Policy-Based Non-interactive Outsourcing of Computation using multikey FHE and CP-ABE*

Michael Clear† and Ciarán McGoldrick
School of Computer Science and Statistics,
Trinity College Dublin
{*clearm, Ciaran.McGoldrick*}*@scss.tcd.ie*

Keywords:
Non-interactive computing delegation, multikey FHE, CP-ABE, homomorphic encryption, access policy composition

Abstract:
We consider the problem of outsourced computation that operates on encrypted inputs supplied by multiple independent parties. To facilitate fine-grained access control, it would be desirable if each party could encrypt her input under an appropriate access policy. Moreover, a party should only be authorized to decrypt the result of a computation performed on a set of encrypted inputs if his credentials satisfy the composition of all input policies. There has been limited success so far achieving homomorphic encryption in the functional setting; that is, for primitives such as Ciphertext-Policy Attribute Based Encryption (CP-ABE) and Identity Based Encryption (IBE). We introduce a new primitive that captures homomorphic encryption with support for access policies and policy composition. We then present a generic construction using CP-ABE and multikey Fully-Homomorphic encryption (FHE). Furthermore, we show that a CP-ABE scheme that is homomorphic for circuits of polylogarithmic depth in some parameter m implies a CP-ABE scheme that is homomorphic for circuits of arity m and unbounded depth.

1 Introduction

With the advent of cloud computing, there is a rapidly expanding interest in using remote data centers to perform large computational tasks. Many organizations do not have the computational resources to perform such tasks and the low cost,scalable and highly available model offered by remote providers present an attractive option to organizations. A significant downside of delegating computing jobs to the cloud is the risk of exposure of the delegator's sensitive data. Indeed, sending such data in an unencrypted form may be strictly prohibited by government and industry policies. A number of cryptographic primitives have been proposed to preserve privacy in computing tasks carried out by untrusted or semi-trusted parties. A well-known example is fully-homomorphic encryption (FHE), which was first realized in 2009 by Gentry (Gentry, 2009). FHE allows us to outsource a computation to a cloud provider in such a way that the cloud provider can carry out the computation without being able to see the inputs and outputs. Gentry's construction is public-key and thus allows public delegatability insofar as the sender(s) of inputs to the cloud need not have access to the secret key needed to decrypt the result. Therefore, multiple encryptors may independently contribute data that is to be (potentially) incorporated into a large remote computation.

1.1 The Problem Domain

In standard public-key FHE, there is only a single target recipient. This may be ill-suited to the needs of a large organization. Consider a scenario where staff have restricted access to data based on their department and position. The organization has opted to avail of the computational resources of a cloud provider for the purpose of delegating sizeable computational tasks. Each sender of data acts independently since they are potentially unaware of other's participation.

To comply with the organization's privacy regulations, each sender must encrypt her data under an appropriate access policy that specifies the

*This is the full version of a paper that appeared at SECRYPT 2013 (Clear and McGoldrick, 2013).

†The author's work is funded by the Irish Research Council EMBARK Initiative.

credentials a staff member must have in order to access the data (or any derivative thereof). We assume such an access policy is feasibly determined from the data source and context.

The computation to be performed, and the inputs to be used, may be decided at a later stage by a subset of the senders, or other delegated authority. The results of the computation are then subsequently returned to the organization, and they should *only* be accessible to a given staff member if her credentials satisfy the cumulative policies associated with *all* the inputs used.

One solution is to use public-key FHE together with a trusted access control system (ACS), which holds the private key for the FHE scheme. The role of the ACS is to grant users (i.e. staff members in the above scenario) access to a plaintext after verifying that their credentials satisfy the policy set associated with the corresponding ciphertext. Access control of this form facilitates expressive policies. However, it must be used in conjunction with a cryptographic primitive such as a non-interactive zero-knowledge proof system as, otherwise, unauthorized users may collude in order to report an incorrect policy.

This approach suffers from a number of drawbacks:

- All parties interested in a result are required to contact the ACS, which must remain online and exhibit high availability in order to guarantee satisfactory responsiveness. The ACS may therefore act as a bottleneck, especially under high load scenarios.
- Adhering to the principle of least privilege, the organization may wish to limit the capabilities of the ACS. In particular, it may have reservations about the ACS being compromised, and potentially providing an attacker access to all results returned from the cloud.
- Remote users, with appropriate valid credentials, cannot directly query the cloud for data and decrypt non-interactively. All requests must be routed via the organization's ACS.

Many of these shortcomings are flexibly addressed through a functional encryption (FE) approach. In the FE setting, a trusted authority (TA) authenticates and authorizes users by issuing them secret keys for certain *capabilities*. For our purposes, we deal with a special case of FE known as ciphertext-policy attribute-based encryption (CP-ABE) where the *capabilities* correspond to credentials or *attributes*. Note that we use the term *attribute* here to refer to (collectively) what some authors describe as a particular set of attributes. A user with a secret key for an attribute a can decrypt any ciphertext encrypted under a policy satisfied by a. A principal advantage of CP-ABE over an ACS-based solution is that once the user is issued a secret key for a, no further interaction with the TA is required (for a certain period of time i.e. a may be time-limited) throughout which the user can decrypt an arbitrary number of ciphertexts non-interactively. The advantages of ABE in distributed environments have been investigated in other work, such as (Pirretti et al., 2010). Although CP-ABE has some deficiencies, such as inherent escrow (which the ACS approach suffers from also) and a lack of support for revocation, it is well-suited to achieving fine-grained access control with minimal interaction.

It is not trivial to reconcile the features of FHE and CP-ABE. After the publication of the preliminary version of this work, the first fully-homomorphic* attribute-based and identity-based encryption schemes were presented (Gentry et al., 2013). This solves a challenging open problem.

1.2 Contributions

In this work, we propose a syntax for a more general primitive which seeks to capture the requirements of the problem space described above, while incorporating properties from FHE and CP-ABE. We call this primitive *policy-based homomorphic encryption* (PBHE). The formulation of PBHE extends the recent definition of multikey FHE by Lopez-Alt, Tromer and Vaikuntanathan (López-Alt et al., 2012). Central to PBHE is the notion of access policy composition, and we define the syntax and the correctness properties of PBHE in terms of an algebraic structure defined on access policies. PBHE can be instantiated by any homomorphic CP-ABE scheme or any standard homomorphic public-key cryptosystem.

Another contribution of this work is the construction of a new PBHE scheme that supports fully-homomorphic evaluation of circuits whose input ciphertexts are encrypted under a bounded number of independently-chosen policies. This scheme fulfills the requirements of the scenario outlined above for a bounded number of senders.

Finally, and leveraging the work of (López-Alt et al., 2012), we prove that if a CP-ABE scheme $\mathcal{E}$ is homomorphic for a class of circuits of polylogarithmic depth in a parameter m (which is polynomial in the security parameter), then there exists a scheme $\mathcal{E}'$ that is homomorphic for all circuits with arity m and with arbitrary depth. This is a significant result as obtaining homomorphic CP-

*More precisely, such schemes are levelled fully-homomorphic; they can evaluate circuits of bounded depth, albeit polynomially sized in the security parameter.

ABE for circuits of unbounded depth has been impeded by the fact that there does not seem to be a way to employ bootstrapping in the functional setting (non-interactively) since bootstrapping requires encryptions of the secret key bits to be available as part of the public key.

We note that our work in this paper is limited to the semi-honest model. In particular, we assume that the cloud is semi-honest. We leave to future work the challenge of securing against malicious adversaries, especially in verifying that a function was evaluated correctly.

1.3 Related Work

Homomorphic encryption in a multi-user setting is considered in (Xiao et al., 2012). The authors propose a new homomorphic symmetric encryption scheme that is shown to be secure under a factoring assumption, provided an attacker is limited to obtaining a bounded number (polynomial in the security parameter) of plaintext-ciphertext pairs. The authors also propose a system model with mutually untrusted components that enables a master key for their scheme to be derived from any user's key by splitting it into shares that are distributed to each component by a dealer at system initialization time. However, their solution requires interaction with a server known as a *key agent* for every request/response to/from the cloud. Furthermore, their solution does not support any level of expressive access control.

The notion of multikey FHE was recently presented in (López-Alt et al., 2012) along with a concrete construction based on NTRU (Hoffstein et al., 1998). In that work, multikey FHE is used to construct an "on-the-fly" multi-party computation (MPC) protocol that is secure in the malicious setting. In such an MPC protocol, a number of parties independently send encrypted inputs to an evaluator without interaction. The evaluator then computes a function F over the inputs and sends the encrypted result to each party. It is then possible for the senders to run an interactive MPC protocol to jointly decrypt the result, verify each other's participation, and verify F was honestly computed by the evaluator. While we make use of multikey FHE for our generic construction in Section 5 and as a basis for the syntax of PBHE in Section 4, we address a different problem than (López-Alt et al., 2012) i.e. we do not target MPC wherein each party wishes to keep his input secret. In our protocol, interactive decryption is avoided at the loss of verifiability. Achieving the latter in a meaningful way is a topic for future work.

Additional related work arises in the ABE setting, such as the construction of CP-ABE (Bethencourt et al., 2007), and in the area of access control facilitating access policy composition (Bonatti et al., 2002; Bruns et al., 2007; Ni et al., 2009). More recently (Rao et al., 2011)'s work on policy composition has targeted real-world access control languages like XACML (Moses et al., 2005). In our context the objects to protect are *data*, and the policies are not enforced by a server but rather by an encryption scheme, so it is important to note that the scope for policy composition is far more restrictive as it is necessary to preserve semantic security.

An independent and concurrent work (Gentry et al., 2013) that appeared after the preliminary version of this paper solves the open problem of identity-based and attribute-based FHE (more precisely leveled FHE). The main advantage of our approach is that it supports composition of distinct attributes/identities/policies whereas the schemes in (Gentry et al., 2013) permits evaluation only on ciphertexts encrypted under the *same* attribute/identity. However, the schemes in (Gentry et al., 2013) have constant-sized ciphertexts and are based on weaker assumptions (in comparison to the assumptions on which the concrete instantiations of our schemes are based).

2 Preliminaries

2.1 Notation

A quantity t is said to be negligible in a parameter κ if it holds that $t < 1/f(\kappa)$ for all polynomials f. We write this as $\mathsf{negl}(\kappa)$.

If D is a random variable, the notation $x \xleftarrow{\$} D$ denotes the fact that x is sampled according to the distribution on D. If instead that D is a set, then the notation is understood to mean that x is uniformly sampled from D.

We use the notation $[t]$ to denote the set of contiguous integers $\{1, \ldots, t\}$.

2.2 Access Policies

An access policy is a predicate that grants or denies permission to access a particular object in some specific manner. Some contexts require rich policies that present multiple outcomes for an access. For example, Bruns, Dantas and Huth (Bruns et al., 2007) represent a policy as a four-valued predicate whose range is {*grant*, *deny*, *unspecified* or *conflict*}. Access control systems with these requirements typically accommodate many modes of access to an object. In our case, the objects correspond to data, and it is meaningful in this context to either grant or deny (mutually

exclusive) access to a datum. Therefore, we naturally represent an access policy as a two-valued predicate.

2.3 CP-ABE Syntax

A CP-ABE scheme for a class of access policies $\mathbb{F}$ defined over a domain of attributes $\mathbb{A}$ with message space $\mathbb{M}$ is a tuple of PPT algorithms $(\mathsf{Setup}, \mathsf{Extract}, \mathsf{Enc}, \mathsf{Dec})$. As mentioned in the introduction, we refer to the entity that an access policy is applied to as an *attribute* instead of a set of attributes as in (Bethencourt et al., 2007). An attribute a in a domain $\mathbb{A}$ may be viewed as a set of "sub-attributes". Accordingly, we express access policies as predicates i.e. $\mathbb{F} \subseteq \mathbb{A} \to \{0,1\}$. The Trusted Authority (TA) runs Setup to generate the public parameters PP and a master secret key MSK. It runs $\mathsf{sk}_a \leftarrow \mathsf{Extract}(\mathsf{MSK}, a)$ to derive a secret key for an attribute $a \in \mathbb{A}$.

There are two main definitions of semantic security which are distinguished by whether the adversary is allowed to make adaptive requests for secret keys. In the non-adaptive (IND-NA-CPA) game, a challenger hands PP to the adversary $\mathcal{A}$ who can make queries to an extraction oracle $\mathcal{X} := \mathsf{Extract}(\mathsf{MSK}, \cdot)$ to obtain secret keys for certain attributes. At the end of this phase, $\mathcal{A}$ chooses a target policy $f^* \in \mathbb{F}$ and two messages $m_0, m_1 \in \mathbb{M}$. The challenger uniformly samples a bit $b \xleftarrow{\$} \{0,1\}$ and gives an encryption of m_b under f^* to $\mathcal{A}$. Finally, $\mathcal{A}$ outputs a guess bit b' and is said to win if $b = b'$.

The advantage of $\mathcal{A}$ is defined as $\mathsf{Pr}[b' = b] - 1/2$. A CP-ABE scheme is said to be IND-NA-CPA secure if every PPT adversary $\mathcal{A}$ has only a negligible advantage in the above game. In the adaptive variant (IND-AD-CPA), $\mathcal{A}$ is also allowed to make queries to $\mathcal{X}$ after the challenge ciphertext is generated provided that the further attributes it queries do not satisfy f'.

2.4 Multikey FHE

Multikey FHE allow multiple independently-generated keys to be used together in a homomorphic evaluation. The syntax of multikey FHE imposes a limit N on the number of such keys that can be supported. Furthermore, the size of the evaluated ciphertext does not depend on the size of the circuit (or number of inputs), but instead on the number of independent keys N that is supported. In order to decrypt, the parties who have the corresponding secret keys must collaborate in an MPC protocol.

Definition 2.1 (Based on Definition 2.1 in (López-Alt et al., 2012))**.** *A multikey $\mathbb{C}$-homomorphic scheme family for a class of circuits $\mathbb{C}$ and message space $\mathbb{M}$ is a family of PPT algorithms $\{\mathcal{E}^{(N)} := (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})\}_{N>0}$ where $\mathcal{E}^{(N)}$ is defined as follows:*

- $\mathsf{MKFHE.Gen}$ *takes as input the security parameter 1^κ and outputs a tuple $(\mathsf{pk}, \mathsf{sk}, \mathsf{vk})$ where pk is a public key, sk is a secret key and vk is an evaluation key.*
- $\mathsf{MKFHE.Enc}$ *takes as input a public key pk and a message $m \in \mathbb{M}$, and outputs an encryption of m under pk.*
- $\mathsf{MKFHE.Dec}$ *takes as input $1 \le k \le N$ secret keys $\mathsf{sk}_1, \ldots, \mathsf{sk}_k$ and a ciphertext c, and outputs a message $m' \in \mathbb{M}$.*
- $\mathsf{MKFHE.Eval}$ *takes as input a circuit $C \in \mathbb{C}$, and ℓ pairs $(c_1, \mathsf{vk}_1), \ldots, (c_\ell, \mathsf{vk}_\ell)$ and outputs a ciphertext $c^\star$.*

Informally, evaluation is only required to be *correct* if at most N keys are used in $\mathsf{MKFHE.Eval}$; that is, $|\{\mathsf{vk}_1, \ldots, \mathsf{vk}_\ell\}| \le N$. Furthermore, the size of an evaluated ciphertext $c^\star$ must only depend polynomially on the security parameter κ and the number of keys N, and not on the size of the circuit.

The IND-CPA security game for multikey homomorphic encryption is the same as that for standard public-key encryption; note that the adversary is given the evaluation key vk.

3 Homomorphic CP-ABE with Bounded Composition

Let $\mathbb{F}$ be a set of valid access policies which accept or reject members of a set of attributes $\mathbb{A}$. An access policy $f \in \mathbb{F}$ is represented as a predicate $\mathbb{A} \to \{0,1\}$. Recall that our goal is to facilitate joint computation on encrypted inputs contributed by multiple independent parties, who may be unaware of each other. Moreover, each party has the liberty to encrypt her inputs under an independently-chosen policy. Accordingly, it is necessary to support composition of these policies. Intuitively, one would expect that the result of the joint computation be decryptable by users with an attribute that satisfies the composite policy. Therefore, we introduce a binary composition operation $\odot$ defined on $\mathbb{F}$.

We begin by giving a precise definition of homomorphic CP-ABE. A CP-ABE scheme is homomorphic for a class of circuits $\mathbb{C}$ if there is an additional algorithm Eval and a composition operation $\odot : \mathbb{F}^2 \to \mathbb{F}$ such that over all choices of $f_1, \ldots, f_\ell \in \mathbb{F}$, $m_1, \ldots, m_\ell \in \mathbb{M}$, $c_1 \leftarrow \mathsf{Enc}(\mathsf{PP}, f_1, m_1), \ldots, c_\ell \leftarrow \mathsf{Enc}(\mathsf{PP}, f_\ell, m_\ell)$ and $C \in \mathbb{C}$, the ciphertext $c' \leftarrow \mathsf{Eval}(\mathsf{PP}, C, c_1, \ldots, c_\ell)$ satisfies

– **Correctness**

$$\mathsf{Dec}(\mathsf{sk}_a, c') = C(m_1, \ldots, m_\ell) \text{ iff } f'(a) = 1 \tag{3.1}$$

for any $a \in \mathbb{A}$ and $\mathsf{sk}_a \leftarrow \mathsf{Extract}(\mathsf{MSK}, a)$.

– **Compactness**

$$|c'| = \mathsf{poly}(\kappa, |f'|) \tag{3.2}$$

where $f' = f_1 \odot \ldots \odot f_\ell$.

The main idea in this paper is to exploit multikey FHE and CP-ABE to construct a new CP-ABE scheme that is homomorphic for a class of circuits $\mathbb{C}$ of bounded arity. However, we can only achieve this for certain policy algebras $(\mathbb{F}, \odot)$. Let $\mathcal{E}_{\mathsf{ABE}}$ be a CP-ABE scheme and let $\mathcal{E}_{\mathsf{MKFHE}}$ be a multikey FHE scheme. Roughly speaking, to encrypt a message m under policy f in our scheme, (1) a key triple $(\mathsf{pk}, \mathsf{vk}, \mathsf{sk})$ is generated for $\mathcal{E}_{\mathsf{MKFHE}}$; (2) m is encrypted with $\mathcal{E}_{\mathsf{MKFHE}}$ under pk; (3) sk is encrypted with $\mathcal{E}_{\mathsf{ABE}}$ under policy f; (4) the two previous ciphertexts along with vk constitute the ciphertext that is produced. Therefore, $\mathcal{E}_{\mathsf{MKFHE}}$ is used for hiding the message and for homomorphic computation whereas $\mathcal{E}_{\mathsf{ABE}}$ enforces the access policies by appropriately hiding the secret keys for $\mathcal{E}_{\mathsf{MKFHE}}$. Technically, it is the number of compositions in our scheme that must be bounded and not the arity of the circuits. However, the former implies the latter due to the syntactic restrictions of homomorphic CP-ABE (See Section 3).

It might seem necessary that $\odot$ be both commutative and associative. However, we only require that these properties hold with respect to semantics. We say that two policies $f, g \in \mathbb{F}$ are *semantically equivalent*, written $f \sim g$, if for all attributes $a \in \mathbb{A}$, we have that $f(a) = g(a)$. Formally, it is required that $\sim$ be a congruence relation with respect to $\odot$ and that $(\mathbb{F}/\sim, \odot)$ be a commutative semigroup. In sum, the properties that $\odot$ must satisfy for any $f, g, h \in \mathbb{F}$ are as follows:

1.
$$f \odot g \sim g \odot f \tag{3.3}$$

2.
$$(f \odot g) \odot h \sim f \odot (g \odot h) \tag{3.4}$$

3.
$$(f \odot g)(a) \Rightarrow f(a) \wedge g(a) \tag{3.5}$$

for any $a \in \mathbb{A}$.

Note that the last property is necessary for semantic security.

We denote the *size* of a policy $f \in \mathbb{F}$ by its length, written $|f| \in \mathbb{N}$. For some algebras, the size of policies do not always grow with composition. Consider the following semilattices (commutative idempotent semigroups).

– the *Kronecker semilattice* where $\odot$ is defined as:

$$f \odot g = \begin{cases} f & \text{if } f = g \\ z & \text{otherwise} \end{cases}$$

and z is a distinguished policy in $\mathbb{F}$ with the property that $z(a) = 0 \quad \forall a \in \mathbb{A}$.

– the *meet semilattice* where $\odot$ is defined as $\wedge$.

Using our approach as described above, we cannot construct homomorphic CP-ABE for idempotent algebras $(\mathbb{F}, \odot)$ because $|f \odot f| = |f|$, which implies that the compactness condition given by 3.2 cannot be satisfied since the ciphertexts in our scheme grow with composition. However, we have obtained the following result. Suppose that $\mathcal{E}_{\mathsf{ABE}}$ is a somewhat-homomorphic CP-ABE scheme with a policy algebra $(\mathbb{F}, \odot)$. More precisely, suppose that $\mathcal{E}_{\mathsf{ABE}}$ is homomorphic for a class of circuits $\mathbb{C}$ of depth that is polylogarithmic in the security parameter. Then there exists a CP-ABE scheme for $(\mathbb{F}, \odot)$ that is homomorphic for a class of circuits of arbitrary depth whose arity is bounded by a fixed polynomial in the security parameter. Informally, the theorem gives us a way to trade "breadth" (arity) for depth.

Theorem 3.1. *Let $\mathcal{E}_{\mathsf{ABE}}$ be a CP-ABE scheme with attribute space $\mathbb{A}$, message space $\mathbb{M}_{\mathsf{ABE}}$ and whose policy algebra $(\mathbb{F}, \odot)$ is an idempotent semigroup. Let κ be the security parameter. Let $m = \mathsf{poly}(\kappa)$. If $\mathcal{E}_{\mathsf{ABE}}$ is homomorphic for circuits of depth $O(\log^2 m)$, then there exists a secure CP-ABE scheme that is homomorphic for all circuits of arbitrary depth with at most m inputs.*

Proof. According to Theorem 4.5 in (López-Alt et al., 2012), there exists a multikey FHE scheme for m keys that is secure under the Ring Learning With Errors (RLWE) and Decisional Small Polynomial Ratio (DSPR) assumptions. Let $\mathcal{E}_{\mathsf{MKFHE}}$ be such a scheme. The dimension parameter of this scheme n is set such that it satisfies

$$m = n^{1-\delta}/log^{O(1)}n$$

for some $\delta \in (0, 1)$. It follows that $n = \mathsf{poly}(m)$ and thus $n = \mathsf{poly}(\kappa)$.

We can use $\mathcal{E}_{\mathsf{MKFHE}}$ and $\mathcal{E}_{\mathsf{ABE}}$ to construct a new CP-ABE scheme that is homomorphic for all m-ary circuits. We denote this class of circuits by $\mathbb{C}_{<m>}$. Indeed, any k-ary circuit with $k < m$ can be modelled as an m-ary circuit. Thus, w.l.o.g. we assume all circuits have m inputs.

The Setup and $\mathsf{Extract}$ algorithms remain unchanged from $\mathcal{E}_{\mathsf{ABE}}$. Encryption proceeds as follows for some $f \in \mathbb{F}$ and $m \in \mathbb{M}$ where $\mathbb{F}$ coincides with the class of access policies supported by $\mathcal{E}_{\mathsf{ABE}}$ and $\mathbb{M}$ coincides with the message space of $\mathcal{E}_{\mathsf{MKFHE}}$:

1. Generate $(\mathsf{pk}, \mathsf{sk}, \mathsf{vk}) \leftarrow \mathsf{MKFHE.Gen}(1^\kappa)$.

2. Compute $\psi \leftarrow \mathsf{ABE.Enc}(\mathsf{PP}, f, \mathsf{sk})$ - assume w.l.o.g. that $\mathsf{sk} \in \mathbb{M}_{\mathsf{ABE}}$.
3. Compute $c^\star \leftarrow \mathsf{MKFHE.Enc}(\mathsf{pk}, m)$.
4. Output $\mathbf{c} := (c^\star, \mathsf{vk}, \psi)$

The Eval algorithm is defined as follows:

1. On input $C \in \mathbb{C}_{<m>}$ and $\mathbf{c}_1, \ldots, \mathbf{c}_\ell$, assume that $\ell \leq m$ since otherwise the inputs $\mathbf{c}_i$ for $i > m$ are superfluous.
2. Parse $\mathbf{c}_i$ as $(c_i^\star, \mathsf{vk}_i, \psi_i)$ for $1 \leq i \leq \ell$.
3. Compute $c' \leftarrow \mathsf{MKFHE.Eval}(C, (c_1^\star, \mathsf{vk}_1), \ldots, (c_\ell^\star, \mathsf{vk}_\ell))$.
4. Set $f' := f_1 \odot \ldots \odot f_\ell$ where f_i is the policy under which the ciphertext $\mathbf{c}_i$ is encrypted for $1 \leq i \leq \ell$.
5. Compute $\psi' \leftarrow \mathsf{ABE.Enc}(\mathsf{PP}, f', c')$.
6. Compute $\psi^\star \leftarrow \mathsf{ABE.Eval}(D, \psi', \psi_1, \ldots, \psi_\ell)$ where D is the decryption circuit of $\mathcal{E}_{\mathsf{MKFHE}}$ (see below).
7. Output $(\perp, \perp, \psi^\star)$.

The Dec algorithm is defined as follows. On input a secret key sk_a for an attribute a and a ciphertext $\mathbf{c}$, perform the steps:

1. If $\mathbf{c}$ is of the form $(\perp, \perp, \psi^\star)$, output $\mathsf{ABE.Dec}(\mathsf{sk}_a, \psi^\star)$.
2. Otherwise, parse $\mathbf{c}$ as $(c^\star, \mathsf{vk}, \psi)$.
3. Compute $\mathsf{sk} \leftarrow \mathsf{ABE.Dec}(\mathsf{sk}_a, \psi)$.
4. If $\mathsf{sk} = \perp$, output $\perp$.
5. Output $\mathsf{MKFHE.Dec}(\mathsf{sk}, c^\star)$.

In the evaluation algorithm, the desired m-ary circuit C is evaluated using the multikey FHE algorithm. Observe that C can be of arbitrary depth since the size of the resultant multikey FHE ciphertext only depends on κ and m. We then encrypt this ciphertext with $\mathcal{E}_{\mathsf{ABE}}$ in order to homomorphically evaluate the decryption circuit of $\mathsf{E}_{\mathsf{MKFHE}}$ using $\mathcal{E}_{\mathsf{ABE}}$. Consequently, we obtain a ciphertext whose size is independent of m as required by the compactness condition for homomorphic CP-ABE. It remains to be shown that $\mathcal{E}_{\mathsf{ABE}}$ has the homomorphic capacity to evaluate D.

Lemma 4.4 in (López-Alt et al., 2012) established that the decryption circuit for the multikey FHE scheme presented therein can be realized as a polynomial-sized circuit of depth $d = O(\log m \cdot (\log\log q + \log n))$ where q is the modulus in the multikey FHE scheme. Now set $q = 2^{m \log m \cdot \log^2 n}$ (this setting allows Theorem 4.5 in (López-Alt et al., 2012) to go through). Recall that $n = \mathsf{poly}(m)$. Therefore, $d = O(log^2(m))$ which is a depth that is supported by $\mathcal{E}_{\mathsf{ABE}}$. This completes the proof. □

Fortunately due to the recent results presented in (Gentry et al., 2013), the precondition for Theorem 3.1 can be satisfied, at least with respect to IBE in the single-identity setting.

4 Policy-Based Homomorphic Encryption

Our approach is applicable to policy algebras $(\mathbb{F}, \odot)$ where the policy size always grows with composition. An example of such an algebra is the free semigroup $\mathfrak{F}^*$ on a set $\mathfrak{F}$. Moreover, our approach can handle at most N compositions where N is the maximum number of independent users supported by the multikey FHE scheme $\mathcal{E}_{\mathsf{MKFHE}}$. Observe that the inputs encrypted by the same user under the same policy need *not* be composed together with $\odot$. Therefore, the scheme can handle more than N inputs, but at most N independent policies. However, the syntax of homomorphic CP-ABE is too limited to capture this exemption. This fact serves to motivate the formulation of a more general primitive which we refer to as policy-based homomorphic encryption (PBHE). Our formulation of PBHE is influenced considerably by the definition of multikey FHE, and inherits many of its properties.

Definition 4.1. *A Policy-Based Homomorphic Encryption (PBHE) scheme for a class of circuits $\mathbb{C}$, a commutative semigroup of access policies $(\mathbb{F}/\sim, \odot)$ and a set of attributes $\mathbb{A}$ is a family of algorithms $\{\mathcal{E}^{(N)} := (\mathsf{Setup}, \mathsf{Extract}, \mathsf{GenKey}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})\}_{N \geq 1}$ where $\mathcal{E}^{(N)}$ is defined as follows:*

- $(\mathsf{PP}, \mathsf{MSK}) \leftarrow \mathsf{Setup}(1^\kappa)$*: Given a security parameter κ, output public parameters* PP *and a master secret key* MSK.
- $\mathsf{sk}_a \leftarrow \mathsf{Extract}(\mathsf{MSK}, a)$*: Given a master secret key* MSK *and an attribute $a \in \mathbb{A}$, output a secret key sk_a for a.*
- $(\mathsf{ek}_f, \mathsf{vk}_f) \leftarrow \mathsf{GenKey}(\mathsf{PP}, f)$*: Given public parameters* PP *and an access policy $f \in \mathbb{F}$, output a pair of encryption and evaluation keys $(\mathsf{ek}_f, \mathsf{vk}_f)$ for f.*
- $c \leftarrow \mathsf{Enc}(\mathsf{PP}, \mathsf{ek}_f, m)$*: Given public parameters* PP*, an encryption key ek_f for policy f, and a plaintext $m \in \mathbb{M}$, output a ciphertext c that encrypts m under policy f.*
- $m \leftarrow \mathsf{Dec}(\mathsf{sk}_a, c)$*: Given a secret key sk_a for attribute $a \in \mathbb{A}$ and a ciphertext c that encrypts a message m under access policy f, output m iff $f(a) = 1$ and $\perp$ otherwise.*
- $c' \leftarrow \mathsf{Eval}(C, (c_1, \mathsf{vk}_1), \ldots, (c_\ell, \mathsf{vk}_\ell))$*: Given a circuit $C \in \mathbb{C}$ and a sequence of ℓ pairs of ci-*

phertext and evaluation keys, output a ciphertext c'.

For every $(\mathsf{PP}, \mathsf{MSK}) \leftarrow \mathsf{Setup}(1^\kappa)$, every collection of $t \leq N$ access policies $f_1, \ldots, f_t \in \mathbb{F}$, and every collection of key-pairs $\mathsf{K} := \{(\mathsf{ek}_i, \mathsf{vk}_i) \leftarrow \mathsf{GenKey}(\mathsf{PP}, f_i)\}_{i\in[t]}$, every sequence of ℓ tuples $\{(c_i, \mathsf{vk}_{v_i}) : v_i \in [t], c_i \leftarrow \mathsf{Enc}(\mathsf{PP}, \mathsf{ek}_{v_i}, m_i)\}_{i\in[\ell]}$ and all attributes $a \in \mathbb{A}$ and secret keys $\mathsf{sk}_a \leftarrow \mathsf{Extract}(\mathsf{MSK}, a)$, and all circuits $C \in \mathbb{C}$, the following properties are satisfied for every $c' = \mathsf{Eval}(C, (c_1, \mathsf{vk}_{v_1}), \ldots, (c_\ell, \mathsf{vk}_{v_\ell}))$ where $f' = \bigodot_{j\in\{v_i,\ldots,v_\ell\}} f_j$:

- **Correctness:**
  1. $\mathsf{vk}_i = \mathsf{vk}_j \Rightarrow f_i = f_j \quad$ for $i, j \in [t]$.
  2. $\mathsf{Dec}(\mathsf{sk}_a, c') = C(m_1, \ldots, m_\ell)$ iff $f'(a) = 1$ and $\perp$ otherwise.
- **Compactness:** $|c'| = \mathsf{poly}(\kappa, |f'|, N)$.[†]

Informally, the first correctness condition requires that evaluation keys be uniquely associated with an access policy. Besides including information necessary for evaluation, which could instead be embedded in the ciphertext, the main role of an evaluation key is to allow ciphertexts produced by the same encryptor to be grouped together into classes. The composition operation is not applied *among* the members of such classes according to the second correctness condition. In other words, composition is performed on equivalence classes where the equivalence relation is defined by equality of evaluation keys. The motivation for this is to compensate for the non-idempotency of an operation $\odot$. For example, it may be the case that the ciphertexts produced by the same encryptor share information that can be exploited to assist homomorphic computation among them. This is exemplified by multikey FHE.

Security The security definition for PBHE is similar to the security definition of CP-ABE. We also refer to this as IND-AD-CPA security. In fact, the security game is the same as that for CP-ABE except that the adversary is also given $(\mathsf{ek}, \mathsf{vk}) \leftarrow \mathsf{PBHE.GenKey}(\mathsf{PP}, f^\star)$ after it chooses a target policy $f^\star$.

5 Construction of PBHE

In this section, we construct a new generic PBHE scheme that can be instantiated by an IND-AD-CPA secure CP-ABE scheme together with any IND-CPA secure multikey FHE scheme.

Remark Concrete constructions of CP-ABE and multi-key FHE already exist which fulfill the properties we need. Examples of the former include (Bethencourt et al., 2007; Waters, 2011) and an example of the latter is the NTRU-based construction from (López-Alt et al., 2012).

Let $\mathcal{E}_{\mathsf{ABE}} = (\mathsf{ABE.Setup}, \mathsf{ABE.Extract}, \mathsf{ABE.Enc}, \mathsf{ABE.Dec})$ be a CP-ABE scheme for a class of policies $\mathbb{F}_{\mathsf{ABE}}$, a set of attributes $\mathbb{A}_{\mathsf{ABE}}$ and a message space $\mathbb{M}_{\mathsf{ABE}}$. Let $\{\mathcal{E}^{(N)}_{\mathsf{MKFHE}} = (\mathsf{MKFHE.Gen}, \mathsf{MKFHE.Enc}, \mathsf{MKFHE.Dec}, \mathsf{MKFHE.Eval})\}_{N>0}$ be a family of multikey fully-homomorphic encryption schemes. In our generic PBHE scheme, the set of attributes $\mathbb{A}$ is defined as $\mathbb{A} \triangleq \mathbb{A}_{\mathsf{ABE}}$. Now we need to define an algebraic structure of access policies $(\mathbb{F}, \odot)$ that obeys the three properties given by 3.3, 3.4 and 3.5.

5.1 Supported Access Policies and Composition

Define a subset $\mathfrak{F} \subseteq \mathbb{F}_{\mathsf{ABE}}$ that is closed under $\wedge$. We define $(\mathbb{F}, \odot)$ as the free semigroup $\mathfrak{F}^*$ on $\mathfrak{F}$ i.e. the set of finite strings composed of elements of $\mathfrak{F}$. For brevity, we will use juxtaposition instead of explicitly writing $\odot$ when representing policies in $\mathbb{F}$. The semantic interpretation of a policy $\mathfrak{f}_1 \ldots \mathfrak{f}_\ell \in \mathbb{F}$ is such that the following holds

$$\mathfrak{f}_1 \ldots \mathfrak{f}_\ell \sim \mathfrak{f}_1 \wedge \ldots \wedge \mathfrak{f}_\ell.$$

Moreover, the size $|f|$ of a policy $f = \mathfrak{f}_1 \ldots \mathfrak{f}_\ell$ is the sum $\sum_{i=1}^{\ell} |\mathfrak{f}_i|$. This is to be distinguished from the *length* of a policy in $f \in \mathbb{F}$, written $\lambda(f)$, which is the length of the corresponding string of elements from $\mathfrak{F}$.

Note that any policy $f \in \mathbb{F}$ can be transformed into a semantically equivalent policy f' with $\lambda(f') = 1$. Thus, we assume without loss of generality that this is what the GenKey algorithm takes as input.

5.2 Key Trees

Now we show how our PBHE scheme enforces access policies in $\mathbb{F}$ and how it handles composition. By abuse of notation, we write $\mathsf{ABE.Enc}(\mathsf{PP}, f, M)$ for $M \notin \mathbb{M}_{\mathsf{ABE}}$ to signify the encryption of multiple elements of $\mathbb{M}_{\mathsf{ABE}}$ in order to "cover" M. The analogous notion is also assumed for decryption.

Let f be a policy in $\mathbb{F}$. Our approach involves mapping f to a binary tree τ_f (which we call a *key tree*) whose nodes are associated with ciphertexts

[†] In the proceedings version, the compactness condition was given as $|c'| = \mathsf{poly}(\kappa, |f'|)$. Dependence on N was omitted on assumption that $N = \mathsf{poly}(\kappa)$. However, to be consistent with (López-Alt et al., 2012), N has been made explicit as a free parameter.

in the CP-ABE scheme. Each leaf node corresponds to an element of $\mathfrak{F}$ while an interior node corresponds to the conjunction of its left and right branches. More precisely, the leaf nodes are encryptions of the secret keys for $\mathcal{E}_{\mathsf{MKFHE}}$ XORed with a random *blinding* string. An interior node encrypts the concatenation of the blinding strings of its child nodes XORed with a new blinding string. Thus, in order to access the secret keys at the leaves, it is necessary to decrypt from the root down. Thus, if a user's attribute satisfies the root policy, she can decrypt every layer and eventually recover the secret keys hidden by the ciphertexts at the leaves. Indeed, satisfying the root policy is a sufficient and necessary condition to recover *any* secret key.

Remark Using a variable-degree tree would make for more space-efficient ciphertexts. However, the algorithms are easier to describe by using a binary tree. Furthermore, the ciphertexts have a unified structure i.e. the result of a multi-stage evaluation is structurally equivalent to that from a single joint evaluation. The construction in Section 7 uses variable-degree trees.

A key tree τ_f for a policy $f := \mathfrak{f}_1 \dots \mathfrak{f}_\ell$ consists of a list of CP-ABE ciphertexts $[\psi_i]_{i \le 2\ell-1}$[‡]. We associate with each leaf node in τ_f a unique key tuple $(\mathsf{pk}, \mathsf{sk}, \mathsf{vk}) \leftarrow \mathsf{MKFHE.Gen}(1^\kappa)$ in the multikey FHE scheme. Roughly speaking, we set the leaf node of τ_f to an encryption of $\mathsf{sk} \oplus r$ under $\mathfrak{f}$ in the CP-ABE scheme where r is a random string of length $|\mathsf{sk}|$. Every tree is associated with such a random value, and thus for convenience, we will sometimes refer to a pair $(r, [\psi_1]_{i \le 2^{h+1}-1})$ as a "tree".

Now to construct a key tree for a policy f, consider an algorithm $\mathsf{MkTree}^\star$ which proceeds as follows:

1. On input $f \in \mathbb{F}$, decompose f into $\mathfrak{f}_1 \dots \mathfrak{f}_k$.
2. For $1 \le i \le k$:
   (a) Set $(\mathsf{pk}_i, \mathsf{sk}_i, \mathsf{vk}_i) \leftarrow \mathsf{MKFHE.Gen}(1^\kappa)$.
   (b) Uniformly sample $r_i \xleftarrow{\$} \{0,1\}^{|\mathsf{sk}_i|}$.
   (c) Compute $\psi_i \leftarrow \mathsf{ABE.Enc}(\mathsf{PP}, \mathfrak{f}_i, r_i \oplus \mathsf{sk}_i)$.
   (d) Set $\tau_i \leftarrow (r_i, [\psi_i])$.
3. For $1 \le i \le \lceil \lg k \rceil$:
   (a) For $1 \le j \le \lceil k/2^i \rceil$:
      i. If $2j > \lceil k/2^{i-1} \rceil$, set $\tau_j \leftarrow \tau_{2j-1}$.
      ii. Else set $\tau_j \leftarrow \mathsf{Combine}(\tau_{2j-1}, \tau_{2j})$
4. Output $(\tau_1, (\mathsf{pk}_1, \mathsf{vk}_1), \dots, (\mathsf{pk}_k, \mathsf{vk}_k))$.

where $\mathsf{Combine}$ is defined below.

[‡]The notation $[a_1, \dots, a_t]$ denotes a "list"; that is, the sequence $a_1, \dots, a_t$.

Note that we denote by MkTree the variant of $\mathsf{MkTree}^\star$ that outputs only the first component of the tuple outputted by $\mathsf{MkTree}^\star$, namely the tree τ_1.

To combine two trees $\tau_f := (r, [\chi_i]_{i \le 2\ell_1 - 1})$ and $\tau_g := (s, [\psi_i]_{i \le 2\ell_2 - 1})$ for policies $f := \mathfrak{f}_1 \dots \mathfrak{f}_{\ell_1}$ ($g := \mathfrak{g}_1 \dots \mathfrak{g}_{\ell_2}$ resp.), the following algorithm is used (we refer to this algorithm as $\mathsf{Combine}$):

1. Uniformly sample $t \xleftarrow{\$} \{0,1\}^{|r \| s|}$.
2. Compute $\omega \leftarrow \mathsf{ABE.Enc}(\mathsf{PP}, \mathfrak{f}' \wedge \mathfrak{g}', t \oplus (r \parallel s))$ where $\mathfrak{f}' = \mathfrak{f}_1 \wedge \dots \wedge \mathfrak{f}_{\ell_1}$ and $\mathfrak{g}' = \mathfrak{g}_1 \wedge \dots \wedge \mathfrak{g}_{\ell_2}$.
3. Construct the tree $(t, [\omega, \chi_1, \dots, \chi_{2\ell_1-1}, \psi_1, \dots, \psi_{2\ell_2-1}])$.

Decrypting a tree with a secret key sk_a for an attribute $a \in \mathbb{A}$ is defined recursively:

- $\mathsf{DecTree}(\mathsf{sk}_a, (t, [\omega])) = \begin{cases} [\pi \oplus t] & \text{if } \pi \neq \bot \\ \bot & \text{otherwise} \end{cases}$
  where $\pi = \mathsf{ABE.Dec}(\mathsf{sk}_a, \omega)$.
- $\mathsf{DecTree}(\mathsf{sk}_a, (t_1 \parallel t_2, [\omega] \parallel [\chi_i]_{i \le 2\ell_1 - 1} \parallel [\psi_i]_{i \le 2\ell_2 - 1})) =$

$$\begin{cases} \mathsf{DecTree}(\mathsf{sk}_a, (r \oplus t_1, [\chi_i]_{i \le 2\ell_1 - 1})) \parallel \\ \quad \mathsf{DecTree}(\mathsf{sk}_a, (s \oplus t_2, [\psi_i]_{i \le 2\ell_2 - 1})) & \text{if } \pi = (r \parallel s) \\ \bot & \text{if } \pi = \bot \end{cases}$$

  where $\pi = \mathsf{ABE.Dec}(\mathsf{sk}_a, \omega)$.

Therefore, $\mathsf{DecTree}$ produces either $\bot$ or a list of secret keys for $\mathcal{E}_{\mathsf{MKFHE}}$.

5.3 Basic Construction

For brevity, we will assume that all policies f passed as input to GenKey satisfy $\lambda(f) = 1$. Our PBHE scheme $\mathcal{E}_{\mathsf{PBHE}}$ is defined as follows:

- $\mathsf{Setup}(1^\kappa)$: Given a security parameter κ, generate $(\mathsf{PP}, \mathsf{MSK}) \leftarrow \mathsf{ABE.Setup}(1^\kappa)$. Output $(\mathsf{PP}, \mathsf{MSK})$.
- $\mathsf{Extract}(\mathsf{MSK}, a)$: Given a master secret key MSK and an attribute $a \in \mathbb{A}$, output $\mathsf{sk}_a \leftarrow \mathsf{ABE.Extract}(\mathsf{MSK}, a)$.
- $\mathsf{GenKey}(\mathsf{PP}, f)$: Given public parameters PP and an access policy $f \in \mathbb{F}$, run:
  1. On assumption (above) f can be parsed as $\mathfrak{f}$ where $\mathfrak{f} \in \mathfrak{F}$.
  2. Compute $(\tau, (\mathsf{pk}, \mathsf{vk})) \leftarrow \mathsf{MkTree}^\star(f)$.
  3. Set $\mathsf{ek} \leftarrow ((\mathsf{pk}, \mathsf{vk}), \tau)$.
  4. Output $(\mathsf{ek}, \mathsf{vk})$.
- $\mathsf{Enc}(\mathsf{PP}, \mathsf{ek}_f, m)$: Given public parameters PP, an encryption key ek_f for policy f, and a plaintext $m \in \mathbb{M}$, run:

1. Parse ek_f as $((\mathsf{pk}, \mathsf{vk}), \tau)$.
2. Compute $c^* \leftarrow \mathsf{MKFHE.Enc}(\mathsf{pk}, m)$.
3. Output $\mathbf{c} := (c^*, \tau)$.

– $\mathsf{Dec}(\mathsf{sk}_a, \mathbf{c})$: Given a secret key sk_a for attribute $a \in \mathbb{A}$ and a ciphertext $\mathbf{c}$ that encrypts a message m under access policy f, run:

1. Parse $\mathbf{c}$ as (c^*, τ).
2. If $\mathsf{DecTree}(\mathsf{sk}_a, \tau) = \perp$, then output $\perp$ and abort.
3. Set $(\mathsf{sk}_1, \ldots, \mathsf{sk}_k) \leftarrow \mathsf{DecTree}(\mathsf{sk}_a, \tau)$
4. Compute $m \leftarrow \mathsf{MKFHE.Dec}(\mathsf{sk}_1, \ldots, \mathsf{sk}_k, c^*)$.
5. Output m.

– $\mathsf{Eval}(C, (\mathbf{c_1}, \mathsf{vk}_1), \ldots, (\mathbf{c}_\ell, \mathsf{vk}_\ell))$: Given a circuit $C \in \mathbb{C}$ and a sequence of ℓ pairs of ciphertext and evaluation keys, perform the following steps:

1. Parse each $\mathbf{c_i}$ as (c_i^*, τ_i).
2. Set $\mathfrak{T} := \{\tau_i\}_{i \in [\ell]}$ (recall that ciphertexts encrypted under the same vk have the same τ component).
3. Run $\mathsf{Combine}$ to recursively build a tree τ from all elements in $\mathfrak{T}$.
4. Set $c^* \leftarrow \mathsf{MKFHE.Eval}(C, (c_1, \mathsf{vk}_1), \ldots, (c_\ell, \mathsf{vk}_\ell))$.
5. Output (c^*, τ).

Theorem 5.1. *If $\mathcal{E}_{\mathsf{ABE}}$ is an IND-AD-CPA-secure CP-ABE scheme and $\mathcal{E}^{(N)}_{\mathsf{MKFHE}}$ is an IND-CPA-secure multikey FHE scheme, then $\mathcal{E}^{(N)}_{\mathsf{PBHE}}$ is IND-AD-CPA-secure.*

Proof. We prove the theorem by means of a hybrid argument.

Hybrid 0 IND-AD-CPA game for $\mathcal{E}_{\mathsf{PBHE}}$.

Hybrid 1 Same as Hybrid 0 except with one difference. Let $f^\star \in \mathbb{F}$ be the target policy chosen by the adversary $\mathcal{A}$. It can be assumed w.l.o.g. that $f^\star = \mathfrak{f}$ for some $\mathfrak{f} \in \mathfrak{F}$. The challenger uses a modified MkTree algorithm to compute the CP-ABE ciphertext corresponding to $\mathfrak{f}$ by running $\psi \leftarrow \mathsf{ABE.Enc}(\mathsf{PP}, \mathfrak{f}, 0^{|\mathsf{sk}|})$ where $0^{|\mathsf{sk}|}$ is a string of zeros whose length is the same as the multikey FHE secret key generated for $\mathfrak{f}$. The algorithm is otherwise unchanged.

We claim that any poly-time $\mathcal{A}$ that can distinguish between Hybrid 0 and Hybrid 1 with a non-negligible advantage can break the IND-AD-CPA security of $\mathcal{E}_{\mathsf{ABE}}$. An adversary $\mathcal{B}$ that uses $\mathcal{A}$ proceeds as follows. When $\mathcal{A}$ chooses a target policy $f^\star := \mathfrak{f}$, $\mathcal{B}$ runs MkTree. Then it gives $\mathfrak{f}$ to its IND-AD-CPA challenger along with two messages $m_0 := \mathsf{sk} \oplus r$ and $m_1 := 0^{|\mathsf{sk}|}$ where r is a random string of length $|\mathsf{sk}|$. Note that we assume for simplicity that both messages are in $\mathbb{M}_{\mathsf{ABE}}$; if multiple messages (say L) are required then the usual hybrid argument can be applied which loses a factor of L. Subsequently, $\mathcal{B}$ embeds the challenge CP-ABE ciphertext as ψ. Therefore, if ψ encrypts m_0, then $\mathcal{B}$ perfectly simulates Hybrid 0. Otherwise, $\mathcal{B}$ perfectly simulates Hybrid 1. Thus, if $\mathcal{A}$ has a non-negligible advantage distinguishing between the hybrids, then $\mathcal{B}$ has a non-negligible advantage attacking the IND-AD-CPA security of $\mathcal{E}_{\mathsf{ABE}}$.

Hybrid 2 Same as Hybrid 1 with the exception that the challenger does not encrypt either message m_0 or m_1 (using $\mathcal{E}_{\mathsf{MKFHE}}$) chosen by $\mathcal{A}$. Instead it encrypts some fixed message $\mu \in \mathbb{M}$. Therefore, the adversary has a zero advantage in this game because the challenge ciphertext contains no information about the challenger's bit.

We now show that if $\mathcal{A}$ can efficiently distinguish between Hybrid 1 and Hybrid 2, then there is a PPT algorithm $\mathcal{G}$ that can use $\mathcal{A}$ to attack the IND-CPA security of $\mathcal{E}_{\mathsf{MKFHE}}$. When $\mathcal{A}$ chooses m_0 and m_1, $\mathcal{G}$ simply gives m_b and μ to its IND-CPA challenger where b is the bit it uniformly samples in its simulation of the IND-AD-CPA challenger. It then sets the first component of the returned ciphertext $\mathbf{c}$ according to the challenge ciphertext $c^\star$ it receives. If $c^\star$ encrypts m_b, then the view of $\mathcal{A}$ is identical to Hybrid 1. Otherwise, the view of $\mathcal{A}$ is identical to Hybrid 2. Therefore, a non-negligible advantage obtained by $\mathcal{A}$ implies a non-negligible advantage for $\mathcal{G}$ in the IND-CPA game, and thus contradicts the IND-CPA security of $\mathcal{E}_{\mathsf{MKFHE}}$. □

6 Simulation Model of Evaluation

Consider ciphertexts $c_1, \ldots, c_\ell$ encrypted under $f_1, \ldots, f_\ell$ respectively. We expect that a ciphertext c' resulting from an evaluation on $c_1, \ldots, c_\ell$ be an encryption under the composite policy $f' = f_1 \odot \ldots \odot f_\ell$. Now suppose an adversary has access to secret keys $\mathsf{sk}_{a_1}, \ldots, \mathsf{sk}_{a_\ell}$ for attributes $a_1, \ldots, a_\ell$ respectively. Furthermore, suppose that for $1 \leq i \leq \ell$, $f_i(a_i) = 1$ and $f'(a_i) = 0$. Then it is required that the adversary be unable to decrypt c'. This property holds for our basic construction in Section 5.3 provided that the adversary does not have access to all of $c_1, \ldots, c_\ell$. To attain the property without such a constraint is more challenging. Ideally, a user that does not have a secret key for an attribute satisfying f' should not learn anything about what c' encrypts. In fact, a stronger requirement is that such a user should not be able to efficiently decide whether c' was produced from $c_1, \ldots, c_\ell$ or an alternative sequence of ℓ ciphertexts $d_1, \ldots, d_\ell$ with d_i encrypted under f_i for

$1 \leq i \leq \ell$.

A PBHE scheme is said to be *EVAL-SIM*-secure if for every adversary $\mathcal{A} := (\mathcal{A}_1, \mathcal{A}_2)$, which is a pair of PPT algorithms, there are no polynomial-time algorithms with a non-negligible advantage distinguishing the following two distributions.

Real Distribution:

1. $(\mathsf{PP}, \mathsf{MSK}) \leftarrow \mathsf{Setup}(1^\lambda)$.
2. $(C, f_1, \ldots, f_k, (v_1, m_1), \ldots, (v_\ell, m_\ell), \mathsf{state}) \leftarrow \mathcal{A}_1^{\mathsf{Extract}(\mathsf{MSK}, \cdot)}(\mathsf{PP})$.
3. Output $\perp$ and abort if $\{v_1, \ldots, v_\ell\} \neq [k]$.
4. Let $a_1, \ldots, a_q$ be the attributes queried by $\mathcal{A}_1$.
5. Let $f' = \bigodot_{j \in \{v_1, \ldots, v_\ell\}} f_j$.
6. Output $\perp$ and abort if for any $1 \leq i \leq q$: $f'(a_i) = 1$.
7. $(\mathsf{ek}_i, \mathsf{vk}_i) \leftarrow \mathsf{GenKey}(\mathsf{PP}, f_i)$ for $1 \leq i \leq k$.
8. $c_j \leftarrow \mathsf{Enc}(\mathsf{PP}, \mathsf{ek}_{v_j}, m_j)$ for $1 \leq j \leq \ell$.
9. $c' \leftarrow \mathsf{Eval}(\mathsf{PP}, C, (c_1, \mathsf{vk}_{v_1}), \ldots, (c_\ell, \mathsf{vk}_{v_\ell}))$.
10. $$\alpha \leftarrow \mathcal{A}_2^{\mathsf{Extract}(\mathsf{MSK}, \cdot)|_{A'}}(\mathsf{state}, c', (\mathsf{ek}_1, \mathsf{vk}_1), \ldots, (\mathsf{ek}_k, \mathsf{vk}_k), c_1, \ldots, c_\ell)$$ where $A' = \{a \in \mathbb{A} \mid f'(a) = 0\}$.
11. Output α.

Ideal Distribution:

1. $(\mathsf{PP}, \mathsf{MSK}) \leftarrow \mathsf{Setup}(1^\lambda)$.
2. $(C, f_1, \ldots, f_k, (v_1, m_1), \ldots, (v_\ell, m_\ell), \mathsf{state}) \leftarrow \mathcal{A}_1^{\mathsf{Extract}(\mathsf{MSK}, \cdot)}(\mathsf{PP})$.
3. Output $\perp$ and abort if $\{v_1, \ldots, v_\ell\} \neq [k]$.
4. Let $a_1, \ldots, a_q$ be the attributes queried by $\mathcal{A}_1$.
5. Let $f' = \bigodot_{j \in \{v_1, \ldots, v_\ell\}} f_j$.
6. Output $\perp$ and abort if for any $1 \leq i \leq q$: $f'(a_i) = 1$.
7. $(\mathsf{ek}_i, \mathsf{vk}_i) \leftarrow \mathsf{GenKey}(\mathsf{PP}, f_i)$ for $1 \leq i \leq k$.
8. $c_j \leftarrow \mathsf{Enc}(\mathsf{PP}, \mathsf{ek}_{v_j}, m_j)$ for $1 \leq j \leq \ell$.
9. $c' \leftarrow \mathcal{S}(\mathsf{PP}, C, f_1, \ldots, f_k)$.
10. $$\alpha \leftarrow \mathcal{A}_2^{\mathsf{Extract}(\mathsf{MSK}, \cdot)|_{A'}}(\mathsf{state}, c', (\mathsf{ek}_1, \mathsf{vk}_1), \ldots, (\mathsf{ek}_k, \mathsf{vk}_k), c_1, \ldots, c_\ell)$$ where $A' = \{a \in \mathbb{A} \mid f'(a) = 0\}$.
11. Output α.

7 Main Construction

Our goal in this section is to present a construction that is EVAL-SIM-secure.

7.1 Prerequisites

Like the construction in Section 5, our construction here relies on both a CP-ABE and multikey FHE scheme. However, some additional properties are needed.

7.1.1 Weak Group Homomorphic CP-ABE

In the case of CP-ABE, we need the scheme to be homomorphic with respect to a group operation. Hence, the message space $\mathbb{G}$ is expected to be a group. We use multiplicative notation to denote the group operation. Furthermore, the operation on the ciphertext space is denoted by $\boxdot$. We require that given any $m \in \mathbb{G}$, $f \in \mathbb{F}$ and any $c \leftarrow E_f(m)$, the distributions

$$\{(c \boxdot E_f(r)) \mid r \xleftarrow{\$} \mathbb{G}\} \underset{C}{\approx} \{(E_f(r_1) \boxdot E_f(r_2) \mid r_1, r_2 \xleftarrow{\$} \mathbb{G}\} \tag{7.1}$$

are computationally indistinguishable, where E_f denotes $\mathsf{ABE.Enc}(\mathsf{PP}, f, \cdot)$. Intuitively, this property means that an evaluation using a known ciphertext produces a ciphertext that "looks like" (computationally) the same evaluation using instead an encryption of a random element of the message space. Group Homomorphic Encryption (GHE) formalized in (Armknecht et al., 2010) trivially satisfies this property, but we don't need such a strong definition. For our purposes, it suffices that the scheme satisfies

$$D_a(E_f(x) \boxdot E_f(y)) = x \cdot y \tag{7.2}$$

for any $x, y \in \mathbb{G}$ and $a \in \mathbb{A}$ s.t. $f(a) = 1$, where D_a denotes the decryption algorithm with a secret key for attribute a. We say that a scheme satisfying both properties is "weakly" group homomorphic (WGH).

Consider a CP-ABE scheme $\mathcal{E}_{\mathsf{ABE}}$ that is WGH. Let t be the length of a secret key for $\mathcal{E}_{\mathsf{MKFHE}}$ when considered as a bit string. We assume there is a natural encoding $T : \{0,1\}^t \to \mathbb{G}$. Suppose the order of $\mathbb{G}$ is too small to accommodate such an encoding. This can be easily addressed. Observe that the binary string can be encoded as an element of $\mathbb{G}^{t'}$ for some t' that is polynomial in the security parameter. Consequently, we can consider a CP-ABE scheme that is WGH with respect to the group $\mathbb{G}^{t'}$ (i.e. the

direct product of t' copies of $\mathbb{G}$). Such a CP-ABE scheme inherits its semantic security by the usual hybrid argument from the semantic security of $\mathcal{E}_{\mathsf{ABE}}$. Therefore, we assume that $\mathcal{E}_{\mathsf{ABE}}$ is a WGH CP-ABE scheme with respect to a group $\mathbb{G}$ that accommodates the encoding T.

7.1.2 Multikey Privacy

We require the multikey FHE scheme $\mathcal{E}_{\mathsf{MKFHE}}$ to satisfy *multikey privacy*. Informally, this means that an attacker cannot distinguish which of two known sets of public keys was used to encrypt a given ciphertext provided both sets have the same cardinality and both sets contain at least one public key whose corresponding secret key is unknown to the attacker. The formal security game is captured in the following experiment.

Let $\mathcal{O}$ be an oracle that returns a key tuple $(\mathsf{pk}, \mathsf{sk}, \mathsf{vk}) \leftarrow \mathsf{GenKey}(1^\lambda)$ when queried for an index $i \in \mathbb{N}$. It returns the same response when queried on the same index. Similarly, let $\mathcal{O}'$ be an oracle that returns a key tuple $(\mathsf{pk}, \mathsf{vk})$ where $(\mathsf{pk}, \mathsf{sk}, \mathsf{vk}) \leftarrow \mathsf{GenKey}(1^\lambda)$. Both oracles generate fresh keys for $\mathcal{E}_{\mathsf{MKFHE}}$ with $\mathcal{O}$ providing both public and secret information associated with the key, and $\mathcal{O}'$ providing only public information.

Experiment $\mathsf{MKPriv}(\mathcal{A}_1, \mathcal{A}_2)$:

1. $$(\mathsf{state}, C, m_1, \ldots, m_\ell, v_{0,1}, \ldots, v_{0,\ell}, v_{1,1}, \ldots, v_{1,\ell}) \leftarrow \mathcal{A}_1^{\mathcal{O},\mathcal{O}'}(1^\lambda)$$
2. Suppose $\mathcal{A}_1$ makes a total of $Q = q + q'$ queries. Assume w.l.o.g. that $\mathcal{A}_1$ queries $\mathcal{O}$ on $1, \ldots, q$ to yield $(\mathsf{pk}_i, \mathsf{sk}_i, \mathsf{vk}_i)$ for $1 \le i \le q$, and it queries $\mathcal{O}'$ on $q+1, \ldots, Q$ to yield $(\mathsf{pk}_i, \mathsf{sk}_i)$ for $q+1 \le i \le Q$.
3. Abort with a random bit unless the following conditions are met for $i \in \{0,1\}$:
   (a) $v_{i,1}, \ldots, v_{i,\ell} \in [Q]$.
   (b) $v_{i,j} > q$ for some j (this implies that $q' \ge 1$ and at least one key to be used in evaluation came from $\mathcal{O}'$).
4. Generate a uniformly random bit $b \xleftarrow{\$} \{0,1\}$.
5. Compute $c_{i,j} \leftarrow \mathsf{Enc}(\mathsf{pk}_{v_{i,j}}, m_j)$ for $i \in \{0,1\}$ and $j \in [\ell]$.
6. Compute $c^* \leftarrow \mathsf{Eval}(C, (c_{b,1}, \mathsf{vk}_{v_{b,1}}), \ldots, (c_{b,\ell}, \mathsf{vk}_{v_{b,\ell}}))$.
7. $b' \leftarrow \mathcal{A}_2(\mathsf{state}, c^*, c_{0,1}, \ldots, c_{0,\ell}, c_{1,1}, \ldots, c_{1,\ell})$.
8. Output 1 if $b' = b$ and output 0 otherwise.

A multikey FHE scheme is said to be *multikey-private* if for any pair of PPT algorithms $(\mathcal{A}_1, \mathcal{A}_2)$, it holds that

$$\mathsf{Pr}[\mathsf{MKPriv}(\mathcal{A}_1, \mathcal{A}_2) \Rightarrow 1] - \frac{1}{2} < \mathsf{negl}(\lambda).$$

Observe that this formulation of multikey FHE privacy requires Eval to be nondeterministic. Otherwise, it is trivial for an adversary to guess the challenger's random coin by merely calling Eval with both sequences of ciphertexts.

Lemma 7.1. *There exists a variant of the multikey FHE scheme from (López-Alt et al., 2012) that is multikey-private under the Decisional Small Polynomial Ratio (DSPR) and Ring Learning With Errors (R-LWE) assumptions.*

Proof. Ciphertexts in the scheme presented in (López-Alt et al., 2012) are polynomials in a ring $R_q = \mathbb{Z}_q[x]/f(x)$. That scheme employs the technique of modulus reduction proposed in (Brakerski et al., 2012). It uses a ladder of modulii $q_0, \ldots, q_L$ with $q_0 > \ldots > q_L$, which can be set such that $q_i = 2^{\Omega(((L+1)-i)\cdot\mu)}$ for some μ. Setting μ to $N \log N \cdot \log^2 n$ satisfies Theorem 4.5 in (López-Alt et al., 2012), and yields a secure bootstrappable scheme for $N < n^\epsilon / \log^{O(1)} n$ users.

Now for appropriate parameters, "fresh" ciphertexts in the scheme are computationally indistinguishable from uniformly random elements in R_{q_0} under the DSPR and R-LWE assumptions provided an attacker does not have access to the secret key.

Let $c_1, c_2 \in R_{q_i}$ be two ciphertexts at "level" i with $0 \le i < L$. Assume that both c_1 and c_2 are uniformly distributed in R_{q_i} (computationally). For the remainder of the proof, we use the term "uniformly distributed" in the "computational" sense with the assumption that an attacker does not have all secret keys to perform decryption.

The scheme describes algorithms to add and multiply ciphertexts. Let $c_{\mathsf{add}} \in R_{q_{i+1}}$ be a ciphertext outputted by the addition algorithm. Similarly, let $c_{\mathsf{mul}} \in R_{q_{i+1}}$ be a ciphertext outputted by the multiplication algorithm. We claim that both c_{add} and c_{mul} are uniformly distributed in $R_{q_{i+1}}$ (computationally) to an attacker who (1) does not have all secret keys needed to decrypt the resultant ciphertexts and (2) is not given both c_1 and c_2. Proving this claim is sufficient to prove the lemma since in our Eval algorithm, we can simply add an encryption of zero to rerandomize the ciphertext. As long as the attacker cannot decrypt the ciphertext, he cannot efficiently learn anything else about which key set was used to generate it.

Now we proceed to prove the claim. Firstly, we obtain $c_0 := c_1 + c_2$ or $c_0 := c_1 \cdot c_2$ depending on

whether the operation is addition or multiplication. Both the addition and multiplication algorithms involve two steps: rekeying and rounding.

Rekeying involves switching the key to that used by the next level. To achieve this, the following process is repeated m times, where m depends on the operation along with the key sets used. The process is started by setting $c'_0 := c_0$. Denote by $c'_{j,k}$ the k-th bit in the binary expansion of c'_j. We compute $c'_j := c'_{j-1}t_j + \sum_{k=0}^{\lfloor \log q_i \rfloor} c'_{j-1,k}u_{j,k}$ for some $t_j \in R_{q_i}$ and $u_{j,k} \in R_{q_i}$ for $1 \leq j \leq m$ and $0 \leq k \leq \lfloor \log q_i \rfloor$ which depend on both the operation and the key set used. We need to show that if c'_{j-1} is uniformly distributed, then so is c'_j. Clearly, the term $c'_{j-1}t_j$ is uniformly distributed. If $c'_{j-1,0} = 1$, we observe that $c'_{j-1}t_j + u_{j,0}$ is also uniformly distributed. By induction on k, it can be easily seen that c'_j is uniformly distributed. Let $c' = c'_m$.

Next we need to show that rounding down from c' in R_{q_i} to $R_{q_{i+1}}$ results in an element of the latter that is computationally indistinguishable from a uniform element in $R_{q_{i+1}}$. Rounding is performed by taking the closest integer vector (polynomials in R_{q_i} can be viewed as n-dimensional integer vectors) $\tilde{c}$ to $(q_{i+1}/q_i) \cdot c'$ satisfying $\tilde{c} \equiv c' \pmod 2$. An upper bound on the statistical distance between the distribution of any component of $\tilde{c}$ and the uniform distribution modulo q_{i+1} given that $c' \xleftarrow{\$} R_{q_i}$ is

$$\frac{q_i}{q_{i+1}} = \frac{1}{2^\mu}.$$

Therefore, $\tilde{c}$ is negligibly close to the uniform distribution on $R_{q_{i+1}}$ in the view of an adversary who (1) does not have both c_1 and c_2 and (2) does not know all secret keys corresponding to the public keys under which $\tilde{c}$ is encrypted. The result follows. □

Given a WGH CP-ABE scheme $\mathcal{E}_{\mathsf{ABE}}$ and a multikey FHE scheme $\mathcal{E}^N_{\mathsf{MKFHE}}$, we can construct an EVAL-SIM-secure PBHE scheme $\mathcal{E}^{N-1}_{\mathsf{PBHE2}}$. We will assume that all policies f passed as input to GenKey satisfy $\lambda(f) = 1$. Our PBHE scheme is defined as follows:

- $\mathsf{Setup}(1^\kappa)$: Given a security parameter κ, generate $(\mathsf{PP}, \mathsf{MSK}) \leftarrow \mathsf{ABE.Setup}(1^\kappa)$. Output $(\mathsf{PP}, \mathsf{MSK})$.
- $\mathsf{Extract}(\mathsf{MSK}, a)$: Given a master secret key MSK and an attribute $a \in \mathbb{A}$, output $\mathsf{sk}_a \leftarrow \mathsf{ABE.Extract}(\mathsf{MSK}, a)$.
- $\mathsf{GenKey}(\mathsf{PP}, f)$: Given public parameters PP and an access policy $f \in \mathbb{F}$, run:
  1. On assumption (above) f can be parsed as $\mathfrak{f}$ where $\mathfrak{f} \in \mathfrak{F}$.
  2. $(\mathsf{pk}, \mathsf{sk}, \mathsf{vk}) \leftarrow \mathsf{MKFHE.Gen}(1^\kappa)$.
  3. Compute $\psi \leftarrow \mathsf{ABE.Enc}(\mathsf{PP}, \mathfrak{f}, T(\mathsf{sk}))$.
  4. Set $\mathsf{ek} \leftarrow ((\mathsf{pk}, \mathsf{vk}), \psi)$.
  5. Output $(\mathsf{ek}, \mathsf{vk})$.
- $\mathsf{Enc}(\mathsf{PP}, \mathsf{ek}_f, m)$: Given public parameters PP, an encryption key ek_f for policy f, and a plaintext $m \in \mathbb{M}$, run:
  1. Parse ek_f as $((\mathsf{pk}, \mathsf{vk}), \psi)$.
  2. Compute $c^* \leftarrow \mathsf{MKFHE.Enc}(\mathsf{pk}, m)$.
  3. Output $\mathbf{c} := (c^*, [\psi])$.
- $\mathsf{Dec}(\mathsf{sk}_a, \mathbf{c})$: Given a secret key sk_a for attribute $a \in \mathbb{A}$ and a ciphertext $\mathbf{c}$ that encrypts a message m under access policy f, run:
  1. If $\mathbf{c}$ is of the form $(c^*, [\psi])$:
  (a) If $\mathsf{ABE.Dec}(\mathsf{sk}_a, \psi) = \perp$, output $\perp$ and abort if
  (b) Else set $\mathsf{sk}_1 \leftarrow T^{-1}(\mathsf{ABE.Dec}(\mathsf{sk}_a, \psi))$.
  (c) Set $h \leftarrow 1$.
  2. Else if $\mathbf{c}$ is of the form $(c^*, [\psi_1^{(1)}, \ldots, \psi_{k+1}^{(1)}, \psi_1^{(0)}, \ldots, \psi_k^{(0)}])$:
  (a) If $\mathsf{ABE.Dec}(\mathsf{sk}_a, \psi_{k+1}^{(1)}) = \perp$, output $\perp$ and abort.
  (b) Set $\mathsf{sk}_{k+1} \leftarrow T^{-1}(\mathsf{ABE.Dec}(\mathsf{sk}_a, \psi_{k+1}^{(1)}))$.
  (c) For $1 \leq i \leq k$:
  i. $r_i \leftarrow \mathsf{ABE.Dec}(\mathsf{sk}_a, \psi_i^{(1)})$.
  ii. Set $\mathsf{sk}_i \leftarrow T^{-1}(\mathsf{ABE.Dec}(\mathsf{sk}_a, \psi_i^{(0)}) \cdot r_i^{-1})$.
  (d) Set $h \leftarrow k + 1$.
  3. Output $\mathsf{MKFHE.Dec}(\mathsf{sk}_1, \ldots, \mathsf{sk}_h, c^*)$.
- $\mathsf{Eval}(C, (\mathbf{c_1}, \mathsf{vk}_1), \ldots, (\mathbf{c}_\ell, \mathsf{vk}_\ell))$: Given a circuit $C \in \mathbb{C}$ and a sequence of ℓ pairs of ciphertext and evaluation keys, perform the following steps:
  1. Parse each $\mathbf{c_i}$ as $(c_i^*, [\psi_i])$.
  2. Set $\Psi := \{\psi_i\}_{i \in [\ell]}$ (recall that ciphertexts encrypted under the same vk have the same ψ component). Let $k = |\Psi|$.
  3. Order the elements of Ψ according to their first occurrence in $\psi_1, \ldots, \psi_\ell$ to obtain $\omega_1, \ldots, \omega_k$.
  4. Sample $r_1, \ldots, r_k \xleftarrow{\$} \mathbb{G}$.
  5. Compute $f' \leftarrow f_1 \odot \ldots \odot f_k$.
  6. For $1 \leq i \leq k$:
  (a) Set f_i to the policy associated with ω_i.
  (b) Set $\omega_i^{(0)} \leftarrow \omega_i \boxdot \mathsf{ABE.Enc}(\mathsf{PP}, f_i, r_i)$.
  (c) Set $\omega_i^{(1)} \leftarrow \mathsf{ABE.Enc}(\mathsf{PP}, f', r_i)$.
  7. Generate $(\mathsf{pk}, \mathsf{sk}, \mathsf{vk}) \leftarrow \mathsf{MKFHE.Setup}(1^\lambda)$.
  8. Set $\omega_{k+1}^{(1)} \leftarrow \mathsf{ABE.Enc}(\mathsf{PP}, f', T(\mathsf{sk}))$.
  9. Derive circuit C' that takes $\ell+1$ inputs and XOR's (any binary operation with an identity would suffice) the $(\ell+1)$-th input with the output of C applied to the first ℓ inputs.

10. Set $c_{\ell+1} \leftarrow \mathsf{MKFHE.Enc}(\mathsf{pk}, 0)$.
11. Set $\mathsf{vk}_{\ell+1} \leftarrow \mathsf{vk}$.
12. Set $c^* \leftarrow \mathsf{MKFHE.Eval}(C', (c_1, \mathsf{vk}_1), \ldots, (c_{\ell+1}, \mathsf{vk}_{\ell+1}))$.
13. Output $(c^*, [\omega_1^{(1)}, \ldots, \omega_{k+1}^{(1)}, \omega_1^{(0)}, \ldots, \omega_k^{(0)}])$.

Theorem 7.1. *$\mathcal{E}_{\mathsf{PBHE2}}$ is EVAL-SIM secure.*

Proof. We can show via a hybrid argument that a simulator $\mathcal{S}$ can be constructed such that the real and ideal distributions defined for EVAL-SIM security are indistinguishable in the view of a poly-bounded adversary. In the following series of Hybrids, adaptations are made to the Eval algorithm in the real world to eliminate its dependence on $(\mathbf{c_1}, \mathsf{vk}_1), \ldots, (\mathbf{c}_\ell, \mathsf{vk}_\ell)$, culminating in an algorithm depending solely on PP, C and $f_1, \ldots, f_k$ as required.

Hybrid 0: This is defined as the real distribution with $\mathcal{E}_{\mathsf{PBHE2}}$.

Hybrid 1: The difference between this hybrid and Hybrid 0 is that step 8 of the Eval algorithm is replaced by

$$\omega_{k+1}^{(1)} \leftarrow \mathsf{ABE.Enc}(\mathsf{PP}, f', s)$$

where $s \xleftarrow{\$} \mathbb{G}$. Thus, no information about the secret key sk is provided in the evaluated ciphertext outputted by Eval.

Indistinguishability of the above hybrids follows from the hypothesized IND-AD-CPA security of $\mathcal{E}_{\mathsf{ABE}}$.

Hybrid 2: In this hybrid, we modify step 12 of Eval as follows. Firstly, ℓ key triples for $\mathcal{E}_{\mathsf{MKFHE}}$ are generated i.e. we have $(\mathsf{pk}'_i, \mathsf{sk}'_i, \mathsf{vk}'_i) \leftarrow \mathsf{MKFHE.Setup}(1^\lambda)$ for $1 \le i \le \ell$. Next encryptions of the messages are obtained:

$$c'_i \leftarrow \mathsf{MKFHE.Enc}(\mathsf{pk}'_i, m_i). \tag{7.3}$$

Note that although $m_1, \ldots, m_\ell$ are not inputs to the Eval algorithm, we make them available in this modification to the real distribution; dependence on the messages will be eliminated in subsequent steps. Finally. c^* is computed:

$$c^* \leftarrow \mathsf{MKFHE.Eval}(C', (c'_1, \mathsf{vk}'_1), \ldots, (c'_\ell, \mathsf{vk}'_\ell), (c_{\ell+1}, \mathsf{vk}_{\ell+1})).$$

It follows from the multikey privacy of $\mathcal{E}_{\mathsf{MKFHE}}$ that Hybrid 1 and Hybrid 2 are computationally indistinguishable.

Next we have the following series of ℓ hybrids. For $1 \le i \le \ell$: **Hybrid** $2 + i$: We replace the encryption of m_i in Hybrid 2 (see 7.3) with $c'_i \leftarrow \mathsf{MKFHE.Enc}(\mathsf{pk}'_i, 0)$ i.e. an encryption of zero.

Hybrids $2+i$ and Hybrids $2+(i-1)$ are indistinguishable by virtue of the IND-CPA security of $\mathcal{E}_{\mathsf{MKFHE}}$.

Next we have the following series of k hybrids. For $1 \le i \le k$:

Hybrid $2+\ell+i$: Step 6.(c) of Eval is modified to

$$\omega_i^{(1)} \leftarrow \mathsf{ABE.Enc}(\mathsf{PP}, f', s_i)$$

where $s_i \xleftarrow{\$} \mathbb{G}$.

Since the adversary does not have a secret key for f', indistinguishability of Hybrid $2+\ell+i$ and Hybrid $2+\ell+(i-1)$ follows from the IND-AD-CPA security of $\mathcal{E}_{\mathsf{ABE}}$.

We complete the proof with the following series of k hybrids.
For $1 \le i \le k$:

Hybrid $2+\ell+k+i$: Step 6. (b) of Eval is modified to

$$\omega_i^{(0)} \leftarrow \mathsf{ABE.Enc}(\mathsf{PP}, f_i, z_i) \boxdot \mathsf{ABE.Enc}(\mathsf{PP}, f_i, r_i)$$

where $z_i \xleftarrow{\$} \mathbb{G}$.

The property of WGH, namely 7.1, implies that an adversary cannot efficiently distinguish between Hybrid $2+\ell+k+i$ and Hybrid $2+\ell+k+(i-1)$.

Observe that in Hybrid $2+\ell+2k$, there is no reliance on $\mathsf{vk}_1, \ldots, \mathsf{vk}_\ell$ or $c_1, \ldots, c_\ell$, and in turn on $m_1, \ldots, m_\ell$. In fact, the modified Eval algorithm in Hybrid $2+\ell+2k$ relies only on PP, a description of a circuit C, and k policies $f_1, \ldots, f_k$. Therefore, the modified algorithm can play the role of the simulator $\mathcal{S}$ in the ideal distribution. Since, the ideal distribution is in all other ways identical to the real distribution, we can conclude that our construction is EVAL-SIM secure. □

7.2 Space Efficiency

One of the drawbacks of our construction is the poor space efficiency of evaluated ciphertexts. This can be improved by using a PRF to generate the random elements used to "blind" the multikey FHE secret keys. Instead of encrypting each of these random elements, a short seed for a PRF may be encrypted. This means that the root of a key tree only ever consists of a single CP-ABE ciphertext (assuming that the seed fits within the message space).

7.3 Open Challenges

The IBE and ABE FHE constructions in (Gentry et al., 2013) facilitate evaluation on ciphertexts encrypted under the same identity or attribute. An open problem is to remove this restriction. This is probably more meaningful in the case of ciphertext policies, since it is natural to want to combine or "compose" such policies, usually by means of conjunction. Therefore, an open problem is to construct a CP-ABE scheme

which allows such composition and which satisfies our definition in Section 3 (recall that the compactness condition requires evaluated ciphertexts to depend polynomially on the "size" of the composite policy, and not the circuit).

Another open challenge is to construct a PBHE scheme that has no fixed bound N i.e. which accommodates an unbounded number of independent users. This could be achieved by solving the open problem of constructing a similarly unbounded multikey FHE mentioned in (López-Alt et al., 2012).

8 Conclusions and Future Work

We have initiated the study of homomorphic encryption with support for fine-grained access control and composition. Furthermore, we have proposed a syntax for a primitive that captures the problem of homomorphic encryption in this setting. An instantiation of this primitive was presented that makes use of both CP-ABE and multikey FHE, and shown to be semantically secure. Given that there are currently no known fully-homomorphic (or even somewhat-homomorphic) CP-ABE schemes supporting composition of distinct policies, it seems that our construction is the only way to achieve the same goal, albeit for a bounded number of independent users. In future work, we hope to move beyond the semi-honest model and tackle the problem of verifiability. We also intend to investigate properties such as circuit privacy.

Acknowledgments

We would like to thank the anonymous reviewers of SECRYPT 2013 for their helpful suggestions.

REFERENCES

Armknecht, F., Katzenbeisser, S., and Peter, A. (2010). Group homomorphic encryption: Characterizations, impossibility results, and applications. Cryptology ePrint Archive, Report 2010/501. `http://eprint.iacr.org/`.

Bethencourt, J., Sahai, A., and Waters, B. (2007). Ciphertext-policy attribute-based encryption. In *Proceedings of the 2007 IEEE Symposium on Security and Privacy*, SP '07, pages 321–334, Washington, DC, USA. IEEE Computer Society.

Bonatti, P., De Capitani di Vimercati, S., and Samarati, P. (2002). An algebra for composing access control policies. *ACM Trans. Inf. Syst. Secur.*, 5(1):1–35.

Brakerski, Z., Gentry, C., and Vaikuntanathan, V. (2012). (leveled) fully homomorphic encryption without bootstrapping. In Goldwasser, S., editor, *ITCS*, pages 309–325. ACM.

Bruns, G., Dantas, D. S., and Huth, M. (2007). A simple and expressive semantic framework for policy composition in access control. In *Proceedings of the 2007 ACM workshop on Formal methods in security engineering*, FMSE '07, pages 12–21, New York, NY, USA. ACM.

Clear, M. and McGoldrick, C. (2013). Policy-Based Non-interactive Outsourcing of Computation using multikey FHE and CP-ABE. *Proceedings of the 10th Internation Conference on Security and Cryptography, SECRYPT 2013*.

Gentry, C. (2009). Fully homomorphic encryption using ideal lattices. *Proceedings of the 41st annual ACM Symposium on Theory of Computing STOC 09*, (September):169.

Gentry, C., Sahai, A., and Waters, B. (2013). Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Canetti, R. and Garay, J. A., editors, *CRYPTO (2013)*, volume 8042 of *Lecture Notes in Computer Science*, pages 75–92. Springer.

Hoffstein, J., Pipher, J., and Silverman, J. H. (1998). NTRU: a ring-based public key cryptosystem. *Lecture Notes in Computer Science*, 1423:267–288.

López-Alt, A., Tromer, E., and Vaikuntanathan, V. (2012). On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In *Proceedings of the 44th symposium on Theory of Computing*, STOC '12, pages 1219–1234, New York, NY, USA. ACM.

Moses, T. et al. (2005). Extensible access control markup language (xacml) version 2.0. *Oasis Standard*, 200502.

Ni, Q., Bertino, E., and Lobo, J. (2009). D-algebra for composing access control policy decisions. In *Proceedings of the 4th International Symposium on Information, Computer, and Communications Security*, ASIACCS '09, pages 298–309, New York, NY, USA. ACM.

Pirretti, M., Traynor, P., McDaniel, P., and Waters, B. (2010). Secure attribute-based systems. *Journal of Computer Security*, 18(5):799–837.

Rao, P., Lin, D., Bertino, E., Li, N., and Lobo, J. (2011). Fine-grained integration of access control policies. *Computers & Security*, 30(23):91 – 107. ¡ce:title¿Special Issue on Access Control Methods and Technologies¡/ce:title¿.

Waters, B. (2011). Ciphertext-policy attribute-based encryption: An expressive, efficient, and provably secure realization. In Catalano, D., Fazio, N., Gennaro, R., and Nicolosi, A., editors, *Public Key Cryptography*, volume 6571 of *Lecture Notes in Computer Science*, pages 53–70. Springer.

Xiao, L., Bastani, O., and Yen, I.-L. (2012). An efficient homomorphic encryption protocol for multi-user systems. *IACR Cryptology ePrint Archive*, 2012:193.