_1 as $\text{SKE.Enc}(\text{SKE.key}_1, \text{kpABE.ct}'_{\text{lb}})$.

Hybrid₂. This hybrid is same as the previous hybrid except the following:

- The challenger samples two SKE secret keys SKE.key_1 and SKE.key_2 .
- For each key query (lb, x) , γ_2 and δ are computed differently as follows:
 - Set $\text{flag} = \begin{cases} 1 & \text{if } f(x) = 1, \\ 0 & \text{otherwise.} \end{cases}$
 - Set $\delta \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{flag})$ and $\gamma_2 \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct}'_{\text{lb}})$ if $\text{flag} = 1$, else $\gamma_2 \leftarrow \mathcal{CT}_{\text{SKE}}$.

Hybrid₃. This hybrid is same as the previous hybrid except that FE.ct in the challenge ciphertext is computed as

$$\text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f, m_0, \perp, 1, \text{SKE.key}_1, \text{SKE.key}_2)).$$

Hybrid₄. This hybrid is same as the previous hybrid except that for each key query (lb, x) $\text{kpABE.ct}'_{\text{lb}}$ is computed as follows:

- If $\text{flag} = 1$, $\text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}_0; r)$, where $r \leftarrow \{0, 1\}^t$,
- Else, if $\text{flag} = 0$, $\text{kpABE.ct}'_{\text{lb}} \leftarrow \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}_0; F(K, (\text{lb}, x)))$. (This is same as in the previous hybrid).

Hybrid₅. This hybrid is same as the previous hybrid except that FE.ct in the challenge ciphertext is computed as

$$\text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f, m_0, K, 2, \perp, \text{SKE.key}_2)).$$

Hybrid₆. This hybrid is same as the previous hybrid except that for all the key queries (lb, x) , $\gamma_1 \leftarrow \mathcal{CT}_{\text{SKE}}$.

Hybrid₇. This hybrid is same as the previous hybrid except that for each such key query (lb, x) where $\text{flag} = 1$, $\text{kpABE.ct}'_{\text{lb}} \leftarrow \mathcal{CT}_{\text{kpABE}}$.

Hybrid₈. This hybrid is same as the previous hybrid except that FE.ct in the challenge ciphertext is computed as

$$\text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f, m_1, K, 2, \perp, \text{SKE.key}_2)).$$

We note that the hybrids hereafter are unwinding the changes made in the previous hybrids.

Hybrid₉. This hybrid is same as the previous hybrid except that for each key query (lb, x) with $\text{flag} = 1$, $\text{kpABE.ct}'_{\text{lb}}$ is changed back to $\text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}_1; r)$, where $r \leftarrow \{0, 1\}^t$.

Hybrid₁₀. This hybrid is same as the previous hybrid except that for each key query (lb, x) , $\gamma_1 \leftarrow \text{SKE.Enc}(\text{SKE.key}_1, \text{kpABE.ct}'_{\text{lb}})$.

Hybrid₁₁. This hybrid is same as the previous hybrid except that FE.ct in the challenge ciphertext is computed as

$$\text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f, m_1, \perp, 1, \text{SKE.key}_1, \text{SKE.key}_2)).$$

Hybrid₁₂. This hybrid is same as the previous hybrid except that for each key query (lb, x) , $\text{kpABE.ct}'_{\text{lb}}$ is computed as $\text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}_1; F(K, (\text{lb}, x)))$.

Hybrid₁₃. This hybrid is same as the previous hybrid except that FE.ct in the challenge ciphertext is computed as

$$\text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f, m_1, K, 0, \perp, \perp)).$$

Hybrid₁₄. This hybrid is same as the previous hybrid except that the γ_2, δ are now sampled uniformly from $\mathcal{CT}_{\text{SKE}}$ for all the key queries.

Hybrid₁₅. This hybrid is same as the previous hybrid except that γ_1 is now sampled uniformly from $\mathcal{CT}_{\text{SKE}}$ for all the key queries. This is the real world where the message m_1 is encrypted in the challenge ciphertext.

Indistinguishability of hybrids: The indistinguishability between the consecutive hybrids is argued in the same way as that in the proof of Theorem 4.2. Therefore, here we give only a brief sketch.

Hybrid₀ \approx_c Hybrid₁ \approx_c Hybrid₂ from SKE security. Hybrid₂ \approx_c Hybrid₃ due to FE security and SKE correctness and Hybrid₃ \approx_c Hybrid₄ follows from PRF security. Hybrid₄ \approx_c Hybrid₅ follows from FE security and SKE correctness and Hybrid₅ \approx_c Hybrid₆ follows again from the SKE security. Hybrid₆ \approx_c Hybrid₇ follows from selective security of ABE. We observe that in

both Hybrid₆ and Hybrid₇, $\text{ABE.ct}'_{\text{lb}}$ is not used when $\text{flag} = 0$ and when $\text{flag} = 1$, it is sampled from $\mathcal{CT}_{\text{kpABE}}$ directly which can be efficiently done without using kpABE.mpk . This lets the reduction go through. The steps of reduction are same as in the proof of Claim 4.2.7. Hybrid₇ \approx_c Hybrid₈ follows again from the security of FE and SKE correctness. In particular, we observe that here the FE challenge messages are $\mu_0 = (\text{kpABE.mpk}, f, m_0, K, 2, \perp, \text{SKE.key}_2)$ and $\mu_1 = (\text{kpABE.mpk}, f, m_1, K, 2, \perp, \text{SKE.key}_2)$. For every $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ function (corresponding to RPE key query (lb, x)) for which FE key is generated, we have the following:

- When $\text{flag} = 0$, this implies $f(x) = 0$, which in turn implies that $\tilde{m}_0 = \tilde{m}_1 = 0$. Hence,

$$\begin{aligned}
& \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_0) \\
&= \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\text{kpABE.mpk}, f, m_0, K, 2, \perp, \text{SKE.key}_2) \\
&= \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, 0; F(K, (\text{lb}, x))) \\
&= \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\text{kpABE.mpk}, f, m_1, K, 2, \perp, \text{SKE.key}_2) \\
&= \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_1).
\end{aligned}$$

- When $\text{flag} = 1$,

$$\begin{aligned}
& \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_0) \\
&= \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\text{kpABE.mpk}, f, m_0, K, 2, \perp, \text{SKE.key}_2) \\
&= \text{SKE.Dec}(\text{SKE.key}_2, \gamma_2) \\
&= \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\text{kpABE.mpk}, f, m_1, K, 2, \perp, \text{SKE.key}_2) \\
&= \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_1).
\end{aligned}$$

This satisfies the admissibility condition for FE security. The rest of the hybrids undo the changes made so far to get to the real world with $\beta = 1$ and the arguments for indistinguishability are same as their counterparts in the first set of hybrids. \square

5 Revocable Mixed Functional Encryption

5.1 Definition

A revocable mixed functional encryption (RMFE) scheme with input domain $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda \in [\mathbb{N}]}$, a function family $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in [\mathbb{N}]}$ where $\mathcal{F}_\lambda = \{f : \mathcal{X}_\lambda \rightarrow \{0, 1\}\}$, a label space $\mathcal{L} = \{\mathcal{L}_\lambda\}_{\lambda \in [\mathbb{N}]}$ has the following syntax.

$\text{Setup}(1^\lambda) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm takes as input the security parameter λ and outputs a master public key mpk and a master secret key msk .

$\text{KeyGen}(\text{msk}, \text{lb}, x) \rightarrow \text{sk}_{\text{lb}, x}$. The key generation algorithm takes as input the master secret key msk , a label $\text{lb} \in \mathcal{L}_\lambda$ and an input $x \in \mathcal{X}_\lambda$. It outputs a secret key $\text{sk}_{\text{lb}, x}$.

$\text{PK-Enc}(\text{mpk}, L) \rightarrow \text{ct}$. The public key encryption algorithm takes as input the master public key mpk and a revocation list $L \subseteq \mathcal{L}_\lambda$ and outputs a ciphertext ct .

$\text{SK-Enc}(\text{msk}, f, L) \rightarrow \text{ct}$. The secret key encryption algorithm takes as input the master secret key msk , a function $f \in \mathcal{F}_\lambda$ and a revocation list $L \subseteq \mathcal{L}_\lambda$, and outputs a ciphertext ct .

$\text{Dec}(\text{sk}_{\text{lb}, x}, L, \text{ct}) \rightarrow \{0, 1\}$. The decryption algorithm takes the secret key $\text{sk}_{\text{lb}, x}$, a revocation list $L \subseteq \mathcal{L}_\lambda$ and a ciphertext ct and outputs a bit.

Definition 5.1 (Correctness). A RMFE scheme is said to be correct if there exists negligible functions $\text{negl}_1(\cdot), \text{negl}_2(\cdot)$ such that for all $\lambda \in \mathbb{N}$, the following holds

$$\Pr \left[\begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda); \\ \text{Dec}(\text{sk}_{\text{lb},x}, L, \text{ct}) = 1 : \text{sk}_{\text{lb},x} \leftarrow \text{KeyGen}(\text{msk}, \text{lb}, x); \\ \text{ct} \leftarrow \text{PK-Enc}(\text{mpk}, L) \end{array} \right] \geq 1 - \text{negl}_1(\lambda).$$

$$\text{lb} \notin L \Rightarrow \Pr \left[\begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda); \\ \text{Dec}(\text{sk}_{\text{lb},x}, L, \text{ct}) = f(x) : \text{sk}_{\text{lb},x} \leftarrow \text{KeyGen}(\text{msk}, \text{lb}, x); \\ \text{ct} \leftarrow \text{SK-Enc}(\text{msk}, f, L) \end{array} \right] \geq 1 - \text{negl}_2(\lambda).$$

Security. Here we define the security requirements of RMFE scheme.

Definition 5.2 (q -query Mode Hiding). Let $q(\cdot)$ be any fixed polynomial. A RMFE scheme satisfies q -query mode hiding security if for every PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for every $\lambda \in \mathbb{N}$,

$$\Pr \left[\begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda); \\ f, L \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot), \text{SK-Enc}(\text{msk}, \cdot, \cdot)}(\text{mpk}); \\ \beta' = \beta : \beta \leftarrow \{0, 1\}; \text{ct}_0 \leftarrow \text{SK-Enc}(\text{msk}, f, L); \\ \text{ct}_1 \leftarrow \text{PK-Enc}(\text{mpk}, L); \\ \beta' \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot), \text{SK-Enc}(\text{msk}, \cdot, \cdot)}(\text{ct}_\beta) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda)$$

where \mathcal{A} can make at most $q(\lambda)$ queries to the $\text{SK-Enc}(\text{msk}, \cdot, \cdot)$ oracle and is admissible only if for all the key queries (lb, x) to the $\text{KeyGen}(\text{msk}, \cdot, \cdot)$ oracle, $f(x) = 1$.

Definition 5.3 (q -query Selective Function Hiding). Let $q(\cdot)$ be any fixed polynomial. A RMFE scheme satisfies q -query selective function hiding security if for every PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for every $\lambda \in \mathbb{N}$,

$$\Pr \left[\begin{array}{l} L \leftarrow \mathcal{A}(1^\lambda); \\ (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda); \\ (f_0, f_1) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot), \text{SK-Enc}(\text{msk}, \cdot, \cdot)}(\text{mpk}); \\ \beta \leftarrow \{0, 1\}; \text{ct}_\beta \leftarrow \text{SK-Enc}(\text{msk}, f_\beta, L); \\ \beta' \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot), \text{SK-Enc}(\text{msk}, \cdot, \cdot)}(\text{ct}_\beta) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda)$$

where \mathcal{A} can make at most $q(\lambda)$ queries to the $\text{SK-Enc}(\text{msk}, \cdot, \cdot)$ oracle and for all the key queries (lb, x) to the $\text{KeyGen}(\text{msk}, \cdot, \cdot)$ oracle, either $f_0(x) = f_1(x)$ or $\text{lb} \in L$.

Remark 5.4. We note that when the function space \mathcal{F}_λ is polynomially small and q is a constant, a variant of Definition 5.3 where the adversary outputs the challenge functions (f_0, f_1) and the SK-Enc query functions $\{\bar{f}_i\}_{i \in [q]}$ at the beginning of the game, before the $\text{Setup}(1^\lambda)$ algorithm is run, is equivalent to Definition 5.3 where the adversary adaptively outputs the challenge functions (f_0, f_1) and can make SK-Enc queries adaptively, with polynomial loss. Similar comment also applies to Definition 5.2. We will use these simplifications in the security proofs.

5.2 Construction

In this section we give a construction of 1-query secure RMFE scheme, with input space $\mathcal{X} = \{\mathcal{X}_\lambda\}_\lambda$, a function family $\mathcal{F} = \{\mathcal{F}_\lambda\}_\lambda$ where $\mathcal{F}_\lambda = \{f : \mathcal{X}_\lambda \rightarrow \{0, 1\}\}$ and a label space $\mathcal{L} = \{\mathcal{L}_\lambda\}_\lambda$. We assume that the size of $|\mathcal{F}_\lambda|$ is bounded by some polynomial in λ , which will suffice for our purpose.

Our scheme uses the following building blocks:

1. A 2-bounded semi-adaptive simulation based function-message private (Definition 2.11) SKFE scheme $\text{SKFE} = (\text{SKFE.Setup}, \text{SKFE.KeyGen}, \text{SKFE.Enc}, \text{SKFE.Dec})$ that supports the function class \mathcal{F} . This can be instantiated from one-way functions (Lemma 2.12).
2. A key-policy ABE scheme $\text{kpABE} = (\text{kpABE.Setup}, \text{kpABE.Enc}, \text{kpABE.KeyGen}, \text{kpABE.Dec})$ for the circuit class $\mathcal{C}_{\ell(\lambda), d(\lambda)}$ with message space $\{0, 1\}^\lambda$ satisfying Sel-IND security (Definition 2.14) and efficiency properties described in Theorem 2.17. We set $\ell(\lambda) = \ell_{\text{lb}} + \log(\lambda) + 1$ and $d(\lambda) = \omega(\log \lambda)$, where ℓ_{lb} is the label length.¹⁶ This can be instantiated from the LWE assumption (Theorem 2.17).
3. A lockable obfuscation scheme $\text{LO} = (\text{LO.Obf}, \text{LO.Eval})$ with lock space $\{0, 1\}^\lambda$ that supports circuits of the form CC defined in Fig. 2. As we will analyze later, the circuit is of fixed polynomial size in λ and $|f|$, where $|f|$ is the description size of the function $f \in \mathcal{F}$. This can be instantiated from the LWE assumption (Theorem 2.24).

Below we describe our construction of a 1-query secure RMFE scheme $\text{RMFE} = (\text{RMFE.Setup}, \text{RMFE.KeyGen}, \text{RMFE.PK-Enc}, \text{RMFE.SK-Enc}, \text{RMFE.Dec})$.

$\text{RMFE.Setup}(1^\lambda) \rightarrow (\text{RMFE.mpk}, \text{RMFE.msk})$. The setup algorithm does the following:

- Generate $\text{SKFE.msk} \leftarrow \text{SKFE.Setup}(1^\lambda)$.
- Generate $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$.
- Output $\text{RMFE.mpk} = \text{kpABE.mpk}$ and $\text{RMFE.msk} = (\text{SKFE.msk}, \text{kpABE.mpk}, \text{kpABE.msk})$.

$\text{RMFE.KeyGen}(\text{RMFE.msk}, \text{lb}, x) \rightarrow \text{RMFE.sk}_{\text{lb}, x}$. The key generation algorithm does the following:

- Parse $\text{RMFE.msk} = (\text{SKFE.msk}, \text{kpABE.mpk}, \text{kpABE.msk})$.
- For all $j \in [\lambda], b \in \{0, 1\}$, sample $K_{j,b}, R_{j,b} \leftarrow \{0, 1\}^\lambda$.
Denote $K = \{K_{j,b}\}_{j \in [\lambda], b \in \{0,1\}}$ and $R = \{R_{j,b}\}_{j \in [\lambda], b \in \{0,1\}}$.
- Compute

$$\text{SKFE.ct} \leftarrow \text{SKFE.Enc}(\text{SKFE.msk}, (x, K, R)).$$
- For all $j \in [\lambda], b \in \{0, 1\}$, compute

$$\text{kpABE.ct}_{\text{lb}, j, b} \leftarrow \text{kpABE.Enc}(\text{kpABE.mpk}, (\text{lb}, j, b), K_{j,b}).$$
- Output $\text{RMFE.sk}_{\text{lb}, x} = (\text{SKFE.ct}, \text{kpABE.mpk}, \{(\text{lb}, j, b), \text{kpABE.ct}_{\text{lb}, j, b}\}_{j \in [\lambda], b \in \{0,1\}})$.

$\text{RMFE.PK-Enc}(\text{RMFE.mpk}, L) \rightarrow \text{RMFE.ct}$. The public key encryption algorithm does the following:

- Computes a simulated code $\text{RMFE.ct} \leftarrow \text{LO.Sim}(1^\lambda, 1^{|\text{CC}|})$ ¹⁷.
- It outputs RMFE.ct as the ciphertext.

$\text{RMFE.SK-Enc}(\text{RMFE.msk}, f, L) \rightarrow \text{RMFE.ct}$. The secret key encryption algorithm does the following:

- Parse $\text{RMFE.msk} = (\text{SKFE.msk}, \text{kpABE.mpk}, \text{kpABE.msk})$, and sample a tag $\mathbf{z} \leftarrow \{0, 1\}^\lambda$ and a lock value $\alpha \leftarrow \{0, 1\}^\lambda$.

¹⁶Concretely, we can choose $d(\lambda) = \Theta(\log \lambda \log \log \lambda)$ for example.

¹⁷Here, CC represents the maximum possible size of $\text{CC}[\cdot, \cdot]$ circuit defined in Figure 2.

- For all $j \in [\lambda]$, compute $\text{kpABE.sk}_{L,j,z_j} \leftarrow \text{kpABE.KeyGen}(\text{kpABE.msk}, C_{L,j,z_j})$, where the function C_{L,j,z_j} has L, j and z_j hardwired and is defined as follows :
On input $(\text{lb}, i, b) \in \mathcal{L}_\lambda \times [\lambda] \times \{0, 1\}$,

$$C_{L,j,z_j}(\text{lb}, i, b) = \begin{cases} 1 & \text{if } (\text{lb} \notin L) \wedge (i = j) \wedge (b = z_j) \\ 0 & \text{otherwise.} \end{cases} \quad (5.1)$$

- Compute $\text{SKFE.sk} \leftarrow \text{SKFE.KeyGen}(\text{SKFE.msk}, P_{f,\mathbf{z},\alpha})$, where the function $P_{f,\mathbf{z},\alpha}$ has f, \mathbf{z}, α hardwired and is defined as follows :
On input $x \in \mathcal{X}_\lambda$, $K = \{K_{j,b}\}_{j \in [\lambda], b \in \{0,1\}}$, $R = \{R_{j,b}\}_{j \in [\lambda], b \in \{0,1\}}$,

$$P_{f,\mathbf{z},\alpha}(x, K, R) = \begin{cases} \bigoplus_j K_{j,z_j} \oplus \alpha & \text{if } f(x) = 0 \\ \bigoplus_j R_{j,z_j} & \text{if } f(x) = 1. \end{cases} \quad (5.2)$$

- Construct function $\text{CC}[\text{SKFE.sk}, \{\text{kpABE.sk}_{L,j,z_j}\}_{j \in [\lambda]}]$, with SKFE.sk and $\{\text{kpABE.sk}_{L,j,z_j}\}_{j \in [\lambda]}$ hardwired and is defined as in Figure 2.
- Output $\text{RMFE.ct} \leftarrow \text{LO.Obf}(\text{CC}[\text{SKFE.sk}, \{\text{kpABE.sk}_{L,j,z_j}\}_{j \in [\lambda]}], \alpha)$.

$\text{RMFE.Dec}(\text{RMFE.sk}_{\text{lb},x}, \text{RMFE.ct}, L) \rightarrow \{0, 1\}$. The decryption algorithm does the following:

- Parse $\text{RMFE.sk}_{\text{lb},x} = (\text{SKFE.ct}, \text{kpABE.mpk}, \{(\text{lb}, j, b), \text{kpABE.ct}_{\text{lb},j,b}\}_{j \in [\lambda], b \in \{0,1\}})$ and $\text{RMFE.ct} = \widetilde{\text{CC}}$, where $\widetilde{\text{CC}}$ is regarded as an obfuscated circuit of LO.
- For all $j \in [\lambda], b \in \{0, 1\}$, compute

$$\text{kpABE.off}_{\text{lb},j,b} \leftarrow \text{kpABE.Dec}^{\text{off}}(\text{kpABE.mpk}, C_{L,j,b}, (\text{lb}, j, b)).$$

- Compute

$$y = \text{LO.Eval}(\widetilde{\text{CC}}, (\text{SKFE.ct}, \{\text{kpABE.ct}_{\text{lb},j,b}, \text{kpABE.off}_{\text{lb},j,b}\}_{j \in [\lambda], b \in \{0,1\}})).$$

- Output 1 if $y = \perp$, else output 0.

Remark 5.5. We note that by performing the part of the ABE decryption that uses $C_{L,j,b}$ outside of CC, we do not need to provide $C_{L,j,b}$ (or L) as input to CC. Instead, we provide $\{\text{kpABE.off}_{\text{lb},j,b}\}_{j \in [\lambda], b \in \{0,1\}}$ whose size is independent of the size of $C_{L,j,b}$ (and thus that of L). This helps us in getting succinct ciphertext.

Correctness. We prove the correctness via the following theorem.

Theorem 5.6. *Suppose kpABE, LO and SKFE are correct and LO is secure, then the above construction of RMFE satisfies correctness as defined in Def. 5.1.*

Proof. We consider the following two cases:

1. Public Encryption Correctness.

For $\text{RMFE.ct} \leftarrow \text{RMFE.PK-Enc}(\text{RMFE.mpk}, L)$, we have that $\text{RMFE.ct} = \text{LO.Sim}(1^\lambda, 1^{|\text{CC}|})$. Firstly, we note that from LO security, RMFE.ct is indistinguishable from $\text{RMFE.ct}'$ computed as $\text{LO.Obf}(C, \alpha)$, for any circuit C of the same size as that used by the simulator and has output of length λ . Now, since $\alpha \leftarrow \{0, 1\}^\lambda$ has high entropy, for all but negligible inputs w , $C(w) \neq \alpha$. So, from the correctness of LO, it follows that with all but negligible probability

$$\text{LO.Eval}(\text{RMFE.ct}, (\text{SKFE.ct}, \{\text{kpABE.ct}_{\text{lb},j,b}, \text{kpABE.off}_{\text{lb},j,b}\}_{j \in [\lambda], b \in \{0,1\}})) = \perp.$$

Hence the RMFE.Dec algorithm outputs 1 with all but negligible probability.

Function $\text{CC}[\text{SKFE.sk}, \{\text{kpABE.sk}_{L,j,z_j}\}_{j \in [\lambda]}]$

Hardwired values: A SKFE secret key SKFE.sk and kpABE keys $\{\text{kpABE.sk}_{L,j,z_j}\}_{j \in [\lambda]}$.

Inputs: A SKFE ciphertext SKFE.ct and kpABE ciphertexts $\{\text{kpABE.ct}_{\text{lb},j,b}, \text{kpABE.off}_{\text{lb},j,b}\}_{j \in [\lambda], b \in \{0,1\}}$.

Output : A binary string $\alpha^* \in \{0, 1\}^\lambda$.

1. For all $j \in [\lambda]$, compute $m_j = \text{kpABE.Dec}^{\text{on}}(\text{kpABE.sk}_{L,j,z_j}, \text{kpABE.ct}_{\text{lb},j,z_j}, \text{kpABE.off}_{\text{lb},j,z_j})$.
Let $M_0 = \bigoplus_j m_j$
2. Compute $M_1 = \text{SKFE.Dec}(\text{SKFE.sk}, \text{SKFE.ct})$.
3. Output $M_1 \oplus M_0$.

Figure 2: Compute and Compare function CC

2. Secret Encryption Correctness.

For $\text{RMFE.ct} \leftarrow \text{RMFE.SK-Enc}(\text{RMFE.msk}, f, L)$, we have $\text{RMFE.ct} = \text{LO.Obf}(\text{CC}[\text{SKFE.sk}, \{\text{kpABE.sk}_{L,j,z_j}\}_{j \in [\lambda]}], \alpha)$. Now consider the following steps of

$\text{LO.Eval}(\text{RMFE.ct}, (\text{SKFE.ct}, \{\text{kpABE.ct}_{\text{lb},j,b}, \text{kpABE.off}_{\text{lb},j,b}\}_{j \in [\lambda], b \in \{0,1\}}))$

- Since $\text{lb} \notin L$, by correctness of kpABE, for all $j \in [\lambda]$, $M_0 = \bigoplus_j m_j = \bigoplus_j K_{j,z_j}$ with all but negligible probability.
- By correctness of SKFE, if $f(x) = 0$, we have $M_1 = \bigoplus_j K_{j,z_j} \oplus \alpha$, else $M_1 = \bigoplus_j R_{j,z_j}$.
- We have $M_0 \oplus M_1 = \alpha$ if $f(x) = 0$.

So, by the correctness of LO, LO.Eval outputs 1 if $f(x) = 0$, \perp otherwise with all but negligible probability. Hence the RMFE.Dec algorithm, by construction, outputs 0 if $f(x) = 0$ and 1 if $f(x) = 1$ with all but negligible probability.

This proves the correctness of the above construction. □

Efficiency. Here we argue that our construction achieves optimal parameters. Namely, we show that the sizes of the parameters are independent of $|L|$. We first observe that $C_{L,j,b}$ as defined in Eq (5.1) can be implemented with depth $d = \omega(\log \lambda)$, since the membership check $\text{lb} \stackrel{?}{\in} L$ can be done with depth $\log(|\text{lb}| \cdot |L|) = \log(\text{poly}(\lambda))$ and the equality check can be done with depth $\log(|j|) = \log \log \lambda$. We then bound the size of parameters.

1. Public key size $|\text{RMFE.mpk}|$: By the efficiency property of kpABE (Theorem 2.17), we have $|\text{RMFE.mpk}| = |\text{kpABE.mpk}| = \text{poly}(\lambda, d, |\text{lb}|) = \text{poly}(\lambda, |\text{lb}|)$.
2. Secret key size $|\text{RMFE.sk}_{\text{lb},x}|$: We have $|\text{RMFE.sk}_{\text{lb},x}| = |\text{SKFE.ct}| + |\text{kpABE.mpk}| + 2 \cdot \lambda(|(\text{lb}, j, b)| + |\text{kpABE.ct}|)$. The first term can be bounded by $\text{poly}(\lambda, |f|, |x|)$, the second is $\text{poly}(\lambda, |\text{lb}|)$, and the last terms is $\text{poly}(\lambda, |\text{lb}|)$. Therefore, the total size is $\text{poly}(\lambda, |f|, |x|, |\text{lb}|)$.

3. Ciphertext size $|\text{RMFE.ct}|$: We first bound the size of the circuit CC defined in Fig. 2. The dominant operations in the circuit is the decryption of SKFE and the online decryption of kpABE. We can see that the former is implemented by a circuit of size $\text{poly}(\lambda, |x|, |f|)$ by the efficiency of SKFE. The latter can be implemented by a circuit of size $\text{poly}(\lambda, \ell, d) = \text{poly}(\lambda, |\text{lb}|)$ by the online efficiency of kpABE (Theorem 2.17). Therefore, the total size CC is $\text{poly}(\lambda, |f|, |x|, |\text{lb}|)$. By the efficiency of LO, the size of the ciphertext is $\text{poly}(\lambda, |f|, |x|, |\text{lb}|)$ as well.
4. Depth of the circuit implementing RMFE.Dec: We also evaluate the depth of the circuit implementing RMFE.Dec and show that it is independent from $|L|$, since it will be used later in Sec. 6. We first observe that the dominant operations in RMFE.Dec are the offline decryption of kpABE and the evaluation of CC. The depth of the former can be bounded by $\text{poly}(\lambda, \ell, \text{depth}(C_{L,j,b})) \leq \text{poly}(\lambda, |\text{lb}|)$ by Theorem 2.17. The depth of the latter can be bounded by its size and thus is $\text{poly}(\lambda, |f|, |x|, |\text{lb}|)$ as we have seen in the previous item. The total depth is thus bounded by $\text{poly}(\lambda, |f|, |x|, |\text{lb}|)$, which is independent from $|L|$.

5.3 Security

In this section we show that our construction of RMFE scheme satisfies all the security requirements.

We will use the following notations in the security proof:

For $X \in \{K, R\}$,

- $X := \{X_{j,b}\}_{j \in [\lambda], b \in \{0,1\}}$, $X_b := \{X_{j,b}\}_{j \in [\lambda]}$, $X_j := \{X_{j,b}\}_{b \in \{0,1\}}$.
- For any vector $\mathbf{z} \in \{0,1\}^\lambda$, $X_{\mathbf{z}} := \{X_{j,z_j}\}_{j \in [\lambda]}$.

Mode hiding

Theorem 5.7. *Assume that SKFE and LO are secure as per definitions 2.11 and 2.23, respectively. Furthermore, assume $|\mathcal{F}_\lambda| \leq \text{poly}(\lambda)$. Then the RMFE construction satisfies 1-query mode hiding security as per Definition 5.2.*

Proof. Recall that for mode hiding, we need

$$\text{RMFE.SK-Enc}(\text{RMFE.msk}, f^*, L^*) \approx_c \text{RMFE.PK-Enc}(\text{RMFE.mpk}, L^*),$$

where for all key queries (lb, x) , $f^*(x) = 1$.

The proof proceeds via the following sequence of hybrid games between the challenger and a PPT adversary \mathcal{A} .

Hybrid₀ : This is the real world with $\beta = 0$, where the challenge ciphertext for (f^*, L^*) is computed using the RMFE.SK-Enc algorithm. We write the complete game here to set up the notations and easy reference in later hybrids.

1. The adversary outputs the challenge function f^* and the SK-Enc query function \bar{f} ¹⁸.
2. The challenger generates $\text{SKFE.msk} \leftarrow \text{SKFE.Setup}(1^\lambda)$,
 $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$, sets $\text{RMFE.mpk} = \text{kpABE.mpk}$ and $\text{RMFE.msk} = (\text{SKFE.msk}, \text{kpABE.mpk}, \text{kpABE.msk})$. It sends RMFE.mpk to \mathcal{A} .

¹⁸To keep the proofs and notations simple, we let f^* and \bar{f} to be given selectively. This is sufficient to achieve security as in Def 5.2, as mentioned in Remark 5.4.

3. **Key Queries:** For each key query (lb, x) , the challenger computes SKFE.ct and kpABE.ct_{lb,j,b} as in the construction and returns RMFE.sk_{lb,x} = (SKFE.ct, kpABE.mpk, $\{(lb, j, b), \text{kpABE.ct}_{lb,j,b}\}_{j \in [\lambda], b \in \{0,1\}}$) to \mathcal{A} .
4. **Challenge Query:** When the adversary outputs L^* for the challenge query, the challenger does the following:
 - Samples a tag $z^* \leftarrow \{0, 1\}^\lambda$ and a lock value $\alpha^* \leftarrow \{0, 1\}^\lambda$.
 - Computes kpABE secret keys $\{\text{kpABE.sk}_{L^*,j,z_j^*}\}_{j \in [\lambda]}$ and a SKFE secret key SKFE.sk^{*} as in the construction.
 - Constructs $\text{CC}[\text{SKFE.sk}^*, \{\text{kpABE.sk}_{L^*,j,z_j^*}\}_{j \in [\lambda]}]$ as defined in Figure 2 and returns $\text{RMFE.ct}^* \leftarrow \text{LO.Obf}(\text{CC}[\text{SKFE.sk}^*, \{\text{kpABE.sk}_{L^*,j,z_j^*}\}_{j \in [\lambda]}], \alpha^*)$ to the adversary \mathcal{A} .
5. **SK-Enc Query:** When the adversary outputs \bar{L} for the SK-Enc query, the challenger does the following:
 - Samples a tag $\bar{z} \leftarrow \{0, 1\}^\lambda$ and a lock value $\bar{\alpha} \leftarrow \{0, 1\}^\lambda$.
 - Computes kpABE secret keys $\{\text{kpABE.sk}_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}$ and a SKFE secret key SKFE. $\bar{\text{sk}}$ as in the construction.
 - Constructs $\text{CC}[\text{SKFE.}\bar{\text{sk}}, \{\text{kpABE.sk}_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}]$ and returns $\text{RMFE.}\bar{\text{ct}} \leftarrow \text{LO.Obf}(\text{CC}[\text{SKFE.}\bar{\text{sk}}, \{\text{kpABE.sk}_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}], \bar{\alpha})$ to the adversary \mathcal{A} .
6. In the end, \mathcal{A} outputs a bit β' .

Hybrid₁ : This hybrid is same as the previous hybrid except the following changes:

1. The challenger samples $\alpha^*, \bar{\alpha}, z^*, \bar{z}$ in the beginning of the game after the adversary outputs f^*, \bar{f} .
2. The challenger then computes $\text{SKFE.sk}^* \leftarrow \text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{|P_{f^*,z^*,\alpha^*}|})$ and $\text{SKFE.}\bar{\text{sk}} \leftarrow \text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{\text{poly}(|P_{\bar{f},\bar{z},\bar{\alpha}}|)})$ in this order using SKFE simulators.
3. The key generation, challenge and SK-Enc queries are answered as follows:
 - For each key query (lb, x) , the SKFE ciphertext in RMFE.sk_{lb,x} is computed as $\text{SKFE.Sim}^{\text{CT}}(\text{SKFE.msk}, \bigoplus_j R_{j,z_j^*}, P_{\bar{f},\bar{z},\bar{\alpha}}(x, K, R))$.
Note that $P_{f^*,z^*,\alpha^*}(x, K, R) = \bigoplus_j R_{j,z_j^*}$, due to admissibility requirement.
 - To answer the challenge and the SK-Enc queries, SKFE.sk^{*} and SKFE. $\bar{\text{sk}}$ generated by SKFE simulators in Step 2 are used for generating RMFE.ct^{*} and RMFE. $\bar{\text{ct}}$, respectively.

Hybrid₂ : This hybrid is same as the previous hybrid except that the challenger uses LO.Sim to generate the challenge ciphertext.

$$\text{RMFE.ct}^* = \text{LO.Sim}(1^\lambda, 1^{|\text{CC}|})^{19}.$$

Hybrid₃ : This hybrid is same as the previous hybrid except that the challenger uses SKFE.Enc and SKFE.KeyGen to generate SKFE ciphertexts and keys respectively. Formally,

$$\text{SKFE.}\bar{\text{sk}} = \text{SKFE.KeyGen}(\text{SKFE.msk}, P_{\bar{f},\bar{z},\bar{\alpha}})$$

¹⁹Here, $|\text{CC}|$ represents the maximum size of the circuit $\text{CC}[\cdot, \cdot]$ defined in Figure 2.

For each key query (lb, x) ,

$$\text{SKFE.ct} = \text{SKFE.Enc}(\text{SKFE.msk}, (K, R, x)),$$

where vectors K and R are freshly sampled for each key as defined in the construction. This is the real world with $\beta = 1$, where the challenge ciphertext is computed using the RMFE.PK-Enc algorithm.

Indistinguishability of hybrids

Claim 5.7.1. *Assume that SKFE is secure (Def. 2.11), then $\text{Hybrid}_0 \approx_c \text{Hybrid}_1$.*

Proof. We show that if there exists a PPT adversary \mathcal{A} who can distinguish between Hybrid_0 and Hybrid_1 with non-negligible advantage ϵ , then there exists a PPT adversary \mathcal{B} against the security of SKFE scheme with the same advantage ϵ . The reduction is as follows.

1. Upon being invoked by the SKFE challenger, \mathcal{B} invokes \mathcal{A} which outputs f^*, \bar{f} .
2. \mathcal{B} samples $\alpha^*, \bar{\alpha}, \mathbf{z}^*, \bar{\mathbf{z}} \leftarrow \{0, 1\}^\lambda$.
3. \mathcal{B} defines the functions $P_{f^*, \mathbf{z}^*, \alpha^*}$ and $P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}}$ as defined in Eq. (5.2) and sends it to the SKFE challenger in this order as key queries. The challenger generates $\text{SKFE.msk} \leftarrow \text{SKFE.Setup}(1^\lambda)$ and samples $\hat{\beta} \leftarrow \{0, 1\}$.
If $\hat{\beta} = 0$, it computes $\text{SKFE.sk}^* = \text{SKFE.KeyGen}(\text{SKFE.msk}, P_{f^*, \mathbf{z}^*, \alpha^*})$ and $\text{SKFE.s}\bar{\mathbf{k}} = \text{SKFE.KeyGen}(\text{SKFE.msk}, P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}})$
else it computes $\text{SKFE.sk}^* = \text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{|P_{f^*, \mathbf{z}^*, \alpha^*}|})$, $\text{SKFE.s}\bar{\mathbf{k}} = \text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{|P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}}|})$ and sends $\{\text{SKFE.sk}^*, \text{SKFE.s}\bar{\mathbf{k}}\}$ to \mathcal{B} .
4. \mathcal{B} generates $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$, sets $\text{RMFE.mpk} = \text{kpABE.mpk}$ and sends RMFE.mpk to \mathcal{A} .
5. **Key Queries:** For each key query (lb, x) , \mathcal{B} does the following:
 - Samples $K_{j,b}, R_{j,b} \leftarrow \{0, 1\}^\lambda, \forall j \in [\lambda], b \in \{0, 1\}$.
 - Computes $\text{kpABE.ct}_{lb,j,b} \leftarrow \text{kpABE.Enc}(\text{kpABE.msk}, (lb, j, b), K_{j,b})$ for $j \in [\lambda], b \in \{0, 1\}$.
 - It sends (x, K, R) as challenge message to the SKFE challenger. The challenger returns $\text{SKFE.ct}_{\hat{\beta}}$, where $\text{SKFE.ct}_0 = \text{SKFE.Enc}(\text{SKFE.msk}, (x, K, R))$ and $\text{SKFE.ct}_1 = \text{SKFE.Sim}^{\text{CT}}(\text{SKFE.msk}, \bigoplus_j R_{j,z_j^*}, P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}}(x, K, R))$.
 - Returns $\text{RMFE.sk}_{lb,x} = (\text{SKFE.ct}_{\hat{\beta}}, \text{kpABE.mpk}, \{(lb, j, b), \text{kpABE.ct}_{lb,j,b}\}_{j,b})$ to \mathcal{A} .
6. **Challenge Query:** When \mathcal{A} outputs L^* for the challenge query, \mathcal{B} does the following:
 - Computes kpABE secret keys $\{\text{kpABE.sk}_{L^*,j,z_j^*}\}_{j \in [\lambda]}$.
 - Constructs $\text{CC}[\text{SKFE.sk}^*, \{\text{kpABE.sk}_{L^*,j,z_j^*}\}_{j \in [\lambda]}]$ as defined in Figure 2 and returns $\text{RMFE.ct}^* \leftarrow \text{LO.Obf}(\text{CC}[\text{SKFE.sk}^*, \{\text{kpABE.sk}_{L^*,j,z_j^*}\}_{j \in [\lambda]}], \alpha^*)$ to \mathcal{A} .
7. **SK-Enc Query:** When the adversary outputs \bar{L} for the SK-Enc query, \mathcal{B} does the following:
 - Computes kpABE secret keys $\{\text{kpABE.sk}_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}$.
 - Constructs $\text{CC}[\text{SKFE.s}\bar{\mathbf{k}}, \{\text{kpABE.sk}_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}]$ and returns $\text{RMFE.c}\bar{\mathbf{t}} \leftarrow \text{LO.Obf}(\text{CC}[\text{SKFE.s}\bar{\mathbf{k}}, \{\text{kpABE.sk}_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}], \bar{\alpha})$ to \mathcal{A} .

8. In the end, \mathcal{A} outputs a bit β' . \mathcal{B} forwards β' as its guess bit to the SKFE challenger.

We observe that if $\hat{\beta} = 0$, then \mathcal{B} simulated Hybrid₀, else Hybrid₁ with \mathcal{A} . Hence, advantage of $\mathcal{B} = |\Pr(\beta' = 1 | \hat{\beta} = 0) - \Pr(\beta' = 1 | \hat{\beta} = 1)| = |\Pr(\beta' = 1 | \text{Hybrid}_0) - \Pr(\beta' = 1 | \text{Hybrid}_1)| = \epsilon$ (by assumption). \square

Claim 5.7.2. Assume that LO is secure (Def. 2.23), then Hybrid₁ \approx_c Hybrid₂.

Proof. We show that if there exists a PPT adversary \mathcal{A} who can distinguish between Hybrid₁ and Hybrid₂ with non-negligible advantage ϵ , then there exists a PPT adversary \mathcal{B} against the security of LO scheme with the same advantage ϵ . The reduction is as follows

1. Upon being invoked by the LO challenger, \mathcal{B} invokes \mathcal{A} . \mathcal{A} outputs f^*, \bar{f} .
2. \mathcal{B} samples $\bar{\alpha}, \mathbf{z}^*, \bar{z} \leftarrow \{0, 1\}^\lambda$. It also defines the function $P_{\bar{f}, \bar{z}, \bar{\alpha}}$.
3. \mathcal{B} generates SKFE.msk \leftarrow SKFE.Setup(1^λ), (kpABE.mpk, kpABE.msk) \leftarrow kpABE.Setup(1^λ) and sets RMFE.mpk = kpABE.mpk and RMFE.msk = (SKFE.msk, kpABE.mpk, kpABE.msk) and sends RMFE.mpk to \mathcal{A} .
4. **Key Queries:** For each key query (lb, x) , \mathcal{B} does the following:
 - Samples $K_{j,b}, R_{j,b} \leftarrow \{0, 1\}^\lambda, \forall j \in [\lambda], b \in \{0, 1\}$.
 - Computes $\text{kpABE.ct}_{lb,j,b} \leftarrow \text{kpABE.Enc}(\text{kpABE.mpk}, (lb, j, b), K_{j,b}) \forall j \in [\lambda], b \in \{0, 1\}$ and $\text{SKFE.ct} \leftarrow \text{SKFE.Sim}^{\text{CT}}(\text{SKFE.msk}, \bigoplus_j R_{j,z_j^*}, P_{\bar{f}, \bar{z}, \bar{\alpha}}(x, K, R))$.
 - Returns $\text{RMFE.sk}_{lb,x} = (\text{SKFE.ct}, \text{kpABE.mpk}, \{(lb, j, b), \text{kpABE.ct}_{lb,j,b}\}_{j,b})$ to \mathcal{A} .
5. **Challenge Query:** When the adversary outputs L^* for the challenge ciphertext, \mathcal{B} does the following:
 - Computes kpABE secret keys $\{\text{kpABE.sk}_{L^*,j,z_j^*}\}_{j \in [\lambda]}$.
 - Constructs a function $\text{CC}[\text{SKFE.sk}^*, \{\text{kpABE.sk}_{L^*,j,z_j^*}\}_{j \in [\lambda]}]$ as defined in Figure 2, where $\text{SKFE.sk}^* \leftarrow \text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{|\text{P}_{f^*, \mathbf{z}^*, \alpha^*}|})$ ²⁰.
 - It sends $\text{CC}[\text{SKFE.sk}^*, \text{kpABE.sk}_{L^*,j,z_j^*}]$ to the LO challenger. The challenger samples a lock value $\alpha^* \leftarrow \{0, 1\}^\lambda$ and $\hat{\beta} \leftarrow \{0, 1\}$, computes and return $\text{Obf}_{\hat{\beta}}$, where $\text{Obf}_0 = \text{LO.Obf}(\text{CC}[\text{SKFE.sk}^*, \{\text{kpABE.sk}_{L^*,j,z_j^*}\}_{j \in [\lambda]}], \alpha^*)$ and $\text{Obf}_1 = \text{LO.Sim}(1^\lambda, 1^{|\text{CC}|})$.
 - It sets $\text{RMFE.ct}^* = \text{Obf}_{\hat{\beta}}$ and sends it to the adversary \mathcal{A} .
6. **SK-Enc Query:** When the adversary outputs \bar{L} for the SK-Enc query, \mathcal{B} does the following:
 - Computes kpABE secret keys $\{\text{kpABE.sk}_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}$.
 - Constructs $\text{CC}[\text{SKFE.s}\bar{k}, \{\text{kpABE.sk}_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}]$, where $\text{SKFE.s}\bar{k} \leftarrow \text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{|\text{P}_{\bar{f}, \bar{z}, \bar{\alpha}}|})$.
 - Computes $\text{RMFE.c}\bar{t} \leftarrow \text{LO.Obf}(\text{CC}[\text{SKFE.s}\bar{k}, \{\text{kpABE.sk}_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}], \bar{\alpha})$ and returns it to the adversary \mathcal{A} .
7. In the end, \mathcal{A} outputs a bit β' . \mathcal{B} forwards β' as its guess bit to the LO challenger.

²⁰The size of $\text{P}_{f^*, \mathbf{z}^*, \alpha^*}$ is independent of any specific lock value and hence can be computed without the knowledge of α^* .

We observe that if $\hat{\beta} = 0$, then \mathcal{B} simulated Hybrid₁, else Hybrid₂ with \mathcal{A} . Hence, advantage of $\mathcal{B} = |\Pr(\beta' = 1|\hat{\beta} = 0) - \Pr(\beta' = 1|\hat{\beta} = 1)| = |\Pr(\beta' = 1|\text{Hybrid}_1) - \Pr(\beta' = 1|\text{Hybrid}_2)| = \epsilon$ (by assumption). \square

Claim 5.7.3. Assume that SKFE is secure (Def. 2.11), then Hybrid₂ \approx_c Hybrid₃.

The proof of this claim is similar to that of Claim 5.7.1, hence omitted. \square

Function Hiding

Theorem 5.8. Assume SKFE is secure (Def. 2.11), kpABE satisfies Sel-IND security (Def. 2.14). Furthermore, assume $|\mathcal{F}_\lambda| \leq \text{poly}(\lambda)$. Then, the RMFE construction satisfies 1-query function hiding as defined in Definition 5.3.

Proof. We prove the theorem via the following sequence of hybrids.

Hybrid₀ : This is the real world with $\beta = 0$. We summarize the steps of the game here to set up the notations used in the later hybrids.

1. The adversary outputs the challenge query f_0, f_1, L^* and the SK-Enc query function \bar{f} in the beginning of the game²¹. The challenger then does the following:
 - Generates $\text{SKFE.msk} \leftarrow \text{SKFE.Setup}(1^\lambda)$, $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$, sets $\text{RMFE.mpk} = \text{kpABE.mpk}$ and $\text{RMFE.msk} = (\text{SKFE.msk}, \text{kpABE.mpk}, \text{kpABE.msk})$.
 - It also samples a tag $\mathbf{z}^* \in \{0, 1\}^\lambda$ and a lock value α^* .
 - Computes $\text{SKFE.sk}^* \leftarrow \text{SKFE.KeyGen}(\text{SKFE.msk}, P_{f_0, \mathbf{z}^*, \alpha^*})$ and $\{\text{kpABE.sk}_{L^*, j, z_j^*}\}_{j \in [\lambda]}$ as defined in the construction.
 - Returns RMFE.mpk and the challenge ciphertext $\text{RMFE.ct}^* \leftarrow \text{LO.Obf}(\text{CC}[\text{SKFE.sk}^*, \{\text{kpABE.sk}_{L^*, j, z_j^*}\}_{j \in [\lambda]}], \alpha^*)$ to \mathcal{A} .
2. **Key Queries:** For each key query (lb, x) , the challenger returns $\text{RMFE.sk}_{\text{lb}, x} = (\text{SKFE.ct}, \{(\text{lb}, j, b), \text{kpABE.ct}_{\text{lb}, j, b}\}_{j \in [\lambda], b \in \{0, 1\}})$, where $\text{SKFE.ct}, \{\text{kpABE.ct}_{\text{lb}, j, b}\}_{j \in [\lambda], b \in \{0, 1\}}$ are computed as in the construction.
3. **SK-Enc Query :** When the adversary outputs \bar{L} for the SK-Enc query, the challenger does the following:
 - Samples a tag $\bar{\mathbf{z}} \leftarrow \{0, 1\}^\lambda$ and a lock value $\bar{\alpha} \leftarrow \{0, 1\}^\lambda$.
 - Computes kpABE secret keys $\{\text{kpABE.sk}_{\bar{L}, j, \bar{z}_j}\}_{j \in [\lambda]}$ and a SKFE secret key $\text{SKFE.s}\bar{\mathbf{k}}$ as in the construction.
 - Constructs $\text{CC}[\text{SKFE.s}\bar{\mathbf{k}}, \{\text{kpABE.sk}_{\bar{L}, j, \bar{z}_j}\}_{j \in [\lambda]}]$ and returns $\text{RMFE.c}\bar{\mathbf{t}} \leftarrow \text{LO.Obf}(\text{CC}[\text{SKFE.s}\bar{\mathbf{k}}, \{\text{kpABE.sk}_{\bar{L}, j, \bar{z}_j}\}_{j \in [\lambda]}], \bar{\alpha})$ to \mathcal{A} .
4. In the end, \mathcal{A} outputs a bit β' .

Hybrid₁ : This hybrid is same as the previous hybrid except the following:

1. The challenger samples $\alpha^*, \bar{\alpha}, \mathbf{z}^*, \bar{\mathbf{z}}$ and defines the functions $P_{f_0, \mathbf{z}^*, \alpha^*}$ and $P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}}$ in the beginning of the game.
2. The SKFE ciphertexts and keys are computed using SKFE simulators as:

²¹To keep the proofs and notations simple, we let f_0, f_1 and \bar{f} to be given selectively. This is sufficient to achieve security as in Def 5.2, as mentioned in Remark 5.4.

- For each key query (lb, x) , the SKFE ciphertext in $\text{RMFE.sk}_{\text{lb},x}$ is computed as $\text{SKFE.Sim}^{\text{CT}}(\text{SKFE.msk}, P_{f_0, \mathbf{z}^*, \alpha^*}(x, K, R), P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}}(x, K, R))$.
- The SKFE secret keys SKFE.sk^* and $\text{SKFE.s}\bar{\mathbf{k}}$ in RMFE.ct^* and $\text{RMFE.c}\bar{\mathbf{t}}$ are computed as $\text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{|P_{f_0, \mathbf{z}^*, \alpha^*}|})$ and $\text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{|P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}}|})$, respectively.

Hybrid₂ : In this hybrid, for each key query (lb, x) with $\text{lb} \in L^*$, for all $j \in [\lambda]$, $\text{kpABE.ct}_{\text{lb},j,1-\bar{z}_j}$ in $\text{RMFE.sk}_{\text{lb},x}$ is computed as $\text{kpABE.Enc}(\text{kpABE.mpk}, (\text{lb}, j, 1 - \bar{z}_j), 0^\lambda)$.

Hybrid₃ : In this hybrid, for each key query (lb, x) with $\text{lb} \in L^*$, K -values are chosen differently as follows: let $i \in [\lambda]$ be the first position where $z_i^* \neq \bar{z}_i$. If no such i exists, then the challenger aborts the game. Else, it samples $\{K_{j,b}\}_{j \in [\lambda] \setminus \{i\}, b \in \{0,1\}}$, K_{i,\bar{z}_i} , $\{R_{j,b}\}_{j \in [\lambda], b \in \{0,1\}}$ uniformly randomly as in the original game. It then sets $K_{i,1-\bar{z}_i} = \bigoplus_{j \in [\lambda] \setminus \{i\}} K_{j,z_j^*} \oplus \alpha^* \oplus \bigoplus_{j \in [\lambda]} R_{j,z_j^*}$.

Hybrid₄ : This hybrid is same as the previous hybrid, except that the SKFE ciphertexts and keys in RMFE secret keys and ciphertexts are now generated with function f_1 in place of f_0 as follows:

- For each key query (lb, x) , the SKFE ciphertext in $\text{RMFE.sk}_{\text{lb},x}$ is computed as $\text{SKFE.Sim}^{\text{CT}}(\text{SKFE.msk}, P_{f_1, \mathbf{z}^*, \alpha^*}(x, K, R), P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}}(x, K, R))$.
- The SKFE secret key SKFE.sk^* in RMFE.ct^* is computed as $\text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{|P_{f_1, \mathbf{z}^*, \alpha^*}|})$.

Now, from here we undo the changes made in the previous hybrids.

Hybrid₅ : In this hybrid, for each key query (lb, x) with $\text{lb} \in L^*$, $K_{i,1-\bar{z}_i}$ (where i is the first position where \mathbf{z}^* and $\bar{\mathbf{z}}$ differ) is also sampled randomly.

Hybrid₆ : In this hybrid, for each such key query (lb, x) with $\text{lb} \in L^*$, $\text{kpABE.Enc}(\text{kpABE.mpk}, (\text{lb}, j, 1-\bar{z}_j), 0^\lambda)$ is changed back to $\text{kpABE.Enc}(\text{kpABE.mpk}, (\text{lb}, j, 1-\bar{z}_j), K_{j,1-\bar{z}_j})$ in $\text{RMFE.sk}_{\text{lb},x}$.

Hybrid₇ : In this hybrid SKFE ciphertexts in RMFE secret keys and SKFE secret keys in RMFE ciphertexts are computed as in the real world. That is,

- For each key query (lb, x) , $\text{RMFE.sk}_{\text{lb},x} = (\text{SKFE.Enc}(\text{SKFE.msk}, K, R), \text{kpABE.mpk}, \{(\text{lb}, j, b), \text{kpABE.Enc}(\text{kpABE.mpk}, (\text{lb}, j, b), K_{j,b})\}_{j \in [\lambda], b \in \{0,1\}})$.
- $\text{RMFE.ct}^* = \text{LO.Obf}(\text{CC}[\text{SKFE.sk}^*, \{\text{kpABE.sk}_{L^*,j,z_j^*}\}_{j \in [\lambda]}], \alpha^*)$, $\text{RMFE.c}\bar{\mathbf{t}} = \text{LO.Obf}(\text{CC}[\text{SKFE.s}\bar{\mathbf{k}}, \{\text{kpABE.sk}_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}], \bar{\alpha})$, where $\text{SKFE.sk}^* = \text{SKFE.KeyGen}(\text{SKFE.msk}, P_{f_1, \mathbf{z}^*, \alpha^*})$ and $\text{SKFE.s}\bar{\mathbf{k}} = \text{SKFE.KeyGen}(\text{SKFE.msk}, P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}})$

Note that this hybrid corresponds to the real world for $\beta = 1$.

Indistinguishability of the hybrids. Now we show that the consecutive hybrids are computationally indistinguishable for any PPT adversary \mathcal{A} .

Claim 5.8.1. Assume that SKFE is secure (Def. 2.11). Then, $\text{Hybrid}_0 \approx_c \text{Hybrid}_1$.

Proof. The proof follows similar steps as the proof for Claim 5.7.1, hence omitted. \square

Claim 5.8.2. Assume that kpABE is selectively secure. Then, $\text{Hybrid}_1 \approx_c \text{Hybrid}_2$.

Proof. To prove the claim we consider the following sub hybrids between Hybrid_1 and Hybrid_2 . Let $L^* = \{\text{lb}_1, \dots, \text{lb}_{|L^*|}\}$ with some fixed ordering between the labels in L^* . Let $L_{[1:k]} \subseteq L^*$ denote the set of first k labels in L^* , i.e. $L_{[1:k]}^* = \{\text{lb}_1, \dots, \text{lb}_k\}$. Then for all $1 \leq i \leq |L^*|$ and $0 \leq \tau \leq \lambda$, define $\text{Hybrid}_{1,(i,\tau)}$: which is same as Hybrid_1 except that for any key query (lb, x) , $\{\text{kpABE.ct}_{\text{lb},j,b}\}_{j \in [\lambda], b \in \{0,1\}}$ is computed differently as follows:

$$\text{kpABE.ct}_{\text{lb},j,b} = \text{kpABE.Enc}(\text{kpABE.mpk}, (\text{lb}, j, b), W),$$

where

$$W = \begin{cases} K_{i,b} & \text{if } (b = \bar{z}_j) \vee (\text{lb} \notin L_{[1:i]}) \vee (\text{lb} = \text{lb}_i \wedge j > \tau), \\ 0^\lambda & \text{otherwise, i.e. } (b = 1 - \bar{z}_j) \wedge (\text{lb} \in L_{[1:i-1]}) \vee (\text{lb} = \text{lb}_i \wedge j \leq \tau). \end{cases}$$

Then, we observe that $\text{Hybrid}_{1,(1,0)} = \text{Hybrid}_1$, $\text{Hybrid}_{1,(|L^*|,\lambda)} = \text{Hybrid}_2$ and $\text{Hybrid}_{1,(i-1,\lambda)} = \text{Hybrid}_{1,(i,0)}$.

Hence, all we need to show is that for all $i \in [|L^*|]$, $\tau \in [\lambda]$,

$$\text{Hybrid}_{1,(i,\tau-1)} \approx_c \text{Hybrid}_{1,(i,\tau)}.$$

This follows from the Sel-IND security of kpABE. In particular, if there is no key query issued for lb_i , then the hybrids are identical. On the other hand, if there is a key query (lb, x) , such that $\text{lb} = \text{lb}_i$, then we show that if \mathcal{A} can distinguish between the two hybrids with non-negligible advantage ϵ then we can design a PPT algorithm \mathcal{B} with the same advantage ϵ against Sel-IND security of kpABE. \mathcal{B} is defined as follows:

1. Upon being invoked by the kpABE challenger, \mathcal{B} invokes \mathcal{A} which outputs f_0, f_1, L^*, \bar{f} . Let $L^* = \{\text{lb}_1, \dots, \text{lb}_{|L^*|}\}$.
2. \mathcal{B} sends $(\text{lb}_i, \tau, 1 - \bar{z}_\tau)$ as challenge attribute to the kpABE challenger. The kpABE challenger samples $\hat{\beta} \leftarrow \{0, 1\}$ and $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$ and sends kpABE.mpk to \mathcal{B} .
3. \mathcal{B} generates $\text{SKFE.msk} \leftarrow \text{SKFE.Setup}(1^\lambda)$ and sets $\text{RMFE.mpk} = \text{kpABE.mpk}$. It also samples $\alpha^*, \bar{\alpha}, \mathbf{z}^*, \bar{\mathbf{z}}$ and defines $P_{f_0, \mathbf{z}^*, \alpha^*}$ and $P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}}$. It then does the following:
 - For each $j \in [\lambda]$, defines the circuit $C_{L^*, j, \mathbf{z}_j^*}$ and sends a key query for $C_{L^*, j, \mathbf{z}_j^*}$ to the kpABE challenger. The kpABE challenger returns $\text{kpABE.sk}_{L^*, j, \mathbf{z}_j^*}$.
 - Computes $\text{SKFE.sk}^* \leftarrow \text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{|P_{f_0, \mathbf{z}^*, \alpha^*}|})$.
 - Computes $\text{RMFE.ct}^* \leftarrow \text{LO.Obf}(\text{CC}[\text{SKFE.sk}^*, \{\text{kpABE.sk}_{L^*, j, \mathbf{z}_j^*}\}_{j \in [\lambda]}])$.
 - Returns RMFE.mpk and RMFE.ct^* to \mathcal{A} .
4. **Key Queries:** For each key query (lb, x) , \mathcal{B} does the following:
 - Computes $\text{SKFE.ct} \leftarrow \text{SKFE.Sim}^{\text{CT}}(\text{SKFE.msk}, P_{f_0, \mathbf{z}^*, \alpha^*}(x, K, R), P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}}(x, K, R))$.
 - To compute kpABE ciphertext, $\text{kpABE.sk}_{\text{lb}, j, b}$ for $j \in [\lambda], b \in \{0, 1\}$,
 - if $(\text{lb} = \text{lb}_i) \wedge j = \tau \wedge b = 1 - \bar{z}_j$ ²², then \mathcal{B} sends challenge query with messages $\mu_0 = K_{j, 1 - \bar{z}_j}$ and $\mu_1 = 0^\lambda$. The kpABE challenger returns $\text{kpABE.ct} = \text{kpABE.Enc}(\text{kpABE.mpk}, (\text{lb}_i, \tau, 1 - \bar{z}_\tau), \mu_{\hat{\beta}})$, which \mathcal{B} sets as $\text{kpABE.ct}_{\text{lb}_i, \tau, 1 - \bar{z}_\tau}$.

²²If there are more than one key queries with $\text{lb} = \text{lb}_i$, then we use multi-challenge version of Sel-IND security of kpABE which can easily be shown equivalent to the one defined in Definition 2.14.

– else, \mathcal{B} computes $\text{kpABE.ct}_{\text{lb},j,b}$ on its own using kpABE.mpk as defined for $\text{Hybrid}_{1,(i,\tau-1)}$ (same for $\text{Hybrid}_{1,(i,\tau)}$).

- Returns $\text{RMFE.sk}_{\text{lb},x} = (\text{SKFE.ct}, \text{kpABE.mpk}, \{(\text{lb}, j, b), \text{kpABE.ct}_{\text{lb},j,b}\}_{j,b})$ to \mathcal{A} .

5. SK-Enc **Query**: When \mathcal{A} outputs \bar{L} as part of SK-Enc query, \mathcal{B} does the following:

- Computes $\text{SKFE.s}\bar{k} \leftarrow \text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{|P_{f,\bar{z},\bar{\alpha}}|})$.
- For each $j \in [\lambda]$, defines circuit $C_{\bar{L},j,\bar{z}_j}$ and sends a kpABE key query for $C_{\bar{L},j,\bar{z}_j}$. The kpABE challenger returns $\text{kpABE.sk}_{\bar{L},j,\bar{z}_j}$.
- Computes $\text{RMFE.c}\bar{t} \leftarrow \text{LO.Obf}(\text{CC}[\text{SKFE.s}\bar{k}, \{\text{kpABE.sk}_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}])$.
- Returns $\text{RMFE.c}\bar{t}$ to \mathcal{A} .

6. In the end, \mathcal{A} outputs a bit β' . \mathcal{B} forwards β' as its guess bit to the kpABE challenger.

We observe that if $\hat{\beta} = 0$, then \mathcal{B} simulated $\text{Hybrid}_{1,(i,\tau-1)}$, else $\text{Hybrid}_{1,(i,\tau)}$ with \mathcal{A} . Hence, advantage of $\mathcal{B} = |\Pr(\beta' = 1 | \hat{\beta} = 0) - \Pr(\beta' = 1 | \hat{\beta} = 1)| = |\Pr(\beta' = 1 | \text{Hybrid}_{1,(i,\tau-1)}) - \Pr(\beta' = 1 | \text{Hybrid}_{1,(i,\tau)})| = \epsilon$ (by assumption).

Admissibility of \mathcal{B} : Firstly, we observe that \mathcal{B} issues key queries for only the following set of circuits: $\{C_{L^*,j,z_j^*}\}_{j \in [\lambda]}$, $\{C_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}$ and the challenge attribute is $(\text{lb}_i, \tau, 1 - \bar{z}_\tau)$, where $\text{lb}_i \in L^*$. Next, we note that

- $C_{L^*,j,z_j^*}(\text{lb}_i, \tau, 1 - \bar{z}_\tau) = 0$ for all $j \in [\lambda]$, because $\text{lb}_i \in L^*$.
- $C_{\bar{L},j,\bar{z}_j}(\text{lb}_i, \tau, 1 - \bar{z}_\tau) = 0$ for all $j \neq \tau$, and $C_{\bar{L},\tau,\bar{z}_\tau}(\text{lb}_i, \tau, 1 - \bar{z}_\tau) = 0$, because $\bar{z}_\tau \neq 1 - \bar{z}_\tau$.

This establishes the admissibility of \mathcal{B} . □

Claim 5.8.3. *Hybrid₂ and Hybrid₃ are statistically indistinguishable.*

Proof. Firstly, we observe that $\Pr[\mathbf{z}^* = \bar{\mathbf{z}}] = 1/2^\lambda$. Hence, with probability $1 - 1/2^\lambda$, the challenger does not abort in Hybrid₃. Next we show that if the game is not aborted, the two hybrids are statistically indistinguishable in the view of the adversary. In this case, the only difference between the two hybrids is the following: for each key query (lb, x) with $\text{lb} \in L^*$, the value of $K_{i,1-\bar{z}_i}$ (i.e. K_{i,z_i^*}), where i is the first index such that $z_i^* \neq \bar{z}_i$, are computed differently. In Hybrid₂, K_{i,z_i^*} is sampled uniformly, while in Hybrid₃, it is computed as $K_{i,z_i^*} = \bigoplus_{j \in [\lambda] \setminus \{i\}} (K_{j,z_j^*} \oplus R_{j,z_j^*}) \oplus R_{i,z_i^*} \oplus \alpha^*$. However, note that if $f_0(x) = 1$, then K_{i,z_i^*} is not used anywhere because of the change we introduced in Hybrid₂ and hence does not affect the adversary's view. On the other hand, if $f_0(x) = 0$, R_{i,z_i^*} is not used anywhere, and hence K_{i,z_i^*} is uniformly random in the adversary's view because of the randomness of R_{i,z_i^*} . □

Claim 5.8.4. *Hybrid₃ and Hybrid₄ are identical in the view of the adversary.*

Proof. Observe that the two hybrids differ only in the computation (simulation) of SKFE ciphertexts and secret keys. In particular,

- $\text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{|P_{f_0,\mathbf{z}^*,\alpha^*}|})$ is changed to $\text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{|P_{f_1,\mathbf{z}^*,\alpha^*}|})$ in the computation of RMFE.ct^* . This is just a conceptual change and both the computations are exactly the same, since $|P_{f_0,\mathbf{z}^*,\alpha^*}| = |P_{f_1,\mathbf{z}^*,\alpha^*}|$.²³

²³we can use padding to make the sizes equal.

- For each key query (lb, x) , $\text{SKFE.Sim}^{\text{CT}}(\text{SKFE.msk}, P_{f_0, \mathbf{z}^*, \alpha^*}(x, K, R), P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}}(x, K, R))$ is changed to $\text{SKFE.Sim}^{\text{CT}}(\text{SKFE.msk}, P_{f_1, \mathbf{z}^*, \alpha^*}(x, K, R), P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}}(x, K, R))$. The change is again only conceptual and does not affect the actual computation as we argue below:
 - For $f_0(x) = f_1(x)$: there is no change.
 - For $f_0(x) \neq f_1(x)$: let $f_0(x) = 1$ and $f_1(x) = 0$. Then, $\text{lb} \in L^*$ and thus

$$\begin{aligned}
 P_{f_0, \mathbf{z}^*, \alpha^*}(x, K, R) &= \bigoplus_{j \in [\lambda]} R_{j, z_j^*} \\
 P_{f_1, \mathbf{z}^*, \alpha^*}(x, K, R) &= \bigoplus_{j \in [\lambda] \setminus \{i\}} K_{j, z_j^*} \oplus \alpha^* \oplus K_{i, z_i^*} \\
 &= \bigoplus_{j \in [\lambda] \setminus \{i\}} K_{j, z_j^*} \oplus \alpha^* \oplus \left(\bigoplus_{j \in [\lambda] \setminus \{i\}} (K_{j, z_j^*} \oplus R_{j, z_j^*}) \oplus R_{i, z_i^*} \right) \oplus \alpha^* \\
 &= \bigoplus_{j \in [\lambda]} R_{j, z_j^*}.
 \end{aligned}$$

- The same argument works for $f_0(x) = 0$ and $f_1(x) = 1$.

□

Indistinguishability between the rest of the hybrids can be argued in the same way as their counterparts in the previous set of hybrids. In particular, proofs for indistinguishability between Hybrid_4 and Hybrid_5 is same as the proof of claim 5.8.3, Hybrid_5 and Hybrid_6 is same as the proof for claim 5.8.2 and Hybrid_6 and Hybrid_7 is same as the proof for claim 5.8.1. □

6 Secret Key RPE from Evasive and Tensor LWE

In this section we construct a secret-key RPE scheme from evasive and tensor LWE, followed by the efficiency and security analysis of our scheme.

6.1 Construction

We give a construction of the secret-key RPE scheme $\text{RPE} = (\text{RPE.Setup}, \text{RPE.KeyGen}, \text{RPE.Broadcast}, \text{RPE.Enc}, \text{RPE.Dec})$ for an attribute space $\mathcal{X} = \{\mathcal{X}_\lambda\}_\lambda$, a function family $\mathcal{F} = \{\mathcal{F}_\lambda\}_\lambda$ where $\mathcal{F}_\lambda = \{f : \mathcal{X}_\lambda \rightarrow \{0, 1\}\}$, a label space $\mathcal{L} = \{\mathcal{L}_\lambda\}_\lambda$ and a message space $\mathcal{M} = \{\mathcal{M}_\lambda\}_\lambda$. We assume that the size of $|\mathcal{F}_\lambda|$ is bounded by some polynomial in λ , which will suffice for our purpose. Our scheme uses the following building blocks:

1. An RMFE scheme $\text{RMFE} = (\text{RMFE.Setup}, \text{RMFE.KeyGen}, \text{RMFE.PK-Enc}, \text{RMFE.SK-Enc}, \text{RMFE.Dec})$ with attribute space \mathcal{X} , label space \mathcal{L} and function family \mathcal{F} . We instantiate it by our construction in Sec. 5.2 based on LWE.
2. A CP-ABE scheme $\text{cpABE} = (\text{cpABE.Setup}, \text{cpABE.Enc}, \text{cpABE.KeyGen}, \text{cpABE.Dec})$ with message space \mathcal{M} satisfying Sel-IND security (Def. 2.14) that supports the circuit class $\mathcal{C}_{\ell(\lambda), d(\lambda)}$ and the efficiency property listed in Theorem 2.18. We set $\ell(\lambda) = |\text{RMFE.sk}_{\text{lb}, x}|$ and $d(\lambda)$ to be the upper bound on the depth of the circuit $C_{L, \text{RMFE.ct}}$ in Eq. (6.1). Looking ahead, we show that the depth is bounded by some polynomial $\text{poly}(\lambda, |f|, |x|, |\text{lb}|)$. We choose d so that it is larger than the value. We can instantiate the scheme by the one proposed by [Wee22] based on tensor and evasive LWE.

We describe our construction below.

$\text{RPE.Setup}(1^\lambda) \rightarrow (\text{RPE.mpk}, \text{RPE.msk})$. The setup algorithm takes as input the security parameter λ and does the following:

- Generates $(\text{RMFE.mpk}, \text{RMFE.msk}) \leftarrow \text{RMFE.Setup}(1^\lambda)$ and $(\text{cpABE.mpk}, \text{cpABE.msk}) \leftarrow \text{cpABE.Setup}(1^\lambda)$.
- Outputs $\text{RPE.mpk} = (\text{RMFE.mpk}, \text{cpABE.mpk})$ and $\text{RPE.msk} = (\text{RMFE.msk}, \text{cpABE.msk})$.

$\text{RPE.KeyGen}(\text{RPE.msk}, \text{lb}, x) \rightarrow \text{RPE.sk}_{\text{lb},x}$. The key generation algorithm takes as input the master secret key RPE.msk , a label $\text{lb} \in \mathcal{L}$ and an attribute $x \in \mathcal{X}$ and does the following:

- Parse $\text{RPE.msk} = (\text{RMFE.msk}, \text{cpABE.msk})$.
- Compute $\text{RMFE.sk}_{\text{lb},x} \leftarrow \text{RMFE.KeyGen}(\text{RMFE.msk}, \text{lb}, x)$.
- Set $\text{att} = ((\text{lb}, x), \text{RMFE.sk}_{\text{lb},x})$.
- Compute $\text{cpABE.sk}_{\text{lb},x} \leftarrow \text{cpABE.KeyGen}(\text{cpABE.msk}, \text{att})$.
- Output $\text{RPE.sk}_{\text{lb},x} = (\text{att}, \text{cpABE.sk}_{\text{lb},x})$.

$\text{RPE.Broadcast}(\text{RPE.mpk}, m, L) \rightarrow \text{RPE.ct}$. The broadcast algorithm takes as input the master public key RPE.mpk , a message m , and a revocation list $L \subseteq \mathcal{L}$ and does the following:

- Parse $\text{RPE.mpk} = (\text{RMFE.mpk}, \text{cpABE.mpk})$.
- Compute $\text{RMFE.ct} \leftarrow \text{RMFE.PK-Enc}(\text{RMFE.mpk}, L)$.
- Construct circuit $C_{L, \text{RMFE.ct}}$ with L and RMFE.ct hardwired, defined as follows:
On input (lb, x) and $\text{RMFE.sk}_{\text{lb},x}$,

$$\begin{aligned} C_{L, \text{RMFE.ct}}((\text{lb}, x), \text{RMFE.sk}_{\text{lb},x}) \\ = (\text{lb} \notin L) \wedge (\text{RMFE.Dec}(\text{RMFE.sk}_{\text{lb},x}, \text{RMFE.ct}, L) = 1). \end{aligned} \quad (6.1)$$

- Compute $\text{cpABE.ct} \leftarrow \text{cpABE.Enc}(\text{cpABE.mpk}, C_{L, \text{RMFE.ct}}, m)$.
- Output $\text{RPE.ct} = (\text{RMFE.ct}, \text{cpABE.ct})$.

$\text{RPE.Enc}(\text{RPE.msk}, f, m, L) \rightarrow \text{RPE.ct}$. The encryption algorithm takes as input the master secret key, a function f , a message m , and a revocation list $L \subseteq \mathcal{L}$ and does the following:

- Parse $\text{RPE.msk} = (\text{RMFE.msk}, \text{cpABE.msk})$.
- Compute $\text{RMFE.ct} \leftarrow \text{RMFE.SK-Enc}(\text{RMFE.msk}, f, L)$.
- Construct circuit $C_{L, \text{RMFE.ct}}$ from L and RMFE.ct as defined in Eq. (6.1).
- Compute $\text{cpABE.ct} \leftarrow \text{cpABE.Enc}(\text{cpABE.mpk}, C_{L, \text{RMFE.ct}}, m)$.
- Output $\text{RPE.ct} = (\text{RMFE.ct}, \text{cpABE.ct})$.

$\text{RPE.Dec}(\text{RPE.sk}_{\text{lb},x}, \text{RPE.ct}, L) \rightarrow m$. The decryption algorithm takes as input the secret key $\text{RPE.sk}_{\text{lb},x}$, a ciphertext RPE.ct , and a revocation list L and does the following:

- Parse $\text{RPE.sk}_{\text{lb},x} = (\text{att} = ((\text{lb}, x), \text{RMFE.sk}_{\text{lb},x}), \text{cpABE.sk}_{\text{lb},x})$ and $\text{RPE.ct} = (\text{RMFE.ct}, \text{cpABE.ct})$.
- Construct circuit $C_{L, \text{RMFE.ct}}$ from L and RMFE.ct as defined in Eq. (6.1).
- Compute and output $\text{cpABE.Dec}(\text{cpABE.mpk}, \text{cpABE.sk}_{\text{lb},x}, \text{att}, \text{cpABE.ct}, C_{L, \text{RMFE.ct}})$.

Correctness. First, we show that $C_{L, \text{RMFE.ct}} \in \mathcal{C}_{\ell, d}$. In particular, it suffices to bound the depth of the circuit d by some fixed polynomial $\text{poly}(\lambda)$. We first observe that checking whether $\text{lb} \in L$ or not can be done with depth $|\log(|\text{lb}| \cdot |L|)| = \log(\text{poly}(\lambda)) \leq \lambda$. We also have that the depth of RMFE.Dec is bounded by $\text{poly}(\lambda, |f|, |x|, |\text{lb}|)$ as we saw in Sec. 5.2. Therefore, the total depth is bounded by $\text{poly}(\lambda, |f|, |x|, |\text{lb}|)$.

Theorem 6.1. *Suppose RMFE and cpABE are correct, then the above construction of secret-key RPE satisfies correctness (Def. 3.2).*

Proof. For any (lb, x) , function $f \in \mathcal{F}_\lambda$ and a revocation list $L \subseteq \mathcal{L}_\lambda$ such that $f(x) = 1, \forall x \in \mathcal{X}_\lambda$ and $\text{lb} \notin L$, consider the following two cases:

1. **Broadcast Correctness:** For any ciphertext $\text{RPE.ct} \leftarrow \text{RPE.Broadcast}(\text{RPE.mpk}, m, L)$, we have $\text{RPE.ct} = (\text{RMFE.ct}, \text{cpABE.ct})$, where $\text{RMFE.ct} \leftarrow \text{RMFE.PK-Enc}(\text{RMFE.mpk}, L)$. So, by the public encryption correctness of RMFE scheme, with all but negligible probability

$$\text{RMFE.Dec}(\text{RMFE.sk}_{\text{lb}, x}, \text{RMFE.ct}, L) = 1.$$

Since $\text{lb} \notin L$, we have $C_{L, \text{RMFE.ct}}(\text{att}) = 1$, where $\text{att} = ((\text{lb}, x), \text{RMFE.sk}_{\text{lb}, x})$. Hence by the correctness of cpABE scheme we have that with all but negligible probability

$$\text{cpABE.Dec}(\text{cpABE.mpk}, \text{cpABE.sk}_{\text{lb}, x}, \text{att}, \text{cpABE.ct}, C_{L, \text{RMFE.ct}}) = m.$$

2. **Encryption Correctness:** For any ciphertext $\text{RPE.ct} \leftarrow \text{RPE.Enc}(\text{RPE.msk}, f, m, L)$, we have $\text{RPE.ct} = (\text{RMFE.ct}, \text{cpABE.ct})$, where $\text{RMFE.ct} \leftarrow \text{RMFE.SK-Enc}(\text{RMFE.msk}, f, L)$. If $(\text{lb} \notin L) \wedge f(x) = 1$, then by the correctness of RMFE.SK-Enc algorithm, we have with all but negligible probability

$$\text{RMFE.Dec}(\text{RMFE.sk}_{\text{lb}, x}, \text{RMFE.ct}, L) = 1.$$

Furthermore, $\text{lb} \notin L$ implies $C_{L, \text{RMFE.ct}}((\text{lb}, x), \text{RMFE.sk}_{\text{lb}, x}) = 1$. Hence by the correctness of cpABE scheme, we have that with all but negligible probability

$$\text{cpABE.Dec}(\text{cpABE.mpk}, \text{cpABE.sk}_{\text{lb}, x}, \text{att}, \text{cpABE.ct}, C_{L, \text{RMFE.ct}}) = m.$$

Hence the above construction of secret-key RPE satisfies correctness. \square

Efficiency. Here we argue that our construction achieves optimal parameters. Namely, we show that all the parameters are independent from $|L|$.

1. **Public key size $|\text{RPE.mpk}|$:** We have $|\text{RPE.mpk}| = |\text{cpABE.mpk}| + |\text{RMFE.mpk}|$. The former is bounded by $\text{poly}(\lambda, \ell, d) = \text{poly}(\lambda, |f|, |x|, |\text{lb}|)$ using Theorem 2.18, where we additionally used $\ell = |\text{RMFE.sk}_{\text{lb}, x}| = \text{poly}(\lambda, |f|, |x|, |\text{lb}|)$ and $d = \text{poly}(\lambda, |f|, |x|, |\text{lb}|)$.
2. **Secret key size $|\text{RPE.sk}_{\text{lb}, x}|$:** We have $|\text{RPE.sk}_{\text{lb}, x}| = |\text{att}| + |\text{cpABE.sk}_{\text{lb}, x}|$. We have $|\text{att}| = |\text{lb}| + |x| + |\text{RMFE.sk}_{\text{lb}, x}| = \text{poly}(\lambda, |f|, |x|, |\text{lb}|)$ and $|\text{cpABE.sk}_{\text{lb}, x}| = \text{poly}(\ell, d) = \text{poly}(\lambda, |f|, |x|, |\text{lb}|)$ by Theorem 2.18. Therefore, the overall size is $\text{poly}(\lambda, |f|, |x|, |\text{lb}|)$.
3. **Ciphertext size $|\text{RPE.ct}|$:** We have $|\text{RPE.ct}| = |\text{RMFE.ct}| + |\text{cpABE.ct}|$. The former is bounded by $\text{poly}(\lambda, |f|, |x|, |\text{lb}|)$ and the latter is bounded by $\text{poly}(\lambda, d) = \text{poly}(\lambda, |f|, |x|, |\text{lb}|)$ by Theorem 2.18. Therefore, the overall size is $\text{poly}(\lambda, |f|, |x|, |\text{lb}|)$.

6.2 Security

In this section we show that our construction of secret-key RPE is secure.

Message Hiding Security

Theorem 6.2. *Assume that cpABE is selectively secure (Def. 2.14), RMFE is correct and $|\mathcal{F}| \leq \text{poly}(\lambda)$. Then RPE scheme satisfies 1-query selective message hiding security (Def. 3.4).*

Proof. Recall that in the message hiding security game, we want

$$\text{RPE.Enc}(\text{RPE.msk}, f, m_0, L) \approx_c \text{RPE.Enc}(\text{RPE.msk}, f, m_1, L),$$

where for all the key queries (lb, x) to the $\text{RPE.KeyGen}(\text{RPE.msk}, \cdot, \cdot)$ oracle, either $f(x) = 0$ or $\text{lb} \in L$. We show that if there exists an adversary \mathcal{A} who has non-negligible advantage ϵ in the selective message hiding security game, then there exists a PPT adversary \mathcal{B} against the security of cpABE scheme with the same advantage ϵ . The reduction is as follows.

1. \mathcal{B} first runs \mathcal{A} and gets f and L ²⁴.
2. \mathcal{B} generates $(\text{RMFE.mpk}, \text{RMFE.msk}) \leftarrow \text{RMFE.Setup}(1^\lambda)$ and computes $\text{RMFE.ct} \leftarrow \text{RMFE.SK-Enc}(\text{RMFE.msk}, f, L)$.
3. \mathcal{B} constructs the circuit $C_{L, \text{RMFE.ct}}$ as defined in the construction and sends it as the challenge function to the cpABE challenger. The cpABE challenger generates $(\text{cpABE.mpk}, \text{cpABE.msk}) \leftarrow \text{cpABE.Setup}(1^\lambda)$ and returns cpABE.mpk to \mathcal{B} . \mathcal{B} sets $\text{RPE.mpk} = (\text{RMFE.mpk}, \text{cpABE.mpk})$ and sends it to \mathcal{A} .
4. **Key Queries:** For each key query (lb, x) , \mathcal{B} does the following:
 - Computes $\text{RMFE.sk}_{\text{lb}, x} \leftarrow \text{RMFE.KeyGen}(\text{RMFE.msk}, \text{lb}, x)$.
 - Sets $\text{att} = ((\text{lb}, x), \text{RMFE.sk}_{\text{lb}, x})$, sends it as a key query to cpABE challenger and gets back $\text{cpABE.sk}_{\text{lb}, x}$.
 - It returns $\text{RPE.sk}_{\text{lb}, x} = (\text{att}, \text{cpABE.sk}_{\text{lb}, x})$ to \mathcal{A} .
5. **Challenge Query:** When \mathcal{A} sends the challenge messages (m_0, m_1) , \mathcal{B} forwards it to the cpABE challenger. The cpABE challenger samples a bit $\hat{\beta} \leftarrow \{0, 1\}$ and returns $\text{cpABE.ct}_{\hat{\beta}} \leftarrow \text{cpABE.Enc}(\text{cpABE.mpk}, C_{L, \text{RMFE.ct}}, m_{\hat{\beta}})$ to \mathcal{B} . \mathcal{B} sends $\text{RPE.ct} = (\text{RMFE.ct}, \text{cpABE.ct}_{\hat{\beta}})$ to \mathcal{A} .
6. **Encryption Query:** When \mathcal{A} makes the encryption query $(\bar{f}, \bar{m}, \bar{L})$, \mathcal{B} does the following:
 - Computes $\text{RMFE.ct} \leftarrow \text{RMFE.SK-Enc}(\text{RMFE.msk}, \bar{f}, \bar{L})$ and constructs the circuit $C_{\bar{L}, \text{RMFE.ct}}$ as defined in the construction.
 - Computes $\text{cpABE.ct} \leftarrow \text{cpABE.Enc}(\text{cpABE.mpk}, C_{\bar{L}, \text{RMFE.ct}}, \bar{m})$.
 - Returns $\text{RPE.ct} = (\text{RMFE.ct}, \text{cpABE.ct})$ to \mathcal{A} .
7. In the end, \mathcal{A} outputs a bit β' .

²⁴To keep the proofs simple, we let f to be given selectively. This is sufficient to achieve security as in Def 3.4, as mentioned in Remark 3.9.

Observe that if the cpABE challenger chose $\hat{\beta} = 0$, then \mathcal{B} simulated the real world where m_0 was encrypted, else it simulated the real world where m_1 was encrypted with \mathcal{A} .

Hence, advantage of \mathcal{B} is $|\Pr[\beta' = 1 | \hat{\beta} = 0] - \Pr[\beta' = 1 | \hat{\beta} = 1]| = |\Pr[\beta' = 1 | \text{RPE.ct} = \text{RPE.Enc}(\text{RPE.msk}, f, m_0, L)] - \Pr[\beta' = 1 | \text{RPE.ct} = \text{RPE.Enc}(\text{RPE.msk}, f, m_1, L)]| = \epsilon$ (by assumption).

Admissibility of \mathcal{B} . We observe that for the challenge circuit $C_{L, \text{RMFE.ct}}$ and for all key queries (lb, x) , queried by \mathcal{B} , we have $C_{L, \text{RMFE.ct}}((\text{lb}, x), \text{RMFE.sk}_{|\text{b}, x}) = 0$ as either (i) $\text{lb} \in L$, or (ii) $\text{lb} \in L$ and $f(x) = 0$, which implies $\text{RMFE.Dec}(\text{RMFE.sk}_{|\text{b}, x}, \text{RMFE.ct}) = 0$ (due to RMFE correctness), by the admissibility condition on \mathcal{A} . So, if \mathcal{A} is admissible then so is \mathcal{B} . \square

Function Hiding Security

Theorem 6.3. *Assume RMFE satisfies 1-query selective function hiding security (Def. 5.3), then the RPE scheme satisfies 1-query selective function hiding security (Def. 3.6).*

Proof. Recall that in the function hiding security game, we want

$$\text{RPE.Enc}(\text{RPE.msk}, f_0, m, L) \approx_c \text{RPE.Enc}(\text{RPE.msk}, f_1, m, L),$$

where for all the key queries (lb, x) to the $\text{RPE.KeyGen}(\text{RPE.msk}, \cdot, \cdot)$ oracle, either $f_0(x) = f_1(x)$ or $\text{lb} \in L$.

We show that if there exists an adversary \mathcal{A} who has non-negligible advantage ϵ in the 1-query selective function hiding security game, then there exists a PPT adversary \mathcal{B} against the 1-query selective function-hiding security of RMFE scheme with the same advantage ϵ . The reduction is as follows.

1. \mathcal{B} runs \mathcal{A} and gets the revocation list L .
2. \mathcal{B} sends L as the challenge revocation list to the RMFE challenger. The challenger generates $(\text{RMFE.mpk}, \text{RMFE.msk})$ and returns RMFE.mpk to \mathcal{B} .
3. \mathcal{B} generates $(\text{cpABE.mpk}, \text{cpABE.msk}) \leftarrow \text{cpABE.Setup}(1^\lambda)$. It sets $\text{RPE.mpk} = (\text{RMFE.mpk}, \text{cpABE.mpk})$ and sends it to \mathcal{A} .
4. **Key Queries:** On each key query (lb, x) , \mathcal{B} does the following:
 - It sends a key query (lb, x) to the RMFE challenger and gets back $\text{RMFE.sk}_{|\text{b}, x}$.
 - Sets $\text{att} = ((\text{lb}, x), \text{RMFE.sk}_{|\text{b}, x})$ and computes $\text{cpABE.sk}_{|\text{b}, x} \leftarrow \text{cpABE.KeyGen}(\text{cpABE.msk}, \text{att})$.
 - It returns $\text{RPE.sk}_{|\text{b}, x} = (\text{att}, \text{cpABE.sk}_{|\text{b}, x})$ to \mathcal{A} .
5. **Challenge Query:** When \mathcal{A} sends the challenge functions (f_0, f_1) and message m , \mathcal{B} does the following:
 - Sends (f_0, f_1) as the challenge functions to the RMFE challenger. The challenger samples $\hat{\beta} \leftarrow \{0, 1\}$, computes $\text{RMFE.ct}_{\hat{\beta}} \leftarrow \text{RMFE.SK-Enc}(\text{RMFE.msk}, f_{\hat{\beta}}, L)$ and returns $\text{RMFE.ct}_{\hat{\beta}}$ to \mathcal{B} .
 - Constructs the circuit $C_{L, \text{RMFE.ct}_{\hat{\beta}}}$ as defined in the construction and computes $\text{cpABE.ct} \leftarrow \text{cpABE.Enc}(\text{cpABE.mpk}, C_{L, \text{RMFE.ct}_{\hat{\beta}}}, m)$.
 - Returns $\text{RPE.ct} = (\text{RMFE.ct}_{\hat{\beta}}, \text{cpABE.ct})$ to \mathcal{A} .

6. **Encryption Query:** When \mathcal{A} makes an encryption query $(\bar{f}, \bar{m}, \bar{L})$, \mathcal{B} does the following:

- Sends a SK-Enc query (\bar{f}, \bar{L}) to the RMFE challenger and gets back RMFE.ct .
- Constructs the circuit $C_{\bar{L}, \text{RMFE.ct}}$ as defined in the construction.
- Computes $\text{cpABE.ct} \leftarrow \text{cpABE.Enc}(\text{cpABE.mpk}, C_{\bar{L}, \text{RMFE.ct}}, \bar{m})$.
- Returns $\text{RPE.ct} = (\text{RMFE.ct}, \text{cpABE.ct})$ to \mathcal{A} .

7. In the end, the adversary outputs a bit β' .

Observe that if the RMFE challenger chose $\hat{\beta} = 0$, then \mathcal{B} simulated the real world where f_0 was encrypted, else it simulated the real world where f_1 was encrypted with \mathcal{A} .

Hence, advantage of \mathcal{B} is $|\Pr[\beta' = 1 | \hat{\beta} = 0] - \Pr[\beta' = 1 | \hat{\beta} = 1]| = |\Pr[\beta' = 1 | \text{RPE.ct} = \text{RPE.Enc}(\text{RPE.msk}, f_0, m, L)] - \Pr[\beta' = 1 | \text{RPE.ct} = \text{RPE.Enc}(\text{RPE.msk}, f_1, m, L)]| = \epsilon$ (by assumption).

Admissibility of \mathcal{B} . First, we note that since \mathcal{A} is allowed to make only one query, $(\bar{f}, \bar{m}, \bar{L})$ to the $\text{RPE.Enc}(\text{msk}, \cdot, \cdot, \cdot)$ oracle, \mathcal{B} also makes only one query, (\bar{f}, \bar{L}) , to the $\text{RMFE.SK-Enc}(\text{msk}, \cdot, \cdot)$ oracle. Next, we observe that since \mathcal{A} is restricted to make key queries (lb, x) such that either $f_0(x) = f_1(x)$ or $\text{lb} \in L$, thus \mathcal{B} also issues key queries (lb, x) to RMFE challenger such that either $f_0(x) = f_1(x)$ or $\text{lb} \in L$. Hence, if \mathcal{A} is admissible, then so is \mathcal{B} . \square

Broadcast Security

Theorem 6.4. *Assume RMFE satisfies 1-query mode hiding security (Def. 5.2), then the RPE scheme satisfies 1-query selective broadcast security (Def. 3.7).*

Proof. Recall that in the broadcast security game, we want

$$\text{RPE.Enc}(\text{RPE.msk}, f, m, L) \approx_c \text{RPE.Broadcast}(\text{RPE.mpk}, m, L),$$

where $f(x) = 1, \forall x \in \mathcal{X}$.

We show that if there exists an adversary \mathcal{A} who has non-negligible advantage ϵ in the broadcast security game, then there exists a PPT adversary \mathcal{B} against the mode-hiding security of RMFE scheme with the same advantage ϵ . The reduction is as follows.

1. \mathcal{B} runs \mathcal{A} and gets the revocation list L .
2. \mathcal{B} sends L as the challenge revocation list to the RMFE challenger. The challenger generates $(\text{RMFE.mpk}, \text{RMFE.msk})$ and returns RMFE.mpk to \mathcal{B} .
3. \mathcal{B} generates $(\text{cpABE.mpk}, \text{cpABE.msk}) \leftarrow \text{cpABE.Setup}(1^\lambda)$. It sets $\text{RPE.mpk} = (\text{RMFE.mpk}, \text{cpABE.mpk})$ and sends it to \mathcal{A} .
4. **Key Queries:** On each key query (lb, x) , \mathcal{B} does the following:
 - It sends a key query (lb, x) to the RMFE challenger and gets back $\text{RMFE.sk}_{\text{lb}, x}$.
 - Sets $\text{att} = ((\text{lb}, x), \text{RMFE.sk}_{\text{lb}, x})$ and computes $\text{cpABE.sk}_{\text{lb}, x} \leftarrow \text{cpABE.KeyGen}(\text{cpABE.msk}, \text{att})$.
 - It returns $\text{RPE.sk}_{\text{lb}, x} = (\text{att}, \text{cpABE.sk}_{\text{lb}, x})$ to \mathcal{A} .
5. **Challenge Query:** When \mathcal{A} sends the challenge function f and message m , \mathcal{B} does the following:

- Sends f as the challenge function to the RMFE challenger. The challenger samples $\hat{\beta} \leftarrow \{0, 1\}$, and returns $\text{RMFE.ct}_{\hat{\beta}}$ to \mathcal{B} , where $\text{RMFE.ct}_0 \leftarrow \text{RMFE.SK-Enc}(\text{RMFE.msk}, f, L)$ and $\text{RMFE.ct}_1 \leftarrow \text{RMFE.PK-Enc}(\text{RMFE.mpk}, L)$.
- Constructs the circuit $C_{L, \text{RMFE.ct}_{\hat{\beta}}}$ as defined in the construction and computes $\text{cpABE.ct} \leftarrow \text{cpABE.Enc}(\text{cpABE.mpk}, C_{L, \text{RMFE.ct}_{\hat{\beta}}}, m)$.
- Returns $\text{RPE.ct} = (\text{RMFE.ct}_{\hat{\beta}}, \text{cpABE.ct})$ to \mathcal{A} .

6. **Encryption Query:** When \mathcal{A} makes an encryption query $(\bar{f}, \bar{m}, \bar{L})$, \mathcal{B} does the following:

- Sends a SK-Enc query (\bar{f}, \bar{L}) to the RMFE challenger and gets back RMFE.ct .
- Constructs the circuit $C_{\bar{L}, \text{RMFE.ct}}$ as defined in the construction.
- Computes $\text{cpABE.ct} \leftarrow \text{cpABE.Enc}(\text{cpABE.mpk}, C_{\bar{L}, \text{RMFE.ct}}, \bar{m})$.
- Returns $\text{RPE.ct} = (\text{RMFE.ct}, \text{cpABE.ct})$ to \mathcal{A} .

7. In the end, the adversary outputs a bit β' .

Observe that if the RMFE challenger chose $\hat{\beta} = 0$, then \mathcal{B} simulated the real world where the ciphertext was computed using Enc algorithm else it simulated the real world where the ciphertext was computed using Broadcast algorithm, with \mathcal{A} .

Hence, the advantage of \mathcal{B} is $|\Pr[\beta' = 1 | \hat{\beta} = 0] - \Pr[\beta' = 1 | \hat{\beta} = 1]| = |\Pr[\beta' = 1 | \text{RPE.ct} = \text{RPE.Enc}(\text{RPE.msk}, f, m, L)] - \Pr[\beta' = 1 | \text{RPE.ct} = \text{RPE.Broadcast}(\text{RPE.msk}, m, L)]| = \epsilon$ (by assumption).

Admissibility of \mathcal{B} . First, we note that since \mathcal{A} is allowed to make only one query, $(\bar{f}, \bar{m}, \bar{L})$ to the $\text{RPE.Enc}(\text{msk}, \cdot, \cdot, \cdot)$ oracle, \mathcal{B} also makes only one query, (f, \bar{L}) , to the $\text{RMFE.SK-Enc}(\text{msk}, \cdot, \cdot)$ oracle. Next, we observe that since \mathcal{A} is restricted to output f , such that $f(x) = 1$ for all $x \in \mathcal{X}$, this implies \mathcal{B} also issues such f as the challenge query to the RMFE challenger in the above mode hiding security game. Thus, if \mathcal{A} is admissible, then so is \mathcal{B} . \square

7 Embedded Identity Trace and Revoke

In this section, we define different variants of an embedded identity trace and revoke system (EITR). Our definitions extend the different notions that Goyal et al. [GKW19] introduced, in the context of embedded identity traitor tracing, to incorporate the revocation list. Concretely, we define three variants of an EITR scheme: 1) Indexed EITR, 2) Bounded EITR, and 3) Unbounded EITR. Our goal is Unbounded EITR and other variants. Other variants are introduced as intermediate goals. We show a construction of EITR from RPE in Sec. 8, bounded EITR from indexed EITR in Sec. 9, and unbounded EITR from bounded EITR in Sec. 10. These implications hold for both secret key and public key settings.

An EITR scheme consists of five polynomial time algorithms—Setup, KeyGen, Enc, Dec and Trace. The syntax of the EITR variants mentioned above differs only in the inputs to the Setup, KeyGen and Trace algorithms. Here we give a unified definition and later specify the distinctness of the three variants.

Consider a general identity space \mathcal{GID} , a label space \mathcal{L} and a message space \mathcal{M} . An embedded identity trace and revoke scheme $\text{EITR} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec}, \text{Trace})$ has the following syntax:

$\text{Setup}(1^\lambda, \text{params}_1) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm takes as input the security parameter λ and parameters params_1 . It outputs a master public key mpk and a master secret key msk .

$\text{KeyGen}(\text{msk}, \text{lb}, \text{gid}) \rightarrow \text{sk}_{\text{lb}, \text{gid}}$. The key generation algorithm takes as input the master secret key msk , a label $\text{lb} \in \mathcal{L}$, and a general identity $\text{gid} \in \mathcal{GID}$. It outputs a secret key $\text{sk}_{\text{lb}, \text{gid}}$.

$\text{Enc}(\text{mpk}, m, L) \rightarrow \text{ct}$. The encryption algorithm takes as input the master public key mpk , a message $m \in \mathcal{M}$, a revocation list $L \subseteq \mathcal{L}$ and outputs a ciphertext ct .

$\text{Dec}(\text{sk}_{\text{lb}, \text{gid}}, \text{ct}, L) \rightarrow y$. The decryption algorithm takes as input a secret key $\text{sk}_{\text{lb}, \text{gid}}$, a ciphertext ct , and a revocation list L and outputs $y \in \mathcal{M} \cup \{\perp\}$.

$\text{Trace}^D(\text{tk}, \text{params}_2, m_0, m_1, L) \rightarrow T$. The tracing algorithm takes as input a tracing key tk , parameter params_2 , two messages m_0, m_1 , a revocation list L and has an oracle access to a decoder D . It outputs a set of traitors T .

The above syntax captures both public key and secret key trace EITR schemes. For any *public tracing* EITR scheme, we have $\text{tk} = \text{mpk}$ and in the *secret tracing* EITR scheme, $\text{tk} = \text{msk}$. We now describe the properties satisfied by an EITR scheme.

Definition 7.1 (Correctness). An EITR scheme is said to be correct if there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, any label $\text{lb} \in \mathcal{L}$, and any revocation list $L \subseteq \mathcal{L}$ such that $\text{lb} \notin L$, the following holds

$$\Pr \left[\begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \text{params}_1); \\ \text{Dec}(\text{sk}_{\text{lb}, \text{gid}}, \text{ct}, L) = m : \text{sk}_{\text{lb}, \text{gid}} \leftarrow \text{KeyGen}(\text{msk}, \text{lb}, \text{gid}); \\ \text{ct} \leftarrow \text{Enc}(\text{mpk}, m, L) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

Definition 7.2 (IND-CPA Security). An EITR scheme is said to be IND-CPA secure if for every stateful PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for every $\lambda \in \mathbb{N}$, the following holds

$$\Pr \left[\begin{array}{l} \text{params}_1 \leftarrow \mathcal{A}(1^\lambda); \\ \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{ct}_\beta) = \beta : \begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \text{params}_1); \\ (m_0, m_1, L) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{mpk}); \\ \beta \leftarrow \{0, 1\}; \text{ct}_\beta \leftarrow \text{Enc}(\text{mpk}, m_b, L) \end{array} \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda)$$

where the adversary has the access to the $\text{KeyGen}(\text{msk}, \cdot, \cdot)$ oracle which has msk hardwired and \mathcal{A} is admissible only if for all the key generation queries (lb, gid) to the KeyGen oracle, $\text{lb} \in L$.

Definition 7.3 (Selective IND-CPA Security). The selective IND-CPA security of an EITR scheme is defined in the same way as Def. 7.2, except that the adversary outputs the revocation list L along with params_1 before the Setup algorithm is run.

In the following, we specify the space \mathcal{GID} , inputs $\text{params}_1, \text{gid}$ and params_2 to the Setup, KeyGen and Trace algorithm, respectively, and then define the secure tracing guarantee of each of the EITR notions separately. We let $\mathcal{ID} = \{0, 1\}^k$ denote the identity space.

Remark 7.4. We assume that there exists an efficiently computable mapping $\text{map} : \mathcal{GID} \rightarrow \mathcal{L}$, that *uniquely* maps an $\text{gid} \in \mathcal{GID}$ to a label $\text{lb} \in \mathcal{L}$. This can be easily ensured, for e.g. by making label lb a part of id . We further note that in real world applications one may also want to ensure that any label lb is associated with at most one id . This can be achieved by using a collision resistant hash function.

Experiment $\text{Expt-Ind-TR}_{\mathcal{A},\epsilon}(\lambda)$

- $1^\kappa, 1^{n_{\text{ind}}} \leftarrow \mathcal{A}(1^\lambda)$
- $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^\kappa, n_{\text{ind}})$
- $(D, m_0, m_1, L) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot, \cdot)}(\text{mpk})$
- $T \leftarrow \text{Trace}^D(\text{tk}, 1^{1/\epsilon(\lambda)}, m_0, m_1, L)$

Here \mathcal{A} has the oracle access to $\text{KeyGen}(\text{msk}, \cdot, \cdot, \cdot)$, which has msk hardwired and on query $(\text{lb}, \text{id}, i)$, it outputs $\text{sk}_{\text{lb}, \text{id}, i}$. The adversary is admissible only if it makes at most one key query for each index.

Figure 3: Expt-Ind-TR

7.1 Indexed Trace and Revoke with Embedded Identity

In an indexed EITR scheme, the key is generated w.r.t. an identity and an index. We have $\mathcal{GID} = \mathcal{ID} \times [n_{\text{ind}}]$, where $[n_{\text{ind}}]$ is the index space for $n_{\text{ind}} \in \mathbb{N}$, $\text{params}_1 = (1^\kappa, n_{\text{ind}})$, $\text{gid} = (\text{id}, i) \in \mathcal{ID} \times [n_{\text{ind}}]$, and $\text{params}_2 = y$ for some $y > 1$.

We now define the secure tracing requirement.

Definition 7.5 (Secure Tracing). Let $\text{Ind-TR} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec}, \text{Trace})$ be an indexed EITR scheme. For any non-negligible function $\epsilon(\cdot)$ and stateful PPT adversary \mathcal{A} , define an experiment $\text{Expt-Ind-TR}_{\mathcal{A},\epsilon}(\lambda)$ as in Figure 3.

Let S be the set of key queries $(\text{lb}, \text{id}, i)$ queried by \mathcal{A} and $S_{\mathcal{ID}} = \{\text{id} : \exists \text{lb} \in \mathcal{L}, i \in [n_{\text{ind}}] \text{ s.t. } (\text{lb}, \text{id}, i) \in S\}$.

Consider the following probabilistic events and their corresponding probabilities:

- **Good-Decoder** : $\Pr[D(\text{ct}) = b : b \leftarrow \{0, 1\}, \text{ct} \leftarrow \text{Enc}(\text{mpk}, m_b, L)] \geq 1/2 + \epsilon(\lambda)$
 $\Pr\text{-Good-Decoder}_{\mathcal{A},\epsilon}(\lambda) = \Pr[\text{Good-Decoder}]$.
- **Cor-Tr** : $|T| > 0, (T \subseteq S_{\mathcal{ID}}) \wedge (T_{\text{lb}} \cap L = \phi)$, where $T_{\text{lb}} = \{\text{lb} \in \mathcal{L} : \exists i \in [n_{\text{ind}}], \exists \text{id} \in T, (\text{lb}, \text{id}, i) \in S\}$.
 $\Pr\text{-Cor-Tr}_{\mathcal{A},\epsilon}(\lambda) = \Pr[\text{Cor-Tr}]$.
- **Fal-Tr** : $(T \not\subseteq S_{\mathcal{ID}}) \wedge (T_{\text{lb}} \cap L = \phi)$
 $\Pr\text{-Fal-Tr}_{\mathcal{A},\epsilon}(\lambda) = \Pr[\text{Fal-Tr}]$.

The Ind-TR scheme is said to satisfy secure tracing property if for every PPT adversary \mathcal{A} and non-negligible function $\epsilon(\cdot)$, there exists negligible functions $\text{negl}_1(\cdot)$ and $\text{negl}_2(\cdot)$, such that for all $\lambda \in \mathbb{N}$, the following holds:

$$\Pr\text{-Fal-Tr}_{\mathcal{A},\epsilon}(\lambda) \leq \text{negl}_1(\lambda), \quad \Pr\text{-Cor-Tr}_{\mathcal{A},\epsilon}(\lambda) \geq \Pr\text{-Good-Decoder}_{\mathcal{A},\epsilon}(\lambda) - \text{negl}_2(\lambda).$$

Definition 7.6 (Selective Secure Tracing). The selective secure tracing of an indexed EITR scheme is defined in the same way as Def. 7.5, except that the adversary outputs the challenge revocation list L along with $(1^\kappa, 1^{n_{\text{ind}}})$ in the beginning of the $\text{Expt-Ind-TR}_{\mathcal{A},\epsilon}(\lambda)$.

Remark 7.7. In the definition above, Setup takes n_{ind} as an input in the binary form rather than in the unary form. This indicates that Setup runs in polynomial time, even if $n_{\text{ind}} = 2^{\text{poly}(\lambda)}$. On

Experiment $\text{Expt-BD-TR}_{\mathcal{A},\epsilon}(\lambda)$

- $1^\kappa, 1^{n_{\text{bd}}} \leftarrow \mathcal{A}(1^\lambda)$
- $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^\kappa, n_{\text{bd}})$
- $(D, m_0, m_1, L) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot)}(\text{mpk})$
- $T \leftarrow \text{Trace}^D(\text{tk}, 1^{1/\epsilon}, m_0, m_1, L)$

Here \mathcal{A} has the oracle access to $\text{KeyGen}(\text{msk}, \cdot, \cdot)$, which has msk hardwired and on query (lb, id) , it outputs $\text{sk}_{|\text{b}, \text{id}}$.

Figure 4: Expt-BD-TR

the other hand, the adversary outputs n_{ind} in the unary form in the security game (Fig. 3). This indicates that we consider the security game only for the case where $n_{\text{ind}} = \text{poly}(\lambda)$. Looking ahead, the former property is necessary when we convert bounded EITR into unbounded EITR in Sec. 10.

7.2 Bounded Trace and Revoke with Embedded Identity

In a bounded EITR scheme, we have $\mathcal{GID} = \mathcal{ID}$, $\text{params}_1 = (1^\kappa, 1^{n_{\text{bd}}})$, where $n_{\text{bd}} \in \mathbb{N}$ is the bound on the number of key queries that an adversary can make in the *correct trace experiment* game, $\text{gid} = \text{id}$ for $\text{id} \in \mathcal{ID}$, and $\text{params}_2 = y$ for some $y > 1$.

We now define the secure tracing requirement.

Definition 7.8 (Secure Tracing). Let $\text{BD-TR} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec}, \text{Trace})$ be a bounded EITR scheme. For any non-negligible function $\epsilon(\cdot)$ and stateful PPT adversary \mathcal{A} , define an experiment $\text{Expt-BD-TR}_{\mathcal{A},\epsilon}(\lambda)$ as in Figure 4.

Let S be the set of key queries made by \mathcal{A} and $S_{\mathcal{ID}} = \{\text{id} : \exists \text{lb} \in \mathcal{L} \text{ s.t. } (\text{lb}, \text{id}) \in S\}$. Consider the following probabilistic events and their corresponding probabilities :

- Good-Decoder : $\Pr[D(\text{ct}) = b : b \leftarrow \{0, 1\}, \text{ct} \leftarrow \text{Enc}(\text{mpk}, m_b, L)] \geq 1/2 + \epsilon(\lambda)$
 $\Pr\text{-Good-Decoder}_{\mathcal{A},\epsilon}(\lambda) = \Pr[\text{Good-Decoder} \wedge |S_{\mathcal{ID}}| \leq n_{\text{ind}}]$.
- Cor-Tr : $|T| > 0, (T \subseteq S_{\mathcal{ID}}) \wedge (T_{|\text{b}} \cap L = \phi)$, where $T_{|\text{b}} = \{\text{lb} \in \mathcal{L} : \exists \text{id} \in T, (\text{lb}, \text{id}) \in S\}$.
 $\Pr\text{-Cor-Tr}_{\mathcal{A},\epsilon}(\lambda) = \Pr[\text{Cor-Tr}]$.
- Fal-Tr : $(T \not\subseteq S_{\mathcal{ID}}) \wedge (T_{|\text{b}} \cap L = \phi)$
 $\Pr\text{-Fal-Tr}_{\mathcal{A},\epsilon}(\lambda) = \Pr[\text{Fal-Tr}]$.

The scheme is said to satisfy secure tracing if for every PPT adversary \mathcal{A} and non-negligible function $\epsilon(\cdot)$, there exists negligible functions $\text{negl}_1(\cdot)$ and $\text{negl}_2(\cdot)$, such that for all $\lambda \in \mathbb{N}$, the following holds:

$$\Pr\text{-Fal-Tr}_{\mathcal{A},\epsilon}(\lambda) \leq \text{negl}_1(\lambda), \quad \Pr\text{-Cor-Tr}_{\mathcal{A},\epsilon}(\lambda) \geq \Pr\text{-Good-Decoder}_{\mathcal{A},\epsilon}(\lambda) - \text{negl}_2(\lambda).$$

Definition 7.9 (Selective Secure Tracing). The selective secure tracing of a bounded EITR scheme is defined in the same way as Def. 7.8, except that the adversary outputs the challenge revocation list L along with $(1^\kappa, 1^{n_{\text{bd}}})$ in the beginning of the $\text{Expt-BD-TR}_{\mathcal{A},\epsilon}(\lambda)$.

Experiment $\text{Expt-TR}_{\mathcal{A},\epsilon,p}(\lambda)$

- $1^\kappa \leftarrow \mathcal{A}(1^\lambda)$.
- $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^\kappa)$.
- $(D, m_0, m_1, L) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot)}(\text{mpk})$.
- $T \leftarrow \text{Trace}^D(\text{tk}, 1^{1/\epsilon(\lambda)}, p(\lambda), m_0, m_1, L)$.

Here \mathcal{A} has the oracle access to $\text{KeyGen}(\text{msk}, \cdot, \cdot)$, which has msk hardwired and on query (lb, id) , it outputs $\text{sk}_{\text{lb}, \text{id}}$.

Figure 5: Expt-TR

Remark 7.10. We point out that the bound on the number of key queries by the adversary is only required for the correct trace guarantee. The false trace guarantee will hold even if the adversary exceeds n_{bd} key queries– this is essential for the transformation from bounded-EITR scheme to unbounded-EITR scheme.

7.3 Unbounded Trace and Revoke with Embedded Identity

In an unbounded EITR scheme, we have $\mathcal{GID} = \mathcal{ID}$, $\text{params}_1 = 1^\kappa$, $\text{gid} = \text{id}$ for $\text{id} \in \mathcal{ID}$, and $\text{params}_2 = (y, Q_{\text{bd}})$ where $y > 1$, $Q_{\text{bd}} \in \mathbb{N}$.

We now define the secure tracing requirement.

Definition 7.11 (Secure Tracing). Let $\text{TR} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec}, \text{Trace})$ be an unbounded EITR scheme. For any non-negligible function $\epsilon(\cdot)$, polynomial $p(\cdot)$ and stateful PPT adversary \mathcal{A} , define the experiment $\text{Expt-TR}_{\mathcal{A},\epsilon,p}(\lambda)$ as in Figure 5.

Let S be the set of key queries made by \mathcal{A} and $S_{\mathcal{ID}} = \{\text{id} : \exists \text{lb} \in \mathcal{L} \text{ s.t. } (\text{lb}, \text{id}) \in S\}$. Consider the following probabilistic events and their corresponding probabilities :

- Good-Decoder : $\Pr[D(\text{ct}) = b : b \leftarrow \{0, 1\}, \text{ct} \leftarrow \text{Enc}(\text{pk}, m_b, L) \geq 1/2 + \epsilon(\lambda)]$
 $\Pr\text{-Good-Decoder}_{\mathcal{A},\epsilon,p}(\lambda) = \Pr[\text{Good-Decoder} \wedge |S_{\mathcal{ID}}| \leq p(\lambda)]$.
- Cor-Tr : $|T| > 0, (T \subseteq S_{\mathcal{ID}}) \wedge (T_{\text{lb}} \cap L = \phi)$, where $T_{\text{lb}} = \{\text{lb} \in \mathcal{L} : \exists \text{id} \in T, (\text{lb}, \text{id}) \in S\}$.
 $\Pr\text{-Cor-Tr}_{\mathcal{A},\epsilon,p}(\lambda) = \Pr[\text{Cor-Tr}]$.
- Fal-Tr : $(T \not\subseteq S_{\mathcal{ID}}) \wedge (T_{\text{lb}} \cap L = \phi)$
 $\Pr\text{-Fal-Tr}_{\mathcal{A},\epsilon,p}(\lambda) = \Pr[\text{Fal-Tr}]$.

The scheme is said to satisfy secure traitor tracing property if for every PPT adversary \mathcal{A} and non-negligible function $\epsilon(\cdot)$, there exists negligible functions $\text{negl}_1(\cdot)$ and $\text{negl}_2(\cdot)$, such that for all $\lambda \in \mathbb{N}$, the following holds:

$$\Pr\text{-Fal-Tr}_{\mathcal{A},\epsilon,p}(\lambda) \leq \text{negl}_1(\lambda), \quad \Pr\text{-Cor-Tr}_{\mathcal{A},\epsilon,p}(\lambda) \geq \Pr\text{-Good-Decoder}_{\mathcal{A},\epsilon,p}(\lambda) - \text{negl}_2(\lambda).$$

Definition 7.12 (Selective Secure Tracing). The selective secure tracing of an unbounded EITR scheme is defined in the same way as Def. 7.11, except that the adversary outputs the challenge revocation list L along with 1^κ in the beginning of the $\text{Expt-TR}_{\mathcal{A},\epsilon}(\lambda)$.

Inputs: an identity id and an index i .

Hardwired Values : Indices $j \in [n_{\text{ind}}], \ell \in [\kappa]$, a bit $b \in \{0, 1\}$.

Output : 0/1.

$$f[j, \ell, b](\text{id}, i) = \begin{cases} 1 & \text{if } (i > j) \vee (i = j \wedge \ell = \perp) \vee (i = j \wedge \text{id}_\ell = 1 - b) \\ 0 & \text{otherwise} \end{cases}$$

Figure 6: Comparison Function $f[j, \ell, b]$

Remark 7.13. [Remark 10.4, [GKW19]]. Note that here the trace algorithm takes an additional parameter Q_{bd} . In the correct trace definition, we require that as long as the tracing algorithm uses a bound greater than the number of keys queried, the tracing algorithm must identify at least one traitor. However, the false trace guarantee should hold for all polynomially bounded Q_{bd} values. In particular, even if the number of keys queried is more than the bound used in tracing, the trace algorithm must not output an identity that was not queried. We can show that this definition implies the ‘standard’ tracing definition where the trace algorithm does not take this bound as input. One simply needs to run this bounded-version of trace with increasing powers of two until the trace algorithm outputs at least one traitor.

8 Indexed Trace and Revoke with Embedded Identity

In this section we construct an indexed (secret/public tracing)-EITR scheme from a (secret/public key)-RPE scheme. We present the construction and proofs for the secret trace setting primarily and also outline the differences in the public trace setting simultaneously.

8.1 Construction

Consider the identity space $\mathcal{ID} = \{0, 1\}^\kappa$ and index bound n_{ind} . For our purpose, it suffices to assume $n_{\text{ind}} \leq 2^{2\lambda}$. Let $\text{RPE} = (\text{RPE.Setup}, \text{RPE.KeyGen}, \text{RPE.Broadcast}, \text{RPE.Enc}, \text{RPE.Dec})$ be a (secret/public key) RPE scheme²⁵ with attribute space $\mathcal{ID} \times [n_{\text{ind}}]$, and supports the function class $\mathcal{F} = \{f[j, \ell, b]\}_{\substack{j \in [n_{\text{ind}}], \ell \in [\kappa], \\ b \in \{0, 1\}}}$ where $f[j, \ell, b] : \mathcal{ID} \times [n_{\text{ind}}] \rightarrow \{0, 1\}$ is as defined in figure 8.1.

We construct an indexed (secret/public tracing)-EITR scheme with identity space $\mathcal{ID} = \{0, 1\}^\kappa$ and index bound n_{ind} as follows :

$\text{Setup}(1^\lambda, 1^\kappa, n_{\text{ind}}) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm does the following:

- Samples $(\text{RPE.mpk}, \text{RPE.msk}) \leftarrow \text{RPE.Setup}(1^\lambda)$.
- Outputs $\text{mpk} = \text{RPE.mpk}$ and $\text{msk} = \text{RPE.msk}$.

$\text{KeyGen}(\text{msk}, (\text{lb}, \text{id}, i)) \rightarrow \text{sk}_{\text{lb}, \text{id}, i}$. The key generation algorithm does the following:

- Parse $\text{msk} = \text{RPE.msk}$.
- Sets $x = (\text{id}, i)$ and runs $\text{RPE.sk}_{\text{lb}, \text{id}, i} \leftarrow \text{RPE.KeyGen}(\text{RPE.msk}, \text{lb}, x)$.
- Outputs $\text{sk}_{\text{lb}, \text{id}, i} = \text{RPE.sk}_{\text{lb}, \text{id}, i}$.

²⁵Public key RPE does not have the Broadcast algorithm.

$\text{Enc}(\text{mpk}, m, L) \rightarrow \text{ct}$. The encryption algorithm does the following:

- Parse $\text{mpk} = \text{RPE.mpk}$.
- Compute $\text{RPE.ct} \leftarrow \text{RPE.Broadcast}(\text{RPE.mpk}, m, L)$. (In public trace setting, the algorithm computes $\text{RPE.ct} \leftarrow \text{RPE.Enc}(\text{RPE.mpk}, f[1, \perp, 0], m, L)$).
- Outputs $\text{ct} = \text{RPE.ct}$.

$\text{Dec}(\text{sk}_{\text{lb}, \text{id}, i}, \text{ct}, L) \rightarrow m'$. The decryption algorithm does the following:

- Parse $\text{sk}_{\text{lb}, \text{id}, i}$ as $\text{RPE.sk}_{\text{lb}, \text{id}, i}$ and ct as RPE.ct .
- Computes and outputs $m' \leftarrow \text{RPE.Dec}(\text{RPE.sk}_{\text{lb}, \text{id}, i}, \text{RPE.ct}, L)$.

$\text{Trace}^D(\text{tk}, y, m_0, m_1, L) \rightarrow T$. Parse tk as $\text{msk} = \text{RPE.msk}$ (in the public trace setting tk is $\text{mpk} = \text{RPE.mpk}$). The tracing algorithm is a two phased process. Briefly the tracing is implemented as follows:

1. First, we run the Index.Trace algorithm as defined in Figure 7, on the index space $[n_{\text{ind}}]$. If the key associated with index $i \in [n_{\text{ind}}]$ is used in constructing the decoder box, then the bit b in the output of the Index.Trace algorithm is 1. We maintain a set T_{index} to store all such indices on which $b = 1$.
2. Next, the ID.Trace algorithm, defined in Figure 8, takes the set T_{index} as input and uses the decoder D to compute the identity associated with each index.

Specifically, the algorithm runs as follows:

- Set $T_{\text{index}} := \phi$. For $i = 1$ to n_{ind} ,
 - Compute $(b, p, q) \leftarrow \text{Index.Trace}(\text{tk}, y, m_0, m_1, L, i)$.
 - If $b = 1$, set $T_{\text{index}} := T_{\text{index}} \cup (i, p, q)$.
- Set $T = \phi$. For $(i, p, q) \in T_{\text{index}}$,
 - Compute $\text{id} \leftarrow \text{ID.Trace}(\text{tk}, y, m_0, m_1, L, (i, p, q))$.
 - Set $T := T \cup \{\text{id}\}$.
- Output T .

Correctness. We prove that the above construction of indexed (secret/public tracing)-EITR scheme satisfies correctness (Def. 7.1) via the following theorem.

Theorem 8.1. *If RPE is a correct (secret/public key)-RPE scheme, then the above construction of indexed (secret/public tracing)-EITR scheme is correct.*

Proof. For the secret trace setting, the correctness follows directly from the broadcast correctness of the underlying (secret-key) RPE scheme.

For the public trace setting, firstly we observe that for any $(\text{id}, i) \in \mathcal{ID} \times [n_{\text{ind}}]$, $f[1, \perp, 0](\text{id}, i) = 1$. Hence, correctness follows directly from the encryption correctness of the underlying (public-key) RPE scheme. □

Algorithm Index.Trace(tk, y , m_0, m_1, L, i)

Inputs: Tracing key tk, a parameter y , messages m_0, m_1 , revocation list L and an index i .

Output: (b, p, q) , where $b \in \{0, 1\}$ and $p, q \in [0, 1] \cup \{\perp\}$.

Let $\epsilon = \lfloor 1/y \rfloor$. It sets $N = \lambda \cdot n_{\text{ind}}/\epsilon$ and $\text{temp}_1 = \text{temp}_2 = 0$. For $j = 1$ to N , it computes the following:

1. It samples $b_j \leftarrow \{0, 1\}$ and computes $\text{ct}_{j,1} \leftarrow \text{RPE.Enc}(\text{tk}, f[i, \perp, 0], m_{b_j}, L)$ and sends $\text{ct}_{j,1}$ to D .
If D outputs b_j , set $\text{temp}_1 = \text{temp}_1 + 1$, else set $\text{temp}_1 = \text{temp}_1 - 1$.
2. It samples $c_j \leftarrow \{0, 1\}$ and computes $\text{ct}_{j,2} \leftarrow \text{RPE.Enc}(\text{tk}, f[i + 1, \perp, 0], m_{c_j}, L)$ and sends $\text{ct}_{j,2}$ to D .
If D outputs c_j , set $\text{temp}_2 = \text{temp}_2 + 1$, else set $\text{temp}_2 = \text{temp}_2 - 1$.

If $\frac{\text{temp}_1 - \text{temp}_2}{N} > \frac{\epsilon}{4n_{\text{ind}}}$, output $(1, \frac{\text{temp}_1}{N}, \frac{\text{temp}_2}{N})$ else output $(0, \perp, \perp)$.

Figure 7: Index Tracing

Algorithm ID.Trace(tk, $y, m_0, m_1, L, (i, p, q)$)

Inputs: Tracing key tk, a parameter y , messages m_0, m_1 , revocation list L , index i , and probabilities p, q .

Output: $\text{id} \in \{0, 1\}^\kappa$

Let $\epsilon = \lfloor 1/y \rfloor$. It sets $N = \lambda \cdot n_{\text{ind}}/\epsilon$ and $\text{temp}_\ell = 0$ for $\ell \in [\kappa]$. For $\ell = 1$ to κ , it does as follows:

1. For $j = 1$ to N
It samples $b_j \leftarrow \{0, 1\}$ and computes $\text{ct}_j \leftarrow \text{RPE.Enc}(\text{tk}, f[i, \ell, 0], m_{b_j}, L)$ and sends ct_j to D .
If D outputs b_j , set $\text{temp}_\ell = \text{temp}_\ell + 1$, else set $\text{temp}_\ell = \text{temp}_\ell - 1$.

Let id be an empty string. For $\ell = 1$ to κ , do the following:

1. If $\frac{p+q}{2} > \frac{\text{temp}_\ell}{N}$, set $\text{id}_\ell = 0$, else set $\text{id}_\ell = 1$.

Output $\text{id} = \text{id}_1 \parallel \dots \parallel \text{id}_\kappa$.

Figure 8: Identity Tracing

Efficiency. We can instantiate the above construction using public key and secret key RPE. The size of each parameter is directly inherited from that of the underlying RPE. We have $|x| = |\text{id}| + \log n_{\text{ind}} \leq |\text{id}| + \lambda$ and $|f| := |f[j, \ell, b]| = \log n_{\text{ind}} + \log \kappa + 1 = O(\lambda)$. We note that here, $|f[j, \ell, b]|$ refers to the description size of the function $|f[j, \ell, b]|$, not the size of the circuit implementing it. In particular, the latter can depend on $|\text{id}|$ while the former is independent from $|\text{id}|$.

Secret Tracing Setting. If we instantiate the scheme with our secret key RPE in Sec. 6, we have

$$|\text{mpk}|, |\text{ct}|, |\text{sk}| = \text{poly}(\lambda, |f|, |x|, |\text{lb}|) = \text{poly}(\lambda, |\text{id}|, |\text{lb}|).$$

Public Tracing Setting. If we instantiate the scheme with our public key RPE in Sec. 4, we have

$$|\text{mpk}|, |\text{ct}| = \text{poly}(\lambda, |f|, |\text{lb}|) = \text{poly}(\lambda, |\text{lb}|), \quad |\text{sk}| = \text{poly}(\lambda, |f|, |x|, |\text{lb}|) = \text{poly}(\lambda, |\text{id}|, |\text{lb}|).$$

8.2 Security

In this section we show that our construction of the indexed (secret/public tracing)-EITR scheme is secure.

IND-CPA security.

Theorem 8.2. *If (secret/public key)-RPE scheme satisfies (0-query-selective/adaptive) message hiding property (Def. 3.4/Def. 3.3) then our construction of indexed (secret/public tracing)-EITR scheme is (selective/adaptive) IND-CPA secure (Def. 7.2/Def. 7.3).*

Proof. Recall that for IND-CPA security, we need

$$\text{Enc}(\text{mpk}, m_0, L) \approx_c \text{Enc}(\text{mpk}, m_1, L).$$

For indexed *secret tracing* EITR scheme, this is equivalent to

$$\text{RPE.Broadcast}(\text{RPE.mpk}, m_0, L) \approx_c \text{RPE.Broadcast}(\text{RPE.mpk}, m_1, L). \quad (8.1)$$

To prove this, we define the following hybrids:

Hybrid₀ : This is the real world with the challenge bit $b = 0$. In particular, the challenger returns the ciphertext $\text{ct} = \text{RPE.Broadcast}(\text{RPE.mpk}, m_0, L)$.

Hybrid₁ : In this hybrid, the challenger returns the ciphertext $\text{ct} = \text{RPE.Enc}(\text{RPE.mpk}, f[1, \perp, 0], m_0, L)$. Indistinguishability from Hybrid₀ follows from the broadcast security of secret-key RPE.

Hybrid₂ : In this hybrid, the challenger returns the ciphertext $\text{ct} = \text{RPE.Enc}(\text{RPE.mpk}, f[1, \perp, 0], m_1, L)$. Indistinguishability from Hybrid₁ follows from the message hiding property of secret-key RPE.

Hybrid₃ : In this hybrid, the challenger returns the ciphertext $\text{ct} = \text{RPE.Broadcast}(\text{RPE.mpk}, m_1, L)$. This is the real world with $b = 1$. Indistinguishability from Hybrid₂ follows from the broadcast security of secret-key RPE.

For indexed *public trace* EITR scheme, IND-CPA security is equivalent to

$$\text{RPE.Enc}(\text{RPE.mpk}, f[1, \perp, 0], m_0, L) \approx_c \text{RPE.Enc}(\text{RPE.mpk}, f[1, \perp, 0], m_1, L). \quad (8.2)$$

The indistinguishability here follows directly from the message hiding property of the underlying RPE scheme. \square

Secure Tracing Analysis. We now show that our construction of indexed (secret/public tracing)-EITR scheme achieves the correct and false trace guarantees. The following analysis is mostly taken from ([GKW19], Section 5.2.2) with appropriate modifications to incorporate the revocation list.

False Trace Guarantee. The false trace guarantee of a trace and revoke scheme ensures that the tracing algorithm does not falsely accuse any user. We prove that there does not exist a PPT adversary who can output a decoder D such that the tracing algorithm, when executed using this decoder D , outputs an identity for which key was not queried by the adversary or the corresponding label is in the revocation list.

Theorem 8.3. *For every stateful PPT adversary \mathcal{A} and non-negligible function $\epsilon(\cdot)$, there exists negligible functions $\text{negl}(\cdot)$, such that for all $\lambda \in \mathbb{N}$*

$$\Pr\text{-Fal-Tr}_{\mathcal{A},\epsilon}(\lambda) \leq \text{negl}(\lambda)$$

where the probability $\Pr\text{-Fal-Tr}_{\mathcal{A},\epsilon}$ is as defined in Def. 7.5.

Proof. Let $S \subseteq \mathcal{L} \times \mathcal{ID} \times [n_{\text{ind}}]$ be the set of label-identity-index pairs on which \mathcal{A} issues key queries. That is, $S = \{(\text{lb}, \text{id}, i) : \mathcal{A} \text{ issues a key query on } (\text{lb}, \text{id}, i)\}$. Let us first recall the assumption that \mathcal{A} can issue at most one key query for any index $i \in [n_{\text{ind}}]$. Next, we set up some notations. For $\text{lb} \in \mathcal{L}$, $i \in [n_{\text{ind}}]$, $\text{id} \in \mathcal{ID}$ and a revocation list L , let

$$\begin{aligned} S_{\text{index}} &= \{i : (\text{lb}, \text{id}, i) \in S \text{ for some } (\text{lb}, \text{id}) \in \mathcal{L} \times \mathcal{ID}\}, \\ L_{\text{index}} &= \{i : (\text{lb}, \text{id}, i) \in S \text{ and } \text{lb} \in L\}, \\ B &= \{(\text{id}, i) : (\text{lb}, \text{id}, i) \in S \text{ and } \text{lb} \notin L\}, \\ B_{\text{index}} &= \{i : (\text{id}, i) \in B\}. \end{aligned}$$

For any decoder box D , messages m_0, m_1 , revocation list L , for any $i \in [n_{\text{ind}} + 1]$, $\ell \in [\kappa]$, $\text{lb} \in \mathcal{L}$ we define

$$p_{i,\perp}^D = \Pr[D(\text{ct}) = b : b \leftarrow \{0, 1\}, \text{ct} \leftarrow \text{RPE.Enc}(\text{tk}, f[i, \perp, 0], m_b, L)]$$

$$p_{i,\ell}^D = \Pr[D(\text{ct}) = b : b \leftarrow \{0, 1\}, \text{ct} \leftarrow \text{RPE.Enc}(\text{tk}, f[i, \ell, 0], m_b, L)]$$

where the probability is taken over the random coins to decoder D and the randomness used during the encryption. We show that $\Pr\text{-Fal-Tr}_{\mathcal{A},\epsilon}(\lambda) \leq \text{negl}(\lambda)$. For $i \in [n_{\text{ind}}]$, $\ell \in [\kappa]$, we also define the following events:

$$\text{Diff-Adv}_i^D : p_{i,\perp}^D - p_{i+1,\perp}^D > \epsilon/8n_{\text{ind}}$$

$$\text{Diff-Adv}_{i,\ell,\text{lwr}}^D : p_{i,\perp}^D - p_{i,\ell}^D > \epsilon/16n_{\text{ind}}$$

$$\text{Diff-Adv}_{i,\ell,\text{upr}}^D : p_{i,\ell}^D - p_{i+1,\perp}^D > \epsilon/16n_{\text{ind}}$$

$$\text{Diff-Adv}^D : \bigvee_{i \in [n_{\text{ind}}] \setminus (S_{\text{index}} \setminus L_{\text{index}})} \text{Diff-Adv}_i^D \quad \bigvee_{\substack{(\text{id}, i) \in B, \ell \in [\kappa], \\ \text{s.t. } \text{id}_\ell = 1}} \text{Diff-Adv}_{i,\ell,\text{lwr}}^D \quad \bigvee_{\substack{(\text{id}, i) \in B, \ell \in [\kappa] \\ \text{s.t. } \text{id}_\ell = 0}} \text{Diff-Adv}_{i,\ell,\text{upr}}^D$$

We will drop the dependence on D for the ease of notation. We have

$$\begin{aligned}
\Pr[\text{Fal-Tr}] &= \Pr[\text{Fal-Tr} \mid \overline{\text{Diff-Adv}}] \Pr[\overline{\text{Diff-Adv}}] + \Pr[\text{Fal-Tr} \mid \text{Diff-Adv}] \Pr[\text{Diff-Adv}] \\
&\leq \Pr[\text{Fal-Tr} \mid \overline{\text{Diff-Adv}}] + \Pr[\text{Diff-Adv}] \\
&= \Pr[\text{Fal-Tr} \mid \overline{\text{Diff-Adv}}] + \sum_{i \in [n_{\text{ind}}]} \Pr[i \notin S_{\text{index}} \setminus L_{\text{index}} \wedge \text{Diff-Adv}_i] \\
&\quad + \sum_{(i,\ell) \in [n_{\text{ind}}] \times [\kappa]} \Pr \left[\exists \text{id} \in \{0,1\}^\kappa \text{ s.t. } (\text{id}, i) \in B \right. \\
&\quad \quad \left. \wedge ([\text{Diff-Adv}_{i,\ell,\text{lwr}} \wedge \text{id}_\ell = 1] \vee [\text{Diff-Adv}_{i,\ell,\text{upr}} \wedge \text{id}_\ell = 0]) \right]
\end{aligned}$$

Now, we argue that each of the terms on the RHS is bounded by a negligible function.

Lemma 8.4. *For every stateful PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}_1(\cdot)$, such that $\forall \lambda \in \mathbb{N}$, we have*

$$\Pr[\text{Fal-Tr} \mid \overline{\text{Diff-Adv}}] \leq \text{negl}_1(\lambda).$$

Proof. Note that the event Fal-Tr occurs if and only if the trace algorithm outputs an identity that was not queried by the adversary or outputs an identity which was queried but also revoked. Since, the tracing scheme is two phased, the false tracing can happen in any of the two stages. First, during the index-tracing procedure, it can happen that $\exists (i, p, q) \in T_{\text{index}}$, such that, $i \notin S_{\text{index}} \setminus L_{\text{index}}$ and secondly, during the identity-tracing procedure, the ID.Trace algorithm outputs an incorrect identity corresponding to some $i \in S_{\text{index}} \setminus L_{\text{index}}$. So, we have

$$\begin{aligned}
&\Pr[\text{Fal-Tr} \mid \overline{\text{Diff-Adv}}] \\
&\leq \sum_{i \in [n_{\text{ind}}]} \Pr[\text{Fal-Tr} \wedge i \notin S_{\text{index}} \setminus L_{\text{index}} \wedge (\exists p, q : (i, p, q) \in T_{\text{index}}) \mid \overline{\text{Diff-Adv}}] \\
&\quad + \sum_{(i,\ell) \in [n_{\text{ind}}] \times [\kappa]} \Pr[\text{Fal-Tr} \wedge \exists \text{id}, \hat{\text{id}} : (\text{id}, i) \in B \wedge (\exists p, q : (i, p, q) \in T_{\text{index}}) \\
&\quad \quad \wedge \hat{\text{id}} \leftarrow \text{ID.Trace}(\text{tk}, 1^y, m_0, m_1, (i, p, q)) \wedge \text{id}_\ell \neq \hat{\text{id}}_\ell \mid \overline{\text{Diff-Adv}}].
\end{aligned}$$

Consider the first term on RHS. If $\overline{\text{Diff-Adv}}$ happens then $\forall i \notin S_{\text{index}} \setminus L_{\text{index}}$ event $\overline{\text{Diff-Adv}}_i$ happens.

We know that $\overline{\text{Diff-Adv}}_i$ implies $p_{i,\perp} - p_{i+1,\perp} \leq \epsilon/8n_{\text{ind}}$. Also the event $(\exists p, q : (i, p, q) \in T_{\text{index}})$ implies $\hat{p}_{i,\perp} - \hat{p}_{i+1,\perp} > \epsilon/4n_{\text{ind}}$, where $\hat{p}_{j,\perp}$ is the estimated probability for $p_{j,\perp}$.

By applying Chernoff bound, we have, for every $i \in [n_{\text{ind}}]$

$$\Pr[i \notin S_{\text{index}} \setminus L_{\text{index}} \wedge (\exists p, q : (i, p, q) \in T_{\text{index}}) \mid \overline{\text{Diff-Adv}}] \leq e^{-\lambda/24} \leq 2^{-O(\lambda)}.$$

Now, in the second term, for a fixed (i, ℓ) the probability term corresponds to the event when the ID.Trace outputs an identity $\hat{\text{id}}$ corresponding to index i , such that $\hat{\text{id}}_\ell \neq \text{id}_\ell$, where id is such that the adversary made a key query for (id, i) . So, if the event $\overline{\text{Diff-Adv}}$ happens then $\forall (\text{id}, i) \in B, \ell \in [\kappa]$, event $\overline{\text{Diff-Adv}}_{i,\ell,X}$ happens where $X = \text{lwr}$ if $\text{id}_\ell = 1$ else $X = \text{upr}$.

Thus, when we have $\text{id}_\ell = 1$, then $\overline{\text{Diff-Adv}}_{i,\ell,\text{lwr}}$ implies $p_{i,\perp} - p_{i,\ell} \leq \epsilon/16n_{\text{ind}}$.

Also, the event $(\exists p, q : (i, p, q) \in T_{\text{index}})$ implies that $\hat{p}_{i,\perp} - \hat{p}_{i+1,\perp} \geq \epsilon/4n_{\text{ind}}$ and the event $\hat{\text{id}} \leftarrow \text{ID.Trace}(\text{tk}, 1^y, m_0, m_1, (i, p, q)) \wedge \text{id}_\ell = 0$ implies $(\hat{p}_{i,\perp} + \hat{p}_{i+1,\perp} > 2\hat{p}_{i,\ell})$. These together imply that $\hat{p}_{i,\perp} - \hat{p}_{i,\ell} > \epsilon/8n_{\text{ind}}$. Similarly, when $\text{id}_\ell = 0$ and $\hat{\text{id}}_\ell = 1$, $\overline{\text{Diff-Adv}}_{i,\ell,\text{upr}}$ implies

$p_{i,\ell} - p_{i+1,\perp} \leq \epsilon/16n_{\text{ind}}$ and following the similar reasoning as above, we get $\hat{p}_{i,\perp} - \hat{p}_{i+1,\perp} \geq \epsilon/4n_{\text{ind}}$ and $\hat{p}_{i,\perp} + \hat{p}_{i+1,\perp} \leq 2\hat{p}_{i,\ell}$ which implies $\hat{p}_{i,\ell} - \hat{p}_{i+1,\perp} \geq \epsilon/8n_{\text{ind}}$. Hence, using Chernoff bound, we get that

$$\Pr[\exists \text{id}, \hat{\text{id}} : (\text{id}, i) \in S \wedge (\exists p, q : (i, p, q) \in T_{\text{index}}) \wedge \hat{\text{id}} \leftarrow \text{ID.Trace}(\text{tk}, 1^y, m_0, m_1, (i, p, q)) \\ \wedge \hat{\text{id}} \in T \wedge \text{id}_\ell \neq \hat{\text{id}}_\ell \mid \overline{\text{Diff-Adv}}] \leq 2^{-O(\lambda)}.$$

Combining the probability of both the RHS terms, we get

$$\Pr[\text{Fal-Tr} \mid \overline{\text{Diff-Adv}}] \leq n_{\text{ind}} \cdot 2^{-O(\lambda)} + n_{\text{ind}} \cdot \kappa \cdot 2^{-O(\lambda)} = \text{negl}_1(\lambda).$$

□

Lemma 8.5. *If the underlying (secret/public key)-RPE satisfies (1-query-selective/adaptive) function hiding security property (Def. 3.6/Def. 3.5), then for every stateful PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}_2(\cdot)$, such that $\forall \lambda \in \mathbb{N}$ and $i \in [n_{\text{ind}}]$, we have*

$$\Pr[i \notin S_{\text{index}} \setminus L_{\text{index}} \wedge \text{Diff-Adv}_i] \leq \text{negl}_2(\lambda).$$

Proof. We give a proof by contradiction. Suppose there exists a PPT adversary \mathcal{A} that outputs a good decoder D along with messages m_0, m_1 and a revocation list L such that there exists $i^* \in [n_{\text{ind}}]$ for which $\Pr[i^* \notin S_{\text{index}} \setminus L_{\text{index}} \wedge \text{Diff-Adv}_{i^*}] \geq \delta$ where δ is some non-negligible function in the security parameter λ . We use this adversary \mathcal{A} to build a reduction \mathcal{B} that can break the function hiding security of the underlying RPE scheme. The reduction is as follows:

1. In the beginning, \mathcal{A} outputs $1^{n_{\text{ind}}}, 1^\kappa, L$ (\mathcal{A} will output L adaptively in the public trace setting).
2. The RPE challenger samples $(\text{RPE.mpk}, \text{RPE.msk}) \leftarrow \text{RPE.Setup}(1^\lambda)$ and sends RPE.mpk to \mathcal{B} . \mathcal{B} sets $\text{mpk} = \text{RPE.mpk}$ and forwards it to \mathcal{A} .
3. When \mathcal{A} issues a secret key query $(\text{lb}, \text{id}, j)$ (recall that \mathcal{A} can make at most one query for each index j), \mathcal{B} sends a key query to the RPE challenger on $(\text{lb}, x = (\text{id}, j))$. The challenger returns $\text{sk}_{\text{lb}, \text{id}, j}$ which \mathcal{B} forwards to \mathcal{A} .
4. In the end, \mathcal{A} outputs a decoder D and messages m_0, m_1 (and a revocation list L in the case of public tracing EITR scheme) to \mathcal{B} . \mathcal{B} then does the following:
 - (a) Samples $i \leftarrow [n_{\text{ind}}] \setminus (S_{\text{index}} \setminus L_{\text{index}})$ and sets $f_0 = f[i, \perp, 0]$ and $f_1 = f[i + 1, \perp, 0]$.
 - (b) Samples $b \leftarrow \{0, 1\}$ and sends (f_0, f_1, m_b) to the RPE challenger. The challenger samples $\alpha \leftarrow \{0, 1\}$, computes $\text{ct}_1 \leftarrow \text{RPE.Enc}(\text{RPE.msk}, f_\alpha, m_b, L)$ and sends ct_1 to \mathcal{B} . (In public trace setting, $\text{ct}_1 \leftarrow \text{RPE.Enc}(\text{RPE.mpk}, f_\alpha, m_b, L)$.)
 - (c) \mathcal{B} samples $\beta \leftarrow \{0, 1\}$ and sends a encryption query on (f_β, m_b, L) and sets the ciphertext returned by the RPE challenger as ct_2 . (In public trace setting, \mathcal{B} computes ct_2 itself as $\text{ct}_2 \leftarrow \text{RPE.Enc}(\text{RPE.mpk}, f_\beta, m_b, L)$.)
 - (d) \mathcal{B} samples $\beta \leftarrow \{0, 1\}$, computes $\text{ct}_2 \leftarrow \text{RPE.Enc}(\text{RPE.mpk}, f_\beta, m_b, L)$.
 - (e) If $D(\text{ct}_1) = D(\text{ct}_2)$, sets $\alpha' = \beta$, else $\alpha' = 1 - \beta$.
 - (f) \mathcal{B} returns α' to the RPE challenger.

\mathcal{B} wins the game if $\alpha = \alpha'$. Note that since $i \leftarrow [n_{\text{ind}}] \setminus (S_{\text{index}} \setminus L_{\text{index}})$, for every key query $(\text{lb}, \text{id}, j)$ that \mathcal{B} issues to the RPE challenger, either $f_0(\text{id}, j) = f_1(\text{id}, j)$ (when $i \notin S_{\text{index}}$) or $\text{lb} \in L$ (when $i \in S_{\text{index}} \cap L_{\text{index}}$). This establishes the admissibility of \mathcal{B} . Now, let us analyse the probability of \mathcal{B} winning against the RPE challenger.

Let $q_{j,b} = \Pr[D(\text{ct}) = b : b \leftarrow \{0, 1\}, \text{ct} \leftarrow \text{RPE.Enc}(\text{tk}, f[j, \perp, 0], m_b, L)]$.

Let E^D be the event that \mathcal{B} wins and E_b^D be the event that \mathcal{B} wins when it samples b as the message bit in step 4b. So, we have

$$\begin{aligned} \Pr[E_b^D] &= \frac{1}{4} (\Pr[E_b^D \mid \alpha = 0, \beta = 0] + \Pr[E_b^D \mid \alpha = 0, \beta = 1] \\ &\quad + \Pr[E_b^D \mid \alpha = 1, \beta = 0] + \Pr[E_b^D \mid \alpha = 1, \beta = 1]) \\ &= \frac{1}{4} (q_{i,b}^2 + (1 - q_{i,b})^2 + 2(q_{i,b}(1 - q_{i+1,b}) + (1 - q_{i,b})q_{i+1,b}) + q_{i+1,b}^2 + (1 - q_{i+1,b})^2) \\ &= \frac{1}{2} + \frac{(q_{i,b} - q_{i+1,b})^2}{2}. \end{aligned}$$

We have assumed $\exists i^* \in [n_{\text{ind}}]$ such that $\Pr[i^* \notin S_{\text{index}} \setminus L_{\text{index}} \wedge \text{Diff-Adv}_{i^*}] \geq \delta$. $\Pr[i = i^*] = 1/n_{\text{ind}}$. So we get $\Pr[i = i^* \wedge i^* \notin S_{\text{index}} \setminus L_{\text{index}} \wedge \text{Diff-Adv}_{i^*}] \geq \delta/n_{\text{ind}}$.

Let F be the event: $\exists i^* \in [n_{\text{ind}}]$ such that $i = i^* \wedge i^* \notin S_{\text{index}} \setminus L_{\text{index}} \wedge \text{Diff-Adv}_{i^*}$. Then when F occurs, we have $p_{i,\perp} - p_{i+1,\perp} > \epsilon/8n_{\text{ind}} \Rightarrow \exists b' \in \{0, 1\}$ s.t $q_{i,b'} - q_{i+1,b'} > \epsilon/8n_{\text{ind}}$. Also, $\Pr[E_{b'}^D] \geq 1/2$ for $b \in \{0, 1\}$, irrespective of occurrence of F . Now,

$$\begin{aligned} \Pr[E_{b'}^D] &= \Pr[E_{b'}^D \mid F] \Pr[F] + \Pr[E_{b'}^D \mid \bar{F}] \Pr[\bar{F}] \\ &\geq (1/2 + \frac{\epsilon^2}{128n_{\text{ind}}^2}) \times (\delta/n_{\text{ind}}) + 1/2 \times (1 - \delta/n_{\text{ind}}) \\ &= 1/2 + \frac{\epsilon^2 \delta}{128n_{\text{ind}}^3}. \end{aligned}$$

Again,

$$\Pr[E^D] = \frac{\Pr[E_{b'}^D]}{2} + \frac{\Pr[E_{b'}^D]}{2} \geq \frac{1}{2} \left(\frac{1}{2} + \frac{\epsilon^2 \delta}{128n_{\text{ind}}^3} \right) + \frac{1}{4} = \frac{1}{2} + \eta, \text{ where } \eta = \frac{\epsilon^2 \delta}{128n_{\text{ind}}^3}.$$

Thus, \mathcal{B} wins against the function-hiding security of the underlying RPE scheme with advantage $\geq \eta$, which is non-negligible for non-negligible ϵ and δ , a contradiction. Hence the lemma follows. \square

Lemma 8.6. *If the underlying (secret/public key)-RPE scheme satisfies (1-query-selective/adaptive) function hiding security property (Def. 3.6/Def. 3.5), then for every stateful PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}_3(\cdot)$, such that $\forall \lambda \in \mathbb{N}$ and $i \in [n_{\text{ind}}], \ell \in [\kappa]$, we have*

$$\begin{aligned} \Pr \left[\exists \text{id} \in \{0, 1\}^\kappa \text{ s.t } (\text{id}, i) \in B \wedge ([\text{Diff-Adv}_{i,\ell,\text{lwr}} \wedge \text{id}_\ell = 1] \vee [\text{Diff-Adv}_{i,\ell,\text{upr}} \wedge \text{id}_\ell = 0]) \right] \\ \leq \text{negl}_3(\lambda). \end{aligned}$$

Proof. The proof is similar to the proof of Lemma 8.5. We show that if there exists a PPT adversary \mathcal{A} who outputs a good decoder D along with messages m_0, m_1 and a revocation list L for which there exists $i^* \in [n_{\text{ind}}], \ell^* \in [\kappa]$ such that $\Pr[\exists \text{id} \in \{0, 1\}^\kappa \text{ s.t } (i^*, \text{id}) \in B \wedge ([\text{Diff-Adv}_{i^*,\ell^*,\text{lwr}} \wedge \text{id}_{\ell^*} = 1] \vee [\text{Diff-Adv}_{i^*,\ell^*,\text{upr}} \wedge \text{id}_{\ell^*} = 0])] \geq \delta$.

We use this adversary \mathcal{A} to build a reduction \mathcal{B} that can break the function hiding security of the underlying RPE scheme. \mathcal{B} is defined in the same way as in the proof of Lemma 8.5, except Step 4a, which is now described as follows:

\mathcal{B} samples $i \leftarrow B_{\text{index}}$, $\ell \in [\kappa]$ and a bit $g \leftarrow \{0, 1\}$ and sets $f_0 = f[i, \perp, 0]$ and $f_1 = f[i, \ell, 0]$, if $g = 0$, else sets $f_0 = f[i, \ell, 0]$ and $f_1 = f[i + 1, \perp, 0]$.

Analysis of \mathcal{B} 's advantage:

Let $q_{j,k,b} = \Pr[D(\text{ct}) = b : b \leftarrow \{0, 1\}, \text{ct} \leftarrow \text{RPE.Enc}(\text{tk}, f[j, k, 0], m_b, L)]$.

Let E^D be the event that \mathcal{B} wins and E_b^D be the event that \mathcal{B} wins when \mathcal{B} samples b as the message bit in step 4b. So, we have

$$\begin{aligned} \Pr[E_b^D] &= \frac{1}{8} (\Pr[E_b^D \mid \alpha = 0, \beta = 0, g = 0] + \Pr[E_b^D \mid \alpha = 0, \beta = 1, g = 0] \\ &\quad + \Pr[E_b^D \mid \alpha = 1, \beta = 0, g = 0] + \Pr[E_b^D \mid \alpha = 1, \beta = 1, g = 0] \\ &\quad + \Pr[E_b^D \mid \alpha = 0, \beta = 0, g = 1] + \Pr[E_b^D \mid \alpha = 0, \beta = 1, g = 1] \\ &\quad + \Pr[E_b^D \mid \alpha = 1, \beta = 0, g = 1] + \Pr[E_b^D \mid \alpha = 1, \beta = 1, g = 1]) \\ &= \frac{1}{8} (q_{i,\perp,b}^2 + (1 - q_{i,\perp,b})^2 + 2(q_{i,\perp,b}(1 - q_{i,\ell,b}) + (1 - q_{i,\perp,b})q_{i,\ell,b}) + q_{i,\ell,b}^2 \\ &\quad + (1 - q_{i,\ell,b})^2 + q_{i,\ell,b}^2 + (1 - q_{i,\ell,b})^2 + 2(q_{i,\ell,b}(1 - q_{i+1,\perp,b}) + \\ &\quad + (1 - q_{i,\ell,b})q_{i+1,\perp,b}) + q_{i+1,\perp,b}^2 + (1 - q_{i+1,\perp,b})^2) \\ &= \frac{1}{2} + \frac{(q_{i,\perp,b} - q_{i,\ell,b})^2}{4} + \frac{(q_{i,\ell,b} - q_{i+1,\perp,b})^2}{4}. \end{aligned}$$

We have assumed that $\exists i^* \in [n_{\text{ind}}], \ell^* \in [\kappa]$ such that $\Pr[\exists \text{id} \in \{0, 1\}^\kappa \text{ s.t. } (i^*, \text{id}) \in B \wedge ([\text{Diff-Adv}_{i^*,\ell^*,\text{lwr}} \wedge \text{id}_{\ell^*} = 1] \vee [\text{Diff-Adv}_{i^*,\ell^*,\text{upr}} \wedge \text{id}_{\ell^*} = 0])] \geq \delta$.

Using $\Pr[i = i^* \wedge \ell = \ell^*] = 1/\kappa n_{\text{ind}}$, we get $\Pr[(i = i^* \wedge \ell = \ell^*) \wedge \exists \text{id} \in \{0, 1\}^\kappa \text{ s.t. } ((i^*, \text{id}) \in B \wedge ([\text{Diff-Adv}_{i^*,\ell^*,\text{lwr}} \wedge \text{id}_{\ell^*} = 1] \vee [\text{Diff-Adv}_{i^*,\ell^*,\text{upr}} \wedge \text{id}_{\ell^*} = 0]))] \geq \frac{\delta}{\kappa n_{\text{ind}}}$

Let F be the event: $\exists i^* \in [n_{\text{ind}}], \ell^* \in [\kappa]$ such that $(\exists \text{id} \in \{0, 1\}^\kappa \text{ s.t. } (i^*, \text{id}) \in B \wedge ([\text{Diff-Adv}_{i^*,\ell^*,\text{lwr}} \wedge \text{id}_{\ell^*} = 1] \vee [\text{Diff-Adv}_{i^*,\ell^*,\text{upr}} \wedge \text{id}_{\ell^*} = 0]))$. Then when F occurs, we have $(p_{i,\perp} - p_{i,\ell} > \epsilon/8n_{\text{ind}}) \vee (p_{i,\ell} - p_{i+1,\perp} > \epsilon/8n_{\text{ind}})$. This implies that there exists $b' \in \{0, 1\}$ s.t. $(q_{i,\perp,b'} - q_{i,\ell,b'} > \epsilon/8n_{\text{ind}}) \vee ((q_{i,\ell,b'} - q_{i+1,\perp,b'} > \epsilon/8n_{\text{ind}}))$. Also, $\Pr[E_{b'}^D] \geq 1/2$ for $b \in \{0, 1\}$, irrespective of occurrence of F . Now,

$$\begin{aligned} \Pr[E_{b'}^D] &= \Pr[E_{b'}^D \mid F] \Pr[F] + \Pr[E_{b'}^D \mid \bar{F}] \Pr[\bar{F}] \\ &\geq (1/2 + \frac{\epsilon^2}{256n_{\text{ind}}^2}) \times (\delta/\kappa n_{\text{ind}}) + 1/2 \times (1 - \delta/\kappa n_{\text{ind}}) \\ &= 1/2 + \frac{\epsilon^2 \delta}{256\kappa n_{\text{ind}}^3}. \end{aligned}$$

Again,

$$\Pr[E^D] = \frac{\Pr[E_{b'}^D]}{2} + \frac{\Pr[E_{b'}^D]}{2} \geq \frac{1}{2} \left(\frac{1}{2} + \frac{\epsilon^2 \delta}{256\kappa n_{\text{ind}}^3} \right) + \frac{1}{4} = \frac{1}{2} + \eta, \text{ where } \eta = \frac{\epsilon^2 \delta}{512n_{\text{ind}}^3 \kappa}.$$

□

Combining the result of Lemmas 8.4, 8.5 and 8.6, we get

$$\Pr\text{-Fal-Tr}_{\mathcal{A},\epsilon}(\lambda) \leq \text{negl}_1(\lambda) + n_{\text{ind}} \cdot \text{negl}_2(\lambda) + n_{\text{ind}} \cdot \kappa \cdot \text{negl}_3(\lambda) = \text{negl}(\lambda).$$

□

Correct trace guarantee. We prove that whenever an adversary outputs a good decoder, the tracing algorithm will output, with all but negligible probability, at least one valid user identity which was queried by the adversary.

Theorem 8.7. *If the underlying (secret/public key)-RPE scheme in the indexed (secret/public tracing)-EITR scheme construction satisfies (1-query-selective broadcast and message/adaptive message) hiding property, then for every stateful PPT adversary \mathcal{A} for the (selective/adaptive)-tracing game (Def. 7.6/Def. 7.5) and non-negligible function ϵ , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, the following holds*

$$\Pr[\text{Cor-Tr}] \geq \Pr[\text{Good-Decoder}] - \text{negl}(\lambda)$$

Proof. Let us define $p_{\text{Broadcast}}^D = \Pr[D(\text{ct}) = b : b \leftarrow \{0, 1\}, \text{ct} \leftarrow \text{RPE.Broadcast}(\text{RPE.mpk}, m_b, L)]$. Then in the secret trace setting, if the event $\text{Good-Decoder}_{\mathcal{A}, \epsilon}$ occurs, then this implies $p_{\text{Broadcast}}^D \geq 1/2 + \epsilon$. Further, from the broadcast security of (secret-key) RPE, we get $p_{1, \perp}^D \geq p_{\text{Broadcast}}^D - \text{negl}_1(\lambda)$, which implies

$$p_{1, \perp}^D \geq 1/2 + \epsilon - \text{negl}_1(\lambda). \quad (8.3)$$

In public trace setting, the event $\text{Good-Decoder}_{\mathcal{A}, \epsilon}$ directly implies $p_{1, \perp}^D \geq 1/2 + \epsilon$ by definition. Also, by message hiding property of the underlying (secret/public key) RPE scheme, we have

$$p_{n_{\text{ind}}+1, \perp}^D \leq 1/2 + \text{negl}_2(\lambda) \quad (8.4)$$

for some negligible function $\text{negl}_2(\cdot)$ with overwhelming probability. This is so, because $f[n_{\text{ind}}, \perp, 0](\text{id}, i) = 0$ for all $(\text{id}, i) \in \mathcal{ID} \times [n_{\text{ind}}]$.

Combining equations (8.3) and (8.4), we get $p_{1, \perp}^D - p_{n_{\text{ind}}+1, \perp}^D > \epsilon/2$.

Let $S^{\text{ind}} = \{i \in [n_{\text{ind}}] \mid p_{i, \perp}^D - p_{i+1, \perp}^D > \epsilon/2n_{\text{ind}}\}$. Then if the event Good-Decoder occurs, $S^{\text{ind}} \neq \phi$.

By Chernoff bound, we have

$$\forall i \in S^{\text{index}}, \quad \Pr[\hat{p}_{i, \perp}^D - \hat{p}_{i+1, \perp}^D \leq \epsilon/4n_{\text{ind}}] \leq 2^{-O(\lambda)} = \text{negl}_3(\lambda)$$

for some negligible function $\text{negl}_3(\cdot)$. Here, \hat{p} denotes the estimate for p computed by tracing algorithm.

So, with all but negligible probability,

$$T_{\text{index}} \neq \phi \text{ and } \forall (i, p, q) \in T_{\text{index}}, p - q > \epsilon/4n_{\text{ind}}$$

where T_{index} and p, q are as defined in the Trace algorithm 8.1.

Note that the ID.Trace algorithm (Figure 8) takes as input $(i, p, q) \in T_{\text{index}}$ and outputs a corresponding id , where $\text{id}_\ell = 1$ if $\hat{p}_{i, \ell}^D > (p + q)/2$ else $\text{id}_\ell = 0$ for $\ell \in [\kappa]$. Then, for every $(i, p, q) \in T_{\text{index}}$, the tracing algorithm outputs an identity. Hence, if $T_{\text{index}} \neq \phi \Rightarrow T \neq \phi$, where T is defined as in the tracing algorithm. So,

$$\begin{aligned} \Pr[T \neq \phi] &\geq \Pr[T \neq \phi \wedge \text{Good-Decoder}] \\ &\geq (1 - \text{negl}_2(\lambda)) \Pr[\text{Good-Decoder}] \\ &\geq \Pr[\text{Good-Decoder}] - \text{negl}(\lambda) \end{aligned}$$

for some negligible function $\text{negl}(\cdot)$. Combining this with the false trace guarantee, we have

$$\Pr[\text{Cor-Tr}] \geq \Pr[\text{Good-Decoder}] - \text{negl}(\lambda).$$

□

9 Bounded Trace and Revoke with Embedded identity

In this section we show how to construct a bounded (secret/public tracing)-EITR scheme from an indexed (secret/public tracing)-EITR scheme. The construction and security analysis in this section is an adaptation of ([GKW19], Section 9) with appropriate modifications to incorporate the revocation list L . We present the construction and proofs for the secret trace setting primarily and also outline the differences in the public trace setting simultaneously.

9.1 Construction

Let $\text{Ind-TR} = (\text{Ind.Setup}, \text{Ind.KeyGen}, \text{Ind.Enc}, \text{Ind.Dec}, \text{Ind.Trace})$ be an indexed (secret/public tracing)-EITR system. Let $\text{Sig} = (\text{Sig.KeyGen}, \text{Sig.Sign}, \text{Sig.Verify})$ be a signature scheme that satisfies unforgeability with signature space $\{0, 1\}^{\ell_s}$. We let n_{bd} denote the bound on the number of key queries that the adversary can make. For our purpose, we can assume $n_{\text{bd}} \leq 2^\lambda$. We construct a bounded (secret/public tracing)-EITR scheme as follows:

$\text{Setup}(1^\lambda, 1^\kappa, n_{\text{bd}}) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm does the following:

1. Set $n_{\text{index}} = 2n_{\text{bd}}^2$.
2. For $j = 1$ to λ , sample $(\text{Ind.mpk}_j, \text{Ind.msk}_j) \leftarrow \text{Ind.Setup}(1^\lambda, 1^{\kappa'}, n_{\text{index}})$, where κ' is $\kappa + \ell_s$.
3. Sample $(\text{sig.sk}, \text{sig.vk}) \leftarrow \text{Sig.KeyGen}(1^\lambda)$.
4. Output $\text{mpk} = (\text{sig.vk}, \{\text{Ind.mpk}_j\}_{j \in [\lambda]})$ and $\text{msk} = (\text{sig.vk}, \text{sig.sk}, \{\text{Ind.mpk}_j, \text{Ind.msk}_j\}_{j \in [\lambda]})$.

$\text{KeyGen}(\text{msk}, \text{lb}, \text{id}) \rightarrow \text{sk}_{\text{lb}, \text{id}}$. The key generation algorithm does the following:

1. Parse msk as $(\text{sig.vk}, \text{sig.sk}, \{\text{Ind.mpk}_j, \text{Ind.msk}_j\}_{j \in [\lambda]})$.
2. Compute $\sigma = \text{Sig.Sign}(\text{sig.sk}, \text{id})$ and let $\text{id}' = (\text{id}, \sigma)$.
3. For $j = 1$ to λ , do the following:
 - (a) Sample $i_j \leftarrow [2n_{\text{bd}}^2]$.
 - (b) $\text{Ind.sk}_{\text{lb}, \text{id}, j} \leftarrow \text{Ind.KeyGen}(\text{Ind.msk}_j, \text{lb}, \text{id}', i_j)$.
4. Return $\text{sk}_{\text{lb}, \text{id}} = \{\text{Ind.sk}_{\text{lb}, \text{id}, j}\}_{j \in [\lambda]}$.

$\text{Enc}(\text{mpk}, m, L) \rightarrow \text{ct}$. The encryption algorithm does the following:

1. Parse mpk as $(\text{sig.vk}, \{\text{Ind.mpk}_j\}_{j \in [\lambda]})$.
2. For $j = 1$ to $\lambda - 1$, randomly sample $r_j \leftarrow \mathcal{M}$ and set $r_\lambda = m \oplus r_1 \oplus \dots \oplus r_{\lambda-1}$.
3. For $j = 1$ to λ , compute $\text{Ind.ct}_j \leftarrow \text{Ind.Enc}(\text{Ind.mpk}_j, r_j, L)$.
4. Return $\text{ct} = \{\text{Ind.ct}_j\}_{j \in [\lambda]}$.

$\text{Dec}(\text{sk}_{\text{lb}, \text{id}}, \text{ct}, L) \rightarrow m'$. The decryption algorithm does the following:

1. Parse $\text{sk}_{\text{lb}, \text{id}} = \{\text{Ind.sk}_{\text{lb}, \text{id}, j}\}_{j \in [\lambda]}$ and $\text{ct} = \{\text{Ind.ct}_j\}_{j \in [\lambda]}$.
2. For $j = 1$ to λ , compute $r'_j = \text{Ind.Dec}(\text{Ind.sk}_{\text{lb}, \text{id}, j}, \text{Ind.ct}_j, L)$.
3. If any of the decryption fails then output \perp , else output $m' = r'_1 \oplus \dots \oplus r'_\lambda$.

$\text{Trace}^D(\text{tk}, y, m_0, m_1, L) \rightarrow T^{\text{final}}$. The trace algorithm uses two algorithms Bnd-isGoodDecoder and Bnd-Subtrace defined in Figures 9 and 10, respectively as subroutines and is defined as follows:

1. Parse tk as $\text{msk} = (\text{sig.vk}, \text{sig.sk}, \{\text{Ind.mpk}_j, \text{Ind.msk}_j\}_{j \in [\lambda]})$ and let $\text{Ind.tk}_j = \text{Ind.msk}_j$ for $j \in [\lambda]$. (In the public trace setting, $\text{tk} = \text{mpk}$ and $\text{Ind.tk}_j = \text{Ind.mpk}_j$).
2. Set $j = 1$.
3. Set $\text{flag} = 0$. For $\text{itr} = 1$ to $\lambda \cdot y$, do the following
 - (a) Choose a random message $r \leftarrow \mathcal{M}$.
 - (b) Run Bnd-isGoodDecoder as $\text{flag} \leftarrow \text{Bnd-isGoodDecoder}^D(\{\text{Ind.mpk}_j\}_{j \in [\lambda]}, 1^y, m_0, m_1, r, L, j)$.
 - (c) If $\text{flag} = 1$, break. Else, continue.
4. If $\text{flag} = 1$, run Bnd-Subtrace as $T \leftarrow \text{Bnd-Subtrace}^D(\{\text{Ind.mpk}_j, \text{Ind.tk}_j\}_{j \in [\lambda]}, 1^y, m_0, m_1, r, L, j)$. Else, set $T = \phi$.
5. If $T = \phi$ and $j < \lambda$, set $j = j + 1$ and go to step 3. Otherwise do the following:
 - Set $T^{\text{temp}} = \phi$. For each $\text{id}' = (\text{id}, \sigma) \in T$, if $\text{Sig.Verify}(\text{sig.vk}, \text{id}, \sigma) = 1$, add id to T^{temp} . Concretely,

$$T^{\text{temp}} = \{\text{id} : \exists \sigma \text{ s.t. } (\text{id}, \sigma) \in T \text{ and } \text{Sig.Verify}(\text{sig.vk}, \text{id}, \sigma) = 1\}.$$

- Recall the function map $: \mathcal{ID} \rightarrow \mathcal{L}$, that maps a given identity id to its corresponding label lb (Remark 7.4). For each $\text{id} \in T^{\text{temp}}$, if $\text{map}(\text{id}) \in L$, then set $T^{\text{temp}} = T^{\text{temp}} \setminus \{\text{id}\}$.
- Set $T^{\text{final}} = T^{\text{temp}}$.
- If $T^{\text{final}} = \phi$ and $j < \lambda$, set $j = j + 1$ and go to step 3. Otherwise exit and return T^{final} .

Correctness. We prove that the above construction of bounded (secret/public tracing)-EITR scheme satisfies correctness (Def. 7.1) via the following theorem.

Theorem 9.1. *Assume Ind-TR is a correct indexed (secret/public tracing)-EITR scheme then the above construction of bounded (secret/public tracing)-EITR scheme is correct.*

Proof. We have, as per the construction, that any message m is split in λ components r_1, \dots, r_λ such that $r_1 \oplus \dots \oplus r_\lambda = m$. Then, from the correctness of Ind-TR , we get $r'_k = r_k$ for all $k \in [\lambda]$, where $r'_k \leftarrow \text{Ind.Dec}(\text{Ind.sk}_{|\text{b}, \text{id}, k}, \text{Ind.ct}_k, L)$, as long as $\text{lb} \notin L$. Hence, the decryption algorithm correctly outputs m as $r_1 \oplus \dots \oplus r_\lambda$. \square

Efficiency. We can instantiate the above construction by the indexed public/secret tracing-EITR scheme in Sec. 8.1. The above construction is basically λ times repetition of the underlying indexed EITR scheme. Additional overhead is induced due to the usage of the signature scheme, where identity becomes longer by $\ell_s = \text{poly}(\lambda)$ bit and the master public key is longer by $|\text{sig.vk}| = \text{poly}(\lambda)$ bit.²⁶ These changes do not alter the dependency on $|\text{id}|$ and $|\text{lb}|$ of the parameter size. Therefore, the parameter size is as follows.

²⁶Parameter sizes of most signature schemes depend on the length of the message space, which is $|\text{id}|$ in our case. However, this dependency can be removed by hashing the message before signing using the collision resistant hash functions.

Algorithm Bnd-isGoodDecoder^D(key, 1^y, m₀, m₁, r, L, i)

Inputs: keys key = {Ind.mpk_j}_{j∈[λ]}, parameter y, messages m₀, m₁, r, revocation list L and position-index i ∈ [λ].

Output : 0/1.

1. Set count = 0. Let $\epsilon = 1/y$.
2. For $j = 1$ to $\lambda \cdot y$:
 - Split r in $\lambda - 1$ shares. That is, randomly sample $\lambda - 1$ messages, $r_1, \dots, r_{i-1}, r_{i+1}, \dots, r_\lambda$ such that $\bigoplus_{k \in [\lambda] \setminus \{i\}} r_k = r$.
 - Sample $b \leftarrow \{0, 1\}$. For $k \in [\lambda] \setminus \{i\}$, compute $\text{Ind.ct}_k \leftarrow \text{Ind.Enc}(\text{Ind.mpk}_k, r_k, L)$. Compute $\text{Ind.ct}_i \leftarrow \text{Ind.Enc}(\text{Ind.mpk}_i, r \oplus m_b, L)$.
 - Query D on $\text{ct} = (\text{Ind.ct}_1, \dots, \text{Ind.ct}_\lambda)$. Let b' be D 's response.
 - If $b' = b$, set $\text{count} = \text{count} + 1$.
3. If $\text{count}/(\lambda \cdot y) \geq 1/2 + \epsilon/3$, then output 1, else output 0.

Figure 9: Algorithm Bnd-isGoodDecoder

Algorithm Bnd-Subtrace^D(key, 1^y, m₀, m₁, r, L, i)

Inputs: keys key = {Ind.mpk_j, Ind.tk_j}_{j∈[λ]}, parameter y, messages m₀, m₁, r, revocation list L and index-position i ∈ [λ].

Output : $T \subseteq \{0, 1\}^\kappa$

1. Define oracle $\tilde{D}[\{\text{Ind.mpk}_j\}_{j \in [\lambda]}, r, L, i]$ as in Figure 11.
2. Output $T \leftarrow \text{Ind.Trace}^{\tilde{D}}(\text{Ind.tk}_i, 4y, m_0 \oplus r, m_1 \oplus r, L)$.

Figure 10: Algorithm Bnd-Subtrace

Algorithm $\tilde{D}^D[\text{key}, r, L, i]$

Hardwired values: keys $\text{key} = \{\text{Ind.mpk}_j\}_{j \in [\lambda]}$, message r , revocation list L and index-position $i \in [\lambda]$.

Inputs: Ind.ct .

Output: $0/1$

Upon input Ind.ct , the \tilde{D} oracle does the following:

- Shares r in $\lambda - 1$ components as follows: it chooses $\lambda - 1$ random messages r_k for $k \in [\lambda] \setminus \{i\}$, such that $\bigoplus_{k \in [\lambda] \setminus \{i\}} r_k = r$.
- For $k \in [\lambda] \setminus \{i\}$, computes $\text{Ind.ct}_k = \text{Ind.Enc}(\text{Ind.mpk}_k, r_k, L)$.
- Sets $\text{ct}_{\text{bd}} = (\text{Ind.ct}_1, \dots, \text{Ind.ct}_{i-1}, \text{Ind.ct}, \text{Ind.ct}_{i+1}, \dots, \text{Ind.ct}_\lambda)$.
- Queries oracle D as $b' \leftarrow D(\text{ct}_{\text{bd}})$.
- Outputs b' .

Figure 11: Oracle \tilde{D}

Secret Tracing Setting. In the secret tracing setting, we have

$$|\text{mpk}|, |\text{ct}|, |\text{sk}| = \text{poly}(\lambda, |\text{id}|, |\text{lb}|).$$

Public Tracing Setting. In the public tracing setting, we have

$$|\text{mpk}|, |\text{ct}| = \text{poly}(\lambda, |\text{lb}|), |\text{sk}| = \text{poly}(\lambda, |\text{id}|, |\text{lb}|).$$

9.2 Security

In this section, we prove that our construction of bounded (secret/public tracing)-EITR scheme is secure.

IND-CPA security.

Theorem 9.2. *If Ind-TR is a (selective/adaptive)-IND-CPA secure indexed (secret/public tracing)-EITR scheme, then the above construction of bounded (secret/public tracing)-EITR scheme is (selective/adaptive)IND-CPA secure.*

Proof. We show that if there exists a PPT adversary that breaks the IND-CPA security of the bounded EITR scheme, then we can use \mathcal{A} to build a PPT algorithm \mathcal{B} that breaks IND-CPA security of the underlying Ind-TR scheme. \square

The reduction is as follows:

1. \mathcal{B} gets $(1^\kappa, 1^{n_{\text{bd}}}, L)$ from the adversary \mathcal{A} . It sets $n_{\text{index}} = 2n_{\text{bd}}^2$ and sends $(1^{\kappa+\ell_s}, 1^{n_{\text{index}}}, L)$ to the Ind-TR challenger. The Ind-TR challenger returns Ind.mpk . (In the public trace setting, \mathcal{A} outputs L adaptively, along with the challenge messages in Step 4).

2. \mathcal{B} sets $\text{Ind.mpk}_1 = \text{Ind.mpk}$, generates $(\text{Ind.mpk}_j, \text{Ind.msk}_j) \leftarrow \text{Ind.Setup}(1^\lambda, 1^{\kappa+\ell_s}, n_{\text{index}})$ for $j \in \{2, \dots, \lambda\}$ and $(\text{sig.sk}, \text{sig.vk}) \leftarrow \text{Sig.KeyGen}(1^\lambda)$. It sends $\text{mpk} = (\text{sig.vk}, \{\text{Ind.mpk}_j\}_{j \in [\lambda]})$ to \mathcal{A} .
3. When \mathcal{A} makes a key query (lb, id) , \mathcal{B} samples $i_1 \leftarrow [n_{\text{index}}]$, computes id' as in the construction and sends a key query $(\text{lb}, \text{id}', i_1)$ to the Ind-TR challenger. The Ind-TR challenger returns $\text{Ind.sk}_{\text{lb}, \text{id}} \leftarrow \text{Ind.KeyGen}(\text{Ind.msk}, \text{lb}, \text{id}', i_1)$. \mathcal{B} sets $\text{Ind.sk}_{\text{lb}, \text{id}, 1} = \text{Ind.sk}_{\text{lb}, \text{id}}$, computes $\text{sk}_{\text{lb}, \text{id}, j}$ for $j \in \{2, \dots, \lambda\}$ itself and returns $\text{sk}_{\text{lb}, \text{id}} = \{\text{Ind.sk}_{\text{lb}, \text{id}, j}\}_{j \in \lambda}$ to \mathcal{A} .
4. When \mathcal{A} sends the challenge messages (m_0, m_1) for the challenge query, \mathcal{B} samples $r_j \leftarrow \mathcal{M}$ for $j \in \{2, \dots, \lambda\}$, sets $m'_b = \bigoplus_{j>1} r_j \oplus m_b$, sends (m'_0, m'_1) for the challenge query to the Ind-TR challenger and gets back Ind.ct . \mathcal{B} sets $\text{ct}_1 = \text{Ind.ct}$, computes $\text{ct}_j \leftarrow \text{Ind.Enc}(\text{Ind.mpk}_j, r_j, L)$ for $j \in \{2, \dots, \lambda\}$ and sends $\text{ct} = \{\text{ct}_j\}_{j \in \lambda}$ to \mathcal{A} .
5. \mathcal{A} outputs a bit b' , \mathcal{B} returns b' to the challenger.

Observe that \mathcal{B} issues key queries to the Ind-TR challenger only when there is a key query from \mathcal{A} . From the admissibility of \mathcal{A} , for any key query of the form $(\text{lb}, \text{id}, i_1)$ that \mathcal{B} issues $\text{lb} \in L$. So, \mathcal{A} wins the IND-CPA security game of the BD-TR scheme with advantage ϵ , then \mathcal{B} also wins the IND-CPA security game of the underlying Ind-TR scheme with the same advantage.

Secure Tracing Analysis. Now we show that our construction satisfies false trace and correct trace guarantees.

False Trace Guarantee. First, we show that the probability of false trace in our scheme is negligible in the security parameter via the following theorem.

Theorem 9.3. *Assume that Sig is unforgeable, then our construction of bounded (secret/public tracing)-EITR scheme satisfies the (selective/adaptive) false trace guarantee (Def. 7.8/Def. 7.9), even if the adversary makes unbounded polynomial number of key queries.*

Proof. Let us first setup some notations. For $\text{lb} \in \mathcal{L}$, $\text{id} \in \mathcal{ID}$ and a revocation list L , let

$$\begin{aligned} S_{\mathcal{ID}} &= \{\text{id} : (\text{lb}, \text{id}) \in S\}, \\ L_{\mathcal{ID}} &= \{\text{id} : (\text{lb}, \text{id}) \in S \wedge \text{lb} \in L\}, \\ T_{\text{lb}}^{\text{final}} &= \{\text{map}(\text{id}) : \text{id} \in T^{\text{final}}\} \end{aligned}$$

where map is as defined in Remark 7.4 and T^{final} is as in the construction.

False trace happens when $T^{\text{final}} \not\subseteq S_{\mathcal{ID}}$ or $T_{\text{lb}}^{\text{final}} \cap L \neq \phi$. Observe that $T_{\text{lb}}^{\text{final}} \cap L = \phi$ by definition. So, all we need to argue is that $T^{\text{final}} \subseteq S_{\mathcal{ID}}$.

Recall that the tracing algorithm uses Bnd-Subtrace algorithm to find a set $T = \{(\text{id}_k, \sigma_k)\}_k$. Identity id_k is added to the set T^{temp} and then to set T^{final} only if $\text{Sig.Verify}(\text{sig.vk}, \text{id}_k, \sigma_k) = 1$. Next we show that if there is an adversary \mathcal{A} who outputs a decoder D along with messages m_0, m_1 and revocation list L such that the tracing algorithm outputs T^{final} where $T^{\text{final}} \not\subseteq S_{\mathcal{ID}}$, then there exists a reduction \mathcal{B} against unforgeability of signature scheme Sig. The reduction is defined as follows:

Upon receiving sig.vk from the Sig challenger, \mathcal{B} does the following:

1. Run \mathcal{A} on input 1^λ to obtain $(1^\kappa, 1^{n_{\text{index}}}, L)$. (In the public trace setting, it outputs L adaptively along with the challenge messages).
2. Samples $(\text{Ind.mpk}_j, \text{Ind.msk}_j) \leftarrow \text{Ind.Setup}(1^\lambda, 1^{\kappa'}, n_{\text{index}})$ for $j \in [\lambda]$ and sends $\text{mpk} = (\text{sig.vk}, \{\text{mpk}_j\}_{j \in [\lambda]})$ to \mathcal{A} .
3. Whenever \mathcal{A} issues a key query (lb, id) , \mathcal{B} sends id to the Sig challenger for signature. The Sig challenger returns $\sigma = \text{Sig}(\text{sig.sk}, \text{id})$. \mathcal{B} generates the secret key $\text{sk}_{\text{lb}, \text{id}}$ using $\text{id}' = (\text{id}, \sigma)$ as in the construction and sends $\text{sk}_{\text{lb}, \text{id}}$ to \mathcal{A} .
4. In the end, \mathcal{A} outputs a decoder D , messages m_0, m_1 .
5. \mathcal{B} runs the trace algorithm with the help of the decoder D and gets a set T^{final} .
6. If there exists an $\text{id}^* \in T^{\text{final}}$ such that $\text{id}^* \notin S_{TD}$, then \mathcal{B} returns a forgery for id^* to Sig challenger as follows: $\text{id}^* \in T^{\text{final}}$ implies that there exists σ^* , such that $(\text{id}^*, \sigma^*) \in T$, and $\text{Sig.Verify}(\text{sig.vk}, \text{id}^*, \sigma^*) = 1$. \mathcal{B} returns (id^*, σ^*) as a forgery to the Sig challenger.

Note that since $\text{id}^* \notin S_{TD}$, \mathcal{B} must not have queried a signature on id^* to the Sig challenger and hence (id^*, σ^*) is a valid forgery. If the false tracing happens with non-negligible probability ϵ , then \mathcal{B} also wins the unforgeability game with probability ϵ .

□

Correct Trace Guarantee. Recall that in the experiment for correct tracing, the adversary first sends $(1^\kappa, 1^{n_{\text{bd}}})$ and a revocation list L (in the public trace setting, it outputs L adaptively), then the challenger sends the public key to the adversary. After that the adversary is allowed to make at most n_{bd} key queries. Let $S = \{(\text{lb}, \text{id})\}$ be the set of label-identity pairs for which the adversary issues key queries. At the end, the adversary outputs a decoder box D along with two messages m_0 and m_1 . If D is a good decoder then for correct tracing we want that the tracing algorithm outputs non empty set of traitors $T^{\text{final}} \subseteq S_{TD} \setminus L_{TD}$. We prove the correct trace guarantee of our scheme via the following theorem.

Theorem 9.4. *Assume that the underlying (secret/public tracing)-Ind-TR scheme satisfies (selective/adaptive) correct trace guarantee (Def. 7.6 /Def. 7.5) and let n_{bd} be the bound on the number of key queries for an admissible adversary, then our bounded (secret/public tracing)-EITR scheme also satisfies the (selective/adaptive) correct trace guarantee (Def. 7.9 /Def. 7.8).*

Proof. Similar to [GKW19] we begin with defining some events of interest. We modify the definition of some events to take into account the constraint that $T^{\text{final}} \subseteq S_{TD} \setminus L_{TD}$. We drop the subscripts \mathcal{A}, ϵ and security parameter λ in the following to keep the notations simple.

- Event Admissible-Adversary: It is defined as the event that the adversary \mathcal{A} makes at most n_{bd} key queries.
- Event Tr (Tracing without correctness): Similar to Corr-Tr, except that we don't need $T^{\text{final}} \subseteq S_{TD} \setminus L_{TD}$, i.e., Tr is the event that $T^{\text{final}} \neq \emptyset$ occurs. Denote $\text{Pr-Tr} := \text{Pr}[\text{Tr}]$.
- Event Dist-Indx (Position with distinct indices for each key): Defined as the event that there exists $i \in [\lambda]$ such that the i -th index of each key is distinct.

- Event Dist-Idx_i : It is defined as the event that $i \in [\lambda]$ is the first position such that the i -th index of each key is distinct. By definition, Dist-Idx_i are disjoint events for all $i \in [\lambda]$ and $\cup_{i \in [\lambda]} \text{Dist-Idx}_i = \text{Dist-Idx}$.
- Event Tr_i (Tracing without correctness in i -th iteration): Let T_i denote the set of (identity, signature) pairs traced in the i -th iteration. The event Tr_i happens if T_i is non-empty.
- Event Corr-Tr-Sig_i (Tracing with same signature as that received in key): The event that T_i is not empty and for all $(\text{id}, \sigma) \in T_i$, (lb, id) was queried for a key and $\text{lb} \notin L$, i.e. $\text{id} \in S_{\text{ID}} \setminus L_{\text{ID}}$ and that the key generation oracle output $\text{Ind.sk}_{\text{lb}, \text{id}, i} \leftarrow \text{Ind.KeyGen}(\text{Ind.msk}_i, \text{lb}, (\text{id}, \sigma), j)$ for some j .
- Event $\text{Found-Good-}r_i$: This event occurs if flag is set to 1 in the i -th iteration of the tracing algorithm.
- Event $\text{Good-}\tilde{D}_i$ (Good decoder \tilde{D} during the Bnd-Subtrace routine execution in i -th iteration): It is defined as the event that in the i -th iteration, the execution reaches step 3 (that is, it found a ‘good’ r in the i -th iteration), and the decoder \tilde{D} constructed is an $\epsilon/4$ good decoder for distinguishing messages $m_0 \oplus r$ and $m_1 \oplus r$. Note that if no good r is found in step 3, $\text{Good-}\tilde{D}_i$ is said to not have happened.

With the above events, the correctness of tracing is argued via the following series of inequalities:

$$\Pr\text{-Corr-Tr}(\lambda) \geq \Pr\text{-Tr} - \text{negl} \quad (9.1)$$

$$\geq \Pr[\text{Tr} \wedge \text{Dist-Idx}] - \text{negl} \quad (9.2)$$

$$= \sum_{i \in [\lambda]} \Pr[\text{Tr} \wedge \text{Dist-Idx}_i] - \text{negl} \quad (9.3)$$

$$\geq \sum_{i \in [\lambda]} \Pr[\text{Corr-Tr-Sig}_i \wedge \text{Dist-Idx}_i] - \text{negl} \quad (9.4)$$

$$\geq \sum_{i \in [\lambda]} \Pr[\text{Good-}\tilde{D}_i \wedge \text{Admissible-Adversary} \wedge \text{Dist-Idx}_i] - \text{negl} \quad (9.5)$$

$$\geq \sum_{i \in [\lambda]} \Pr[\text{Good-}\tilde{D}_i \wedge \text{Found-Good-}r_i \wedge \text{Good-Decoder} \wedge \text{Admissible-Adversary} \wedge \text{Dist-Idx}_i] - \text{negl} \quad (9.6)$$

$$\geq \sum_{i \in [\lambda]} \Pr[\text{Found-Good-}r_i \wedge \text{Good-Decoder} \wedge \text{Admissible-Adversary} \wedge \text{Dist-Idx}_i] - \text{negl} \quad (9.7)$$

$$\geq \sum_{i \in [\lambda]} \Pr[\text{Good-Decoder} \wedge \text{Admissible-Adversary} \wedge \text{Dist-Idx}_i] - \text{negl} \quad (9.8)$$

$$= \Pr[\text{Good-Decoder} \wedge \text{Admissible-Adversary} \wedge \text{Dist-Idx}] - \text{negl} \quad (9.9)$$

$$\geq \Pr[\text{Good-Decoder} \wedge \text{Admissible-Adversary}] - \text{negl} \quad (9.10)$$

Explanation for each of the inequalities is the same as in [GKW19] and is omitted. Here, we argue the transitions from equations (9.3) to (9.4) and (9.4) to (9.5) only, since our definition of event Corr-Tr-Sig_i is slightly modified.

- Transition from (9.3) to (9.4) follows from the observation that the event Corr-Tr-Sig_i implies the event Tr because by definition, if Corr-Tr-Sig_i happens then T_i is non empty

and for each (id, σ) pair in T_i , $\text{Sig.Verify}(\text{sig.vk}, id, \sigma)$ verifies and $id \in S_{TD} \setminus L_{TD}$. Hence, the set of traitors T^{final} obtained from T_i will also be non empty.

- Transition from (9.4) to (9.5) is argued via the following claim:

Claim 9.5. *Assume that the underlying Ind-TR scheme satisfies correct trace guarantee. Then for all $i \in [\lambda]$*

$$\Pr[\text{Corr-Tr-Sig}_i \wedge \text{Dist-Idx}_i] \geq \Pr[\text{Good-}\tilde{D}_i \wedge \text{Admissible-Adversary} \wedge \text{Dist-Idx}_i] - \text{negl}.$$

Proof. We prove the claim by contradiction. We assume that there exists an adversary \mathcal{A} who outputs a good decoder D along with messages m_0, m_1 and a revocation list L such that $\Pr[\text{Good-}\tilde{D}_i \wedge \text{Admissible-Adversary} \wedge \text{Dist-Idx}_i] - \Pr[\text{Corr-Tr-Sig}_i \wedge \text{Dist-Idx}_i]$ is non-negligible. Then we use \mathcal{A} to build an adversary \mathcal{B} against correct trace guarantee of Ind-TR, defined as follows:

1. At the start of the game, \mathcal{A} outputs $(1^\kappa, 1^{n_{\text{bd}}}, L)$. (In the public trace setting, it outputs L adaptively along with the challenge message).
2. \mathcal{B} sets $n_{\text{index}} = 2n_{\text{bd}}^2$, sends $(1^{\kappa+\ell_s}, 1^{n_{\text{index}}}, L)$ to the Ind-TR challenger and gets back Ind.mpk .
3. \mathcal{B} sets $\text{Ind.mpk}_i = \text{Ind.mpk}$ and does the following:
 - Samples $(\text{sig.sk}, \text{sig.vk}) \leftarrow \text{Sig.KeyGen}(1^\lambda)$.
 - For $k \in [\lambda] \setminus \{i\}$, sample $(\text{Ind.mpk}_k, \text{Ind.msk}_k) \leftarrow \text{Ind.Setup}(1^{\kappa+\ell_s}, 1^{n_{\text{index}}})$.
 - Set $\text{mpk} = (\text{sig.vk}, \{\text{Ind.mpk}_k\}_{k \in [\lambda]})$ and sends mpk to \mathcal{A} .
4. \mathcal{A} issues at most n_{bd} key queries. For each key query (lb, id) by \mathcal{A} , \mathcal{B} does the following:
 - Computes $\text{Sig.Sign}(\text{sig.sk}, \text{id})$ and sets $\text{id}' = (\text{id}, \sigma)$.
 - For $k \in [\lambda]$, randomly chooses $j_k \leftarrow [n_{\text{index}}]$ and sends a key query $(\text{lb}, \text{id}', j_i)$ to the Ind-TR challenger and gets back $\text{Ind.sk}_{\text{lb}, \text{id}, j_i}$. \mathcal{B} sets $\text{Ind.sk}_{\text{lb}, \text{id}, i} = \text{Ind.sk}_{\text{lb}, \text{id}, j_i}$ and for $k \in [\lambda] \setminus \{i\}$, it computes $\text{Ind.sk}_{\text{lb}, \text{id}, k} \leftarrow \text{Ind.KeyGen}(\text{Ind.msk}_k, \text{lb}, \text{id}', j_k)$.
 - Sends $\text{sk}_{\text{lb}, \text{id}} = (\text{Ind.sk}_{\text{lb}, \text{id}, 1}, \dots, \text{sk}_{\text{lb}, \text{id}, \lambda})$ to \mathcal{A} .
5. If i is not the first position for which the i -th position indices (j_i 's) are distinct for all the key queries, then \mathcal{B} outputs a random decoder and quits the game.
6. In the end, \mathcal{A} outputs a decoder D along with messages m_0, m_1 .
7. \mathcal{B} runs $\text{Bnd-isGoodDecoder}(\{\text{Ind.mpk}\}_{j \in [\lambda]}, 1^y, m_0, m_1, L, r)$ for uniformly and independently sampled r until it finds a r for which Bnd-isGoodDecoder algorithm outputs 1. If Bnd-isGoodDecoder algorithm does not output 1 even after $\lambda \cdot y$ attempts, then \mathcal{B} outputs a random decoder and quits.
8. \mathcal{B} constructs decoder \tilde{D} as defined in Figure 11 and sets $\tilde{m}_b = m_b \oplus r$ for $b \in \{0, 1\}$ and sends $(\tilde{D}, \tilde{m}_0, \tilde{m}_1, L)$ to the Ind-TR challenger.

Now let us analyze the probability that \mathcal{B} outputs a $1/4y$ good decoder box. This happens if (i) i is the first position such that the i -th position indices are different for all the key queries (ii) Bnd-isGoodDecoder outputs 1 for some r and \tilde{D} is $\epsilon/4 = 1/4y$ good decoder for distinguishing $m_0 \oplus r, m_1 \oplus r$. First event is same as Dist-Idx_i and second event is same as $\text{Good-}\tilde{D}_i$. Hence,

the probability that \mathcal{B} outputs a $1/4\gamma$ good decoder box is $\Pr[\text{Good-}\tilde{\mathcal{D}}_i \wedge \text{Dist-Ind}\mathcal{X}_i]$. Then, by correct trace guarantee of Ind-TR,

$$\begin{aligned} \Pr[\text{Corr-Tr-Sig}_i \wedge \text{Dist-Ind}\mathcal{X}_i] &\geq \Pr[\text{Good-}\tilde{\mathcal{D}}_i \wedge \text{Dist-Ind}\mathcal{X}_i] - \text{negl} \\ &\geq \Pr[\text{Good-}\tilde{\mathcal{D}}_i \wedge \text{Admissible-Adversary} \wedge \text{Dist-Ind}\mathcal{X}_i] - \text{negl}. \end{aligned}$$

This contradicts our assumption that $\Pr[\text{Good-}\tilde{\mathcal{D}}_i \wedge \text{Admissible-Adversary} \wedge \text{Dist-Ind}\mathcal{X}_i] - \Pr[\text{Corr-Tr-Sig}_i \wedge \text{Dist-Ind}\mathcal{X}_i]$ is non-negligible, hence the proof. \square

\square

10 Unbounded Trace and Revoke with Embedded Identities

In this section we show how to construct unbounded (secret/public tracing)-EITR scheme from a bounded (secret/public tracing)-EITR scheme. The transformation technique and the security analysis in this section is adapted from [GKW19] with modifications to incorporate the revocation list. We present the construction and proofs for *secret-key* trace setting primarily and also outline the differences in the public-key trace setting simultaneously.

10.1 Construction

Let $\text{BD-TR} = (\text{BD.Setup}, \text{BD.KeyGen}, \text{BD.Enc}, \text{BD.Dec}, \text{BD.Trace})$ be a bounded (secret/public tracing)-EITR scheme for identity space $\mathcal{ID} = \{0, 1\}^\kappa$. We construct an unbounded (secret/public tracing)-EITR scheme as follows.

$\text{Setup}(1^\lambda, 1^\kappa) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm does the following:

1. For $j = 1$ to λ , sample $(\text{BD.mpk}_j, \text{BD.msk}_j) \leftarrow \text{BD.Setup}(1^\lambda, 1^\kappa, n_{\text{bd}} = 2^j)$.
2. Output $\text{mpk} = \{\text{BD.mpk}_j\}_{j \in [\lambda]}$ and $\text{msk} = \{\text{BD.mpk}_j, \text{BD.msk}_j\}_{j \in [\lambda]}$.

$\text{KeyGen}(\text{msk}, \text{lb}, \text{id}) \rightarrow \text{sk}_{\text{lb}, \text{id}}$. The KeyGen algorithm does the following:

1. Parse $\text{msk} = \{\text{BD.mpk}_j, \text{BD.msk}_j\}_{j \in [\lambda]}$.
2. For $j = 1$ to λ , it computes $\text{BD.sk}_j \leftarrow \text{BD.KeyGen}(\text{BD.msk}_j, \text{lb}, \text{id})$.
3. Returns $\text{sk}_{\text{lb}, \text{id}} = \{\text{BD.sk}_j\}_{j \in [\lambda]}$.

$\text{Enc}(\text{mpk}, m, L) \rightarrow \text{ct}$. The encryption algorithm does the following:

1. Parse $\text{mpk} = \{\text{BD.mpk}_j\}_{j \in [\lambda]}$.
2. Secret share m in λ shares as follows. For $j = 1$ to $\lambda - 1$, randomly sample $r_j \leftarrow \mathcal{M}$ and set $r_\lambda = m \oplus r_1 \oplus \dots \oplus r_{\lambda-1}$.
3. For $j = 1$ to λ , compute $\text{BD.ct}_j = \text{BD.Enc}(\text{BD.mpk}_j, r_j, L)$.
4. Output $\text{ct} = \{\text{BD.ct}_j\}_{j \in [\lambda]}$.

$\text{Dec}(\text{sk}_{\text{lb}, \text{id}}, \text{ct}, L) \rightarrow m'$. The decryption algorithm does the following:

1. Parse $\text{sk}_{\text{lb}, \text{id}} = \{\text{BD.sk}_j\}_{j \in [\lambda]}$ and $\text{ct} = \{\text{BD.ct}_j\}_{j \in [\lambda]}$.

Algorithm isGoodDecoder^D(key, 1^y, m₀, m₁, r, L, i)

Inputs: keys key = {BD.mpk_j}_{j∈[λ]}, parameter y, messages m₀, m₁, r, revocation list L and a position-index i ∈ [λ].

Output : 0/1.

1. Set count = 0. Let $\epsilon = 1/y$.
2. For $j = 1$ to $\lambda \cdot y$:
 - Sample $\lambda - 1$ messages, $r_1, \dots, r_{i-1}, r_{i+1}, \dots, r_\lambda$ randomly such that $\bigoplus_{k \in [\lambda] \setminus \{i\}} r_k = r$.
 - Sample $b \leftarrow \{0, 1\}$, and compute ciphertexts $\text{BD.ct}_k = \text{BD.Enc}(\text{BD.mpk}_k, r_k, L)$ for $k \in [\lambda] \setminus \{i\}$ and $\text{BD.ct}_i = \text{BD.Enc}(\text{BD.mpk}_i, r \oplus m_b, L)$.
 - Query D on $\text{ct} = (\text{BD.ct}_1, \dots, \text{BD.ct}_\lambda)$. Let b' be the response of D .
 - If $b' = b$, set $\text{count} = \text{count} + 1$.
3. If $\text{count}/(\lambda \cdot y) \geq 1/2 + \epsilon/3$, then output 1, else output 0.

Figure 12: Algorithm isGoodDecoder for Unbounded EITR

2. For $j = 1$ to λ , compute $r'_j = \text{BD.Dec}(\text{BD.sk}_j, \text{BD.ct}_j, L)$.
3. If any of the decryption fails then output $m' = \perp$, else output $m' = \bigoplus_{j \in [\lambda]} r'_j$.

$\text{Trace}^D(\text{tk}, y, Q_{\text{bd}}, m_0, m_1, L) \rightarrow T$. The trace algorithm uses two algorithms isGoodDecoder and SubTrace defined in Figures 12 and 13, respectively as subroutines. The tracing algorithm is as follows.

1. Parse tk as $\text{msk} = \{\text{BD.mpk}_j, \text{BD.msk}_j\}_{j \in [\lambda]}$ and let $\text{BD.tk}_j = \text{BD.msk}_j$ for $j \in [\lambda]$. (For the public trace setting $\text{tk} = \text{mpk}$ and $\text{BD.tk}_j = \text{BD.mpk}_j$)
2. Set $j = \lceil \log Q_{\text{bd}} \rceil$.
3. Set flag = 0. For $\text{itr} = 1$ to $\lambda \cdot y$, do the following
 - (a) Choose a random message $r \leftarrow \mathcal{M}$.
 - (b) Run isGoodDecoder as $\text{flag} \leftarrow \text{isGoodDecoder}^D(\{\text{BD.mpk}_j\}_{j \in [\lambda]}, 1^y, m_0, m_1, r, L, j)$
 - (c) If flag = 1, break. Else, continue.
4. If flag = 1, run SubTrace as $T \leftarrow \text{SubTrace}^D(\{\text{BD.mpk}_j, \text{BD.tk}_j\}_{j \in [\lambda]}, 1^y, m_0, m_1, r, L, j)$. Else, set $T = \phi$.
5. Output T .

Correctness. We show that the above construction of bounded (secret/public tracing)-EITR satisfies correctness (Def. 7.1) via the following theorem.

Theorem 10.1. *Assume BD-TR is a correct bounded (secret/public tracing)-EITR scheme then the above construction of unbounded (secret/public tracing)-EITR scheme is correct.*

Algorithm $\text{SubTrace}^D(\text{key}, 1^y, m_0, m_1, r, L, i)$

Inputs: keys $\text{key} = \{\text{BD.mpk}_j, \text{BD.tk}_j\}_{j \in [\lambda]}$, parameter y , messages m_0, m_1, r , revocation list L and a position-index $i \in [\lambda]$.

Output: $T \subseteq \{0, 1\}^\kappa$.

1. Define oracle $\tilde{D}[\{\text{BD.mpk}_j\}_{j \in [\lambda]}, r, L, i]$ as in Figure 14.
2. Output $T \leftarrow \text{BD.Trace}^{\tilde{D}}(\text{BD.tk}_i, 4y, m_0 \oplus r, m_1 \oplus r, L)$.

Figure 13: Algorithm: SubTrace for Unbounded EITR

Algorithm $\tilde{D}^D[\text{key}, r, L, i]$

Hardwired values: keys $\text{key} = \{\text{BD.mpk}_j\}_{j \in [\lambda]}$, message r , revocation list L and a position-index $i \in [\lambda]$.

Inputs: BD.ct .

Output: $0/1$

On input BD.ct , the \tilde{D} oracle does the following:

- It first shares r in $\lambda - 1$ components as follows: it chooses $\lambda - 1$ random messages r_k for $k \in [\lambda] \setminus \{i\}$, such that $\bigoplus_{k \in [\lambda] \setminus \{i\}} r_k = r$.
- It computes $\text{BD.ct}_k = \text{BD.Enc}(\text{BD.mpk}_k, r_k, L)$ for $k \in [\lambda] \setminus \{i\}$.
- Sets $\text{ct} = (\text{BD.ct}_1, \dots, \text{BD.ct}_{i-1}, \text{BD.ct}, \text{BD.ct}_{i+1}, \dots, \text{BD.ct}_\lambda)$.
- Outputs $b' \leftarrow D(\text{ct})$.

Figure 14: Oracle \tilde{D} for Unbounded EITR

Proof. For all $k \in [\lambda]$, if $\text{BD.ct}_k \leftarrow \text{BDEnc}(\text{BD.mpk}_k, r_k, L)$ and $\text{BD.sk}_k \leftarrow \text{BDKeyGen}(\text{BD.msk}_k, (\text{lb}, \text{id}))$, we have $r_k \leftarrow \text{BD.Dec}(\text{BD.sk}_k, \text{BD.ct}_k, L)$, as long as $\text{id} \notin L$, from the correctness of the underlying BD-TR scheme. Hence, the decryption of $\text{ct} = (\text{BD.ct}_1, \dots, \text{BD.ct}_\lambda)$ correctly outputs m as $r_1 \oplus \dots \oplus r_\lambda$. \square

Efficiency. We can instantiate the above construction by the bounded public/secret tracing-EITR scheme in Sec. 9. Since the above construction is simple λ times repetition of the underlying bounded EITR scheme, the parameter size of the scheme is as follows.

Secret Tracing Setting. In the secret tracing setting, we have

$$|\text{mpk}|, |\text{ct}|, |\text{sk}| = \text{poly}(\lambda, |\text{id}|, |\text{lb}|).$$

Public Tracing Setting. In the public tracing setting, we have

$$|\text{mpk}|, |\text{ct}| = \text{poly}(\lambda, |\text{lb}|), |\text{sk}| = \text{poly}(\lambda, |\text{id}|, |\text{lb}|).$$

10.2 Security

In this section, we prove that our construction of unbounded (secret/public tracing)-EITR scheme is secure.

IND-CPA security.

Theorem 10.2. *If the underlying BD-TR scheme is (selective/adaptive) IND-CPA secure bounded (secret/public tracing)-EITR scheme, then the above construction of unbounded (secret/public tracing)-EITR is (selective/adaptive) IND-CPA secure.*

Proof. We show that if there exists a PPT adversary that breaks the IND-CPA security of unbounded EITR scheme, then we can use \mathcal{A} to build a PPT algorithm \mathcal{B} that breaks IND-CPA security of the underlying BD-TR scheme. \square

The reduction is defined as follows:

1. \mathcal{B} first gets $1^\kappa, L$ from the adversary \mathcal{A} (in the public trace setting, \mathcal{A} outputs L adaptively along with the challenge messages). It then sets $n_{\text{bd}} = 2$ and sends $(1^\kappa, 1^{n_{\text{bd}}}, L)$ to the BD-TR challenger. The BD-TR challenger returns BD.mpk .
2. \mathcal{B} sets $\text{BD.mpk}_1 = \text{BD.mpk}$, generates $((\text{BD.mpk}_j, \text{BD.msk}_j) \leftarrow \text{BD.Setup}(1^\lambda, 1^\kappa, 2^j))$ for $j \in \{2, \dots, \lambda\}$ and sends $\text{mpk} = \{\text{BD.mpk}_j\}_{j \in [\lambda]}$ to \mathcal{A} .
3. When \mathcal{A} makes a key query (lb, id) , \mathcal{B} sends (lb, id) to the BD-TR challenger as a key query and sets the secret key it receives as BD.sk_1 . \mathcal{B} then computes $\text{BD.sk}_j \leftarrow \text{BD.KeyGen}(\text{BD.msk}_j, \text{lb}, \text{id})$ for $j \in \{2, \dots, \lambda\}$ and sends $\text{sk}_{\text{lb}, \text{id}} = \{\text{BD.sk}_j\}_{j \in [\lambda]}$ to \mathcal{A} .
4. When \mathcal{A} sends the challenge query (m_0, m_1) , \mathcal{B} samples $r_j \leftarrow \mathcal{M}$ for $j \in \{2, \dots, \lambda\}$, sets $m'_b = \bigoplus_{j>1} r_j \oplus m_b$ for $b \in \{0, 1\}$ and sends (m'_0, m'_1) as challenge query to the BD challenger and sets the returned ciphertext as BD.ct_1 . \mathcal{B} then computes $\text{BD.ct}_j \leftarrow \text{BD.Enc}(\text{BD.mpk}_j, r_j, L)$ for $j \in \{2, \dots, \lambda\}$ and sends $\text{ct} = \{\text{BD.ct}_j\}_{j \in [\lambda]}$ to \mathcal{A} .
5. \mathcal{A} outputs a bit b' , \mathcal{B} sends b' to the BD-TR challenger.

We observe that \mathcal{B} issues a key query (lb, id) to the BD-TR challenger only when \mathcal{A} makes a key query on (lb, id) . So, by the admissibility of \mathcal{A} , we have $\text{lb} \in L$ and thus \mathcal{B} is admissible in the IND-CPA game of BD-TR. Also, if \mathcal{A} has an advantage ϵ in distinguishing the encryptions of m_0, m_1 , then clearly \mathcal{B} has the same advantage in distinguishing the encryptions of m'_0, m'_1 and thus breaking the IND-CPA security of BD-TR.

Secure Tracing Analysis. Now we show that our construction satisfies false trace and correct trace guarantees.

False Trace Guarantee.

Theorem 10.3. *Assume that the underlying (secret/public tracing)-BD-TR scheme satisfies (selective/adaptive) false trace guarantee, then our construction of unbounded (secret/public tracing)-EITR satisfies (selective/adaptive) false trace guarantee as defined in Def. 7.12/Def. 7.11.*

Proof. Suppose there exists a PPT adversary \mathcal{A} , polynomial $p(\lambda)$ and non-negligible functions $\epsilon(\cdot), \delta(\cdot)$ such that $\Pr[\text{Fal-Tr}]_{\mathcal{A}, \epsilon, p} \geq \delta(\lambda)$, then we can build a PPT reduction \mathcal{B} that can break the false trace guarantee of BD-TR. The reduction is as follows:

1. \mathcal{B} first receives $1^\kappa, L$ from the adversary \mathcal{A} . (In the public trace setting, \mathcal{A} outputs L adaptively along with the challenge messages). It sets $i = \lceil \log p(\lambda) \rceil$, sends $(1^\kappa, 1^{2^i}, L)$ to the BD-TR challenger and sets the public key it receives as BD.mpk_i . \mathcal{B} generates $(\text{BD.mpk}_j, \text{BD.msk}_j) \leftarrow \text{BD.Setup}(1^\lambda, 1^\kappa, 2^j)$ for $j \in [\lambda] \setminus \{i\}$ and sends $\text{mpk} = \{\text{BD.mpk}_j\}_{j \in [\lambda]}$ to \mathcal{A} .
2. When \mathcal{A} makes a key query (lb, id) , \mathcal{B} forwards it as a key query to the BD-TR challenger and sets the secret key it receives as BD.sk_i . \mathcal{B} generates $\text{BD.sk}_j \leftarrow \text{BD.KeyGen}(\text{BD.msk}_j, \text{lb}, \text{id})$ for $j \in [\lambda] \setminus \{i\}$ and sends $\text{sk}_{\text{lb}, \text{id}} = \{\text{BD.sk}_j\}_{j \in [\lambda]}$ to \mathcal{A} .
3. Finally, when \mathcal{A} outputs a decoder box D and messages m_0, m_1 , \mathcal{B} runs $\text{isGoodDecoder}((\text{BD.mpk}_j)_{j \in [\lambda]}, 1^{1/\epsilon}, m_0, m_1, r, L, i)$ as defined in Figure 12, for $\lambda \cdot y$ many choices of r , until it finds an r s.t isGoodDecoder outputs 1.
4. \mathcal{B} constructs a decoding box \tilde{D} as defined in Figure 13, sends $(\tilde{D}, m_0 \oplus r, m_1 \oplus r)$ to the BD-TR challenger.

We observe that \tilde{D} uses the decoder D as a subroutine and it returns the response of D as its output. So, if \mathcal{A} outputs (D, m_0, m_1, L) such that the false trace guarantee does not hold with non-negligible probability, then \mathcal{B} breaks the false trace guarantee of the underlying BD-TR scheme. \square

Correct Trace Guarantee.

Theorem 10.4. *Assume that the underlying (secret/public tracing)-BD-TR scheme satisfies (selective/adaptive) correct trace guarantee, then our construction of unbounded (secret/public tracing)-EITR satisfies (selective/adaptive) correct trace guarantee as defined in Def. 7.12/Def. 7.11.*

Proof. Let $i = \lceil \log p(\lambda) \rceil$ and let S, S_{TD}, T_{lb} be as defined in Def. 7.11. Consider the following events

Event Cor-Tr_i : is defined as the event that the SubTrace algorithm, when run for position i outputs a correct traitor set T , i.e. $|T| > 0, (T \subseteq S_{ID}) \wedge (T_{\text{lb}} \cap L = \phi)$.

Event $\text{Good-}\tilde{D}_i$: is defined as the event that the flag is set to 1 in step 3 and the decoder \tilde{D} defined in Fig 14 is $\epsilon/4$ good decoder.

Event $\text{Found-Good-}r_i$: is defined as the event that the isGoodDecoder algorithm, when run for position i outputs 1.

With these definitions, the correctness is argued via following series of inequalities:

$$\Pr\text{-Corr-Tr}(\lambda) = \Pr[\text{Corr-Tr}_i] \quad (10.1)$$

$$\geq \Pr[\text{Corr-Tr}_i \wedge \text{Found-Good-}r_i] \quad (10.2)$$

$$\geq \Pr[\text{Good-}\tilde{D}_i \wedge \text{Found-Good-}r_i \wedge p(\lambda) \geq |S_{ID}|] - \text{negl}(\lambda) \quad (10.3)$$

$$\geq \Pr[\text{Good-}\tilde{D}_i \wedge \text{Found-Good-}r_i \wedge \text{Good-Decoder} \wedge p(\lambda) \geq |S_{ID}|] - \text{negl}(\lambda) \quad (10.4)$$

$$\geq \Pr[\text{Found-Good-}r_i \wedge \text{Good-Decoder} \wedge p(\lambda) \geq |S_{ID}|] - \text{negl}(\lambda) \quad (10.5)$$

$$\geq \Pr[\text{Good-Decoder} \wedge p(\lambda) \geq |S_{ID}|] - \text{negl}(\lambda) \quad (10.6)$$

Equation (10.1) and (10.2) follow directly from the definition of the events involved. Equation (10.3) follows from the following claim:

Claim 10.4.1. *If BD-TR guarantees correct tracing then*

$$\Pr \left[\text{Good-}\tilde{D}_i \wedge \text{Found-Good-}r_i \wedge p(\lambda) \geq |S_{ID}| \right] - \Pr[\text{Corr-Tr}_i \wedge \text{Found-Good-}r_i] \leq \text{negl}(\lambda).$$

Proof. We show that if there exists an adversary \mathcal{A} who outputs a good decoder along with messages m_0, m_1 and a revocation list L such that $\Pr[\text{Good-}\tilde{D}_i \wedge \text{Found-Good-}r_i \wedge p(\lambda) \geq |S_{ID}|] - \Pr[\text{Cor-Tr}_i]$ is non negligible then we can use \mathcal{A} to construct an adversary \mathcal{B} against correct trace guarantee of the underlying BD-TR. The reduction is defined as follows:

1. \mathcal{B} first gets $(1^\kappa, L)$ from \mathcal{A} (in the public trace setting, \mathcal{A} outputs L adaptively along with the challenge messages). It sends $(1^\kappa, 1^{2^i}, L)$ to the BD-TR challenger and sets the public key it receives as BD.mpk_i .
2. For $j \in [\lambda] \setminus \{i\}$, \mathcal{B} generates $(\text{BD.mpk}_j, \text{BD.msk}_j) \leftarrow \text{BD.Setup}(1^\kappa, 2^j)$; sets $\text{mpk} = \{\text{BD.mpk}_j\}_{j \in [\lambda]}$, and sends mpk to \mathcal{A} .
3. For each key query (lb, id) from \mathcal{A} , \mathcal{B} sends (lb, id) as a key query to the BD-TR challenger and sets the key received as BD.sk_i . For $j \in [\lambda] \setminus \{i\}$, it computes BD.sk_j itself using BD.msk_j . It sends $\text{sk}_{\text{lb}, \text{id}} = \{\text{BD.sk}_j\}_{j \in [\lambda]}$ to \mathcal{A} .
4. In the end, \mathcal{A} outputs a decoder D along with messages m_0, m_1 .
5. \mathcal{B} does the following:
 - If the number of key queries by \mathcal{A} exceeds $p(\lambda)$, then \mathcal{B} outputs a random decoder and quits.
 - Else, \mathcal{B} runs $\text{isGoodDecoder}(\{\text{BD.mpk}_j\}_{j \in [\lambda]}, 1^y, m_0, m_1, r, L, i)$ for uniformly and independently sampled r until it finds a r for which isGoodDecoder algorithm outputs 1. If isGoodDecoder algorithm does not output 1 even after $\lambda \cdot y$ attempts, then \mathcal{B} outputs a random decoder and quits.

- Else, \mathcal{B} constructs decoder \tilde{D} as defined in Figure 14 and sets $\tilde{m}_b = m_b \oplus r$ for $b \in \{0, 1\}$.

6. \mathcal{B} sends $(\tilde{D}, \tilde{m}_0, \tilde{m}_1)$ to the BD-TR challenger.

Now we analyze the probability that \mathcal{B} outputs a good decoder. Observe that \mathcal{B} does not abort and outputs a genuine decoder when both Found-Good- r_i and $|S_{ID}| \leq p(\lambda)$ happens. Since the decoder returned by \mathcal{B} is the same decoder as defined in the SubTrace algorithm, we get that the probability that \mathcal{B} outputs a good decoder is $\Pr[\text{Good-}\tilde{D}_i \wedge \text{Found-Good-}r_i \wedge p(\lambda) \geq |S_{ID}|]$. Furthermore, the probability that the BD.Trace algorithm outputs correct set of traitors T (using the decoder returned by \mathcal{B}) is same as $\Pr[\text{Corr-Tr}_i \wedge \text{Found-Good-}r_i]$. Hence, if $\Pr[\text{Good-}\tilde{D}_i \wedge \text{Found-Good-}r_i \wedge p(\lambda) \geq |S_{ID}|] - \Pr[\text{Corr-Tr}_i \wedge \text{Found-Good-}r_i]$ is non negligible, then it breaks the security of correct trace guarantee of the underlying BD-TR scheme. \square

Equation (10.4) again follows directly from the definition. Argument for transition from equation (10.4) to (10.5) and (10.5) to (10.6) is same as that of transition from (9.8) to (9.9) and (9.9) to (9.10), respectively. This completes the proof. \square

11 Extension to Super-Polynomial Size Revocation List

While our main focus in the paper is on the case where the revocation list L is of polynomial size, we can consider an extension where it is of super-polynomial size. In particular, we consider the setting where L has efficient representation by a circuit C_L of polynomial size. Namely, we have $C_L(\text{lb}) = 0$ for revoked label lb and $C_L(\text{lb}) = 1$ for non-revoked label lb . We assume that the depth of C_L is bounded by some polynomial \bar{d} . Namely, $C_L \in \mathcal{C}_{|\text{lb}|, \bar{d}}$. This assumption is necessary because we will use kpABE (resp., cpABE) with the same restriction to generate a secret key (resp., a ciphertext) associated with C_L .

EITR scheme for this setting can be obtained both in the secret and the public tracing settings very similarly to the case where the revocation list is of polynomial size. Namely, we first construct (secret key/public key) RPE with super-polynomial revocation list and then apply the chain of conversions (Sec. 8, 9, and 10). The conversions work almost without change with the natural adaptation of replacing L with C_L and the membership check $\text{lb} \stackrel{?}{\in} L$ with $C(\text{lb}) \stackrel{?}{\in} 0$. The constructions of (public key/secret key) RPE are also almost the same as those in (Sec. 4/Sec. 5 and 6) with the natural adaptation of generating kpABE secret key for C_L ²⁷ in Sec. 4 and 5 and replacing the condition $\text{lb} \notin L$ in Eq. (6.1) defining $C_{L, \text{RMFE.ct}}$ with $C_L(\text{lb}) = 1$. The main difference is that we have to assume sub-exponential LWE assumption instead of (polynomial) LWE assumption for both secret key and public key settings here, because we need adaptive security for the underlying kpABE. We give further details in the following.

- In the public key setting, indistinguishability of Hybrid_{6,a} and Hybrid₇ shown in Claim 4.2.7, where post-challenge key queries are dealt with, cannot be proven any more if we only assume selective security for kpABE. The reason why the original proof does not go through is that we have to deal with the kpABE queries in the order of key first and ciphertext later. With polynomial size L , this does not pose a problem because when the adversary chooses L , all the labels for which we use kpABE security are in L and we can perform a hybrid argument over these labels. However, this is not possible for super-polynomial size L . To deal with the queries in this order, we assume adaptive

²⁷Originally, we start from the revocation list L and then construct the circuit C_L that hardwires L into it. Here, we directly use the circuit C_L that efficiently represents the super-polynomial set L .

security (as per Definition 2.16) for kpABE. Then, the indistinguishability of the games can easily be shown by changing the post-challenge kpABE ciphertexts associated with lb with $C_L(lb) = 0$ one by one.

- In the secret key setting, we also encounter similar difficulty. In particular, indistinguishability of Hybrid₁ and Hybrid₂ shown in Claim 5.8.2 cannot be proven any more if we only assume selective security for kpABE by exactly the same reason. We can overcome the problem by assuming adaptive security for kpABE.

Finally, we briefly discuss the parameter size of the resulting ETR scheme. The parameter size of the resulting scheme is the same as the case of polynomial size revocation list except that they rely on \bar{d} . Notably, they are independent from the size of the circuits being supported inheriting the succinctness properties of the underlying kpABE and cpABE.

Acknowledgements. This work was partly supported by the DST “Swarnajayanti” fellowship, Cybersecurity Center of Excellence, IIT Madras, National Blockchain Project, European Union Horizon 2020 Research and Innovation Program Grant 780701, BPIFrance in the context of the national project RISQ (P141580), and the ANR AMIRAL project (ANR-21-ASTR-0016). The fourth author is partially supported by JST AIP Acceleration Research JPMJCR22U5 and JSPS KAKENHI Grant Number 19H01109, Japan.

References

- [ABSV15] Prabhanjan Ananth, Zvika Brakerski, Gil Segev, and Vinod Vaikuntanathan. From selective to adaptive security in functional encryption. In *CRYPTO*, 2015.
- [Agr19] Shweta Agrawal. Indistinguishability obfuscation without multilinear maps: New techniques for bootstrapping and instantiation. In *Eurocrypt*, 2019.
- [AJ15] Prabhanjan Ananth and Abhishek Jain. Indistinguishability obfuscation from compact functional encryption. In *CRYPTO*, 2015.
- [AJS15] Prabhanjan Ananth, Abhishek Jain, and Amit Sahai. Achieving compactness generically: Indistinguishability obfuscation from non-compact functional encryption. IACR Cryptology ePrint Archive, 2015:730, 2015.
- [ALP11] Nuttapon Attrapadung, Benoît Libert, and Elie de Panafieu. Expressive key-policy attribute-based encryption with constant-size ciphertexts. In *PKC*. Springer, 2011.
- [AM18] Shweta Agrawal and Monosij Maitra. Fe and io for turing machines from minimal assumptions. In *TCC*. Springer, 2018.
- [AV19] Prabhanjan Ananth and Vinod Vaikuntanathan. Optimal bounded-collusion secure functional encryption. In *TCC*, 2019.
- [AWY20] Shweta Agrawal, Daniel Wichs, and Shota Yamada. Optimal broadcast encryption from lwe and pairings in the standard model. In *TCC*, 2020.
- [AY20] Shweta Agrawal and Shota Yamada. Optimal broadcast encryption from pairings and lwe. In *EUROCRYPT*, 2020.

- [BGG⁺14] Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In *EUROCRYPT*, 2014.
- [BGI⁺01] B. Barak, O. Goldreich, R. Impagliazzo, S. Rudich, A. Sahai, S. Vadhan, and K. Yang. On the (im)possibility of obfuscating programs. In *CRYPTO*, 2001.
- [BGW05] Dan Boneh, Craig Gentry, and Brent Waters. Collusion resistant broadcast encryption with short ciphertexts and private keys. In *CRYPTO*, 2005.
- [BKS16] Zvika Brakerski, Ilan Komargodski, and Gil Segev. Multi-input functional encryption in the private-key setting: Stronger security from weaker assumptions. In *EUROCRYPT*. Springer, 2016.
- [BS18] Zvika Brakerski and Gil Segev. Function-private functional encryption in the private-key setting. *Journal of Cryptology*, 31(1):202–225, 2018.
- [BSW07] John Bethencourt, Amit Sahai, and Brent Waters. Ciphertext-policy attribute-based encryption. In *IEEE Symposium on Security and Privacy*, pages 321–334, 2007.
- [BV15] Nir Bitansky and Vinod Vaikuntanathan. Indistinguishability obfuscation from functional encryption. *FOCS*, 2015.
- [BW06] Xavier Boyen and Brent Waters. Anonymous hierarchical identity-based encryption (without random oracles). In *CRYPTO*, 2006.
- [BWZ14] Dan Boneh, Brent Waters, and Mark Zhandry. Low overhead broadcast encryption from multilinear maps. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology – CRYPTO 2014*, 2014.
- [BZ17] Dan Boneh and Mark Zhandry. Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. *Algorithmica*, 79(4):1233–1285, 2017.
- [CFN94] Benny Chor, Amos Fiat, and Moni Naor. Tracing traitors. In *CRYPTO*, 1994.
- [CIJ⁺13] Angelo De Caro, Vincenzo Iovino, Abhishek Jain, Adam O’Neill, Omer Paneth, and Giuseppe Persiano. On the achievability of simulation-based security for functional encryption. In *CRYPTO*, 2013.
- [CLT13] Jean-Sébastien Coron, Tancrede Lepoint, and Mehdi Tibouchi. Practical multilinear maps over the integers. In *CRYPTO*, 2013.
- [CVW⁺18] Yilei Chen, Vinod Vaikuntanathan, Brent Waters, Hoeteck Wee, and Daniel Wichs. Traitor-tracing from lwe made simple and attribute-based. In *TCC*, 2018.
- [DQV⁺21] Lalita Devadas, Willy Quach, Vinod Vaikuntanathan, Hoeteck Wee, and Daniel Wichs. Succinct lwe sampling, random polynomials, and obfuscation. In *TCC*, 2021.
- [FN93] Amos Fiat and Moni Naor. Broadcast encryption. In Douglas R. Stinson, editor, *CRYPTO*, 1993.

- [GGH13] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In *EUROCRYPT*, 2013.
- [GKW17] Rishab Goyal, Venkata Koppula, and Brent Waters. Lockable obfuscation. In *FOCS*, 2017.
- [GKW18] Rishab Goyal, Venkata Koppula, and Brent Waters. Collusion resistant traitor tracing from learning with errors. In *STOC*, 2018.
- [GKW19] Rishab Goyal, Venkata Koppula, and Brent Waters. New approaches to traitor tracing with embedded identities. In *TCC*, 2019.
- [GP21] Romain Gay and Rafael Pass. Indistinguishability obfuscation from circular security. In *STOC*, 2021.
- [GPS16] Sanjam Garg, Omkant Pandey, and Akshayaram Srinivasan. Revisiting the cryptographic hardness of finding a nash equilibrium. In *CRYPTO*, 2016.
- [GPSZ17] Sanjam Garg, Omkant Pandey, Akshayaram Srinivasan, and Mark Zhandry. Breaking the sub-exponential barrier in obfustopia. In *EUROCRYPT*, 2017.
- [GQWW19] Rishab Goyal, Willy Quach, Brent Waters, and Daniel Wichs. Broadcast and trace with n^ϵ ciphertext size from standard assumptions. In *Crypto*, 2019. <https://eprint.iacr.org/2019/636>.
- [GS16] Sanjam Garg and Akshayaram Srinivasan. Single-key to multi-key functional encryption with polynomial loss. In *TCC*, 2016.
- [GVW13] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute based encryption for circuits. In *STOC*, 2013.
- [GVW19] Rishab Goyal, Satyanarayana Vusirikala, and Brent Waters. Collusion resistant broadcast and trace from positional witness encryption. In *PKC*, 2019.
- [JLS21] Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from well-founded assumptions. In *STOC*, 2021.
- [JLS22] Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from lpn over large fields, dlin, and constant depth prgs. In *EUROCRYPT*, 2022.
- [KS20] Ilan Komargodski and Gil Segev. From minicrypt to obfustopia via private-key functional encryption. *Journal of Cryptology*, 33(2):406–458, 2020.
- [KW20] Sam Kim and David J Wu. Collusion resistant trace-and-revoke for arbitrary identities from standard assumptions. In *ASIACRYPT*, 2020.
- [LPST16] Huijia Lin, Rafael Pass, Karn Seth, and Sidharth Telang. Output-compressing randomized encodings and applications. In *TCC*, 2016.
- [NP10] Moni Naor and Benny Pinkas. Efficient trace and revoke schemes. *International Journal of Information Security*, 9(6):411–424, 2010.
- [NWZ16] Ryo Nishimaki, Daniel Wichs, and Mark Zhandry. Anonymous traitor tracing: how to embed arbitrary information in a key. In *EUROCRYPT*, 2016.

- [Tak14] Katsuyuki Takashima. Expressive attribute-based encryption with constant-size ciphertexts from the decisional linear assumption. In *SCN*, 2014.
- [Tsa19] Rotem Tsabary. Fully secure attribute-based encryption for t-CNF from LWE. In *CRYPTO*, 2019.
- [Tsa22] Rotem Tsabary. Candidate witness encryption from lattice techniques. In *Crypto*, 2022.
- [Wee22] Hoeteck Wee. Optimal broadcast encryption and cp-abe from evasive lattice assumptions. In *Eurocrypt*, 2022.
- [WW21] Hoeteck Wee and Daniel Wichs. Candidate obfuscation via oblivious lwe sampling. In *EUROCRYPT*. Springer, 2021.
- [WZ17] Daniel Wichs and Giorgos Zirdelis. Obfuscating compute-and-compare programs under LWE. In *FOCS*, 2017.
- [Zha20] Mark Zhandry. New techniques for traitor tracing: Size $N^{1/3}$ and more from pairings. In *CRYPTO*, 2020.