Attribute-Based Encryption for Circuits

SERGEY GORBUNOV, University of Waterloo

VINOD VAIKUNTANATHAN, MIT and University of Toronto

HOETECK WEE, CNRS – ENS

In an attribute-based encryption (ABE) scheme, a ciphertext is associated with an ℓ -bit *public index* ind and a message m , and a secret key is associated with a Boolean predicate P . The secret key allows decrypting the ciphertext and learning m if and only if $P(\text{ind}) = 1$. Moreover, the scheme should be secure against collusions of users, namely, given secret keys for polynomially many predicates, an adversary learns nothing about the message if none of the secret keys can individually decrypt the ciphertext.

We present attribute-based encryption schemes for circuits of any arbitrary polynomial size, where the public parameters and the ciphertext grow linearly with the depth of the circuit. Our construction is secure under the standard learning with errors (LWE) assumption. Previous constructions of attribute-based encryption were for Boolean formulas, captured by the complexity class NC^1 .

In the course of our construction, we present a new framework for constructing ABE schemes. As a by-product of our framework, we obtain ABE schemes for polynomial-size branching programs, corresponding to the complexity class $LOGSPACE$, under quantitatively better assumptions.

Categories and Subject Descriptors: E.3 [Data Encryption]: Public key cryptosystems

General Terms: Theory

Additional Key Words and Phrases: Cryptography, fine-grained access control, attribute-based encryption, lattices, learning with errors

ACM Reference Format:

Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. 2015. Attribute-based encryption for circuits. J. ACM 62, 6, Article 45 (December 2015), 33 pages.

DOI: <http://dx.doi.org/10.1145/2824233>

1. INTRODUCTION

Attribute-based encryption [Sahai and Waters 2005; Goyal et al. 2006] is an emerging paradigm for public-key encryption which enables fine-grained control of access to encrypted data. In traditional public-key encryption, access to the encrypted data is all or nothing: given the secret key, one can decrypt and read the entire message, but without it, nothing about the message is revealed (other than its length). In attribute-based encryption, an encryption of a message m is labeled with a *public* attribute vector ind (also called the “index”), and secret keys are associated with predicates P . A secret key sk_P decrypts the ciphertext and recovers the message m if and only if ind satisfies the predicate, namely, if and only if $P(\text{ind}) = 1$.

S. Gorbunov was supported by an Ontario Graduate Scholarship (OGS). V. Vaikuntanathan was supported by an NSERC Discovery Grant and by DARPA under agreement number FA8750-11-2-0225. The views and conclusions contained herein are those of the author and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA or the U.S. Government. H. Wee was supported by NSF CAREER Award CNS-1237429.

Authors’ addresses: S. Gorbunov (corresponding author) and V. Vaikuntanathan, MIT Computer Science and Artificial Intelligence Laboratory, Cambridge, MA; H. Wee, CNRS-DIENS, École Normale Supérieure, Paris, France; corresponding author’s address: sergey@mit.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2015 ACM 0004-5411/2015/12-ART45 \$15.00

DOI: <http://dx.doi.org/10.1145/2824233>

Attribute-based encryption captures as a special case previous cryptographic notions such as identity-based encryption (IBE) [Shamir 1984; Boneh and Franklin 2001; Cocks 2001] and fuzzy IBE [Sahai and Waters 2005]. It has also found applications in scenarios that demand complex policies to control access to encrypted data (such as medical or multimedia data [Akinyele et al. 2011; Li et al. 2013; Papanis et al. 2014]), as well as in designing cryptographic protocols for verifiably outsourcing computations [Parno et al. 2012] (see Section 2.3 for further discussion).

The crucial component in the security requirement for attribute-based encryption stipulates that it resists *collusion attacks*, namely, any group of users collectively learns nothing about the message m if none of them is individually authorized to decrypt the ciphertext.

In the past few years, there has been significant progress in attribute-based encryption in terms of efficiency, security guarantees, and diversifying security assumptions [Goyal et al. 2006; Waters 2009; Lewko and Waters 2010; Lewko et al. 2010; Cash et al. 2012; Agrawal et al. 2010a; Okamoto and Takashima 2010]. On the other hand, little progress has been made in terms of supporting larger classes of predicates. The state of the art is Boolean formulas [Goyal et al. 2006; Lewko et al. 2010; Okamoto and Takashima 2010], which is a subclass of log-space computations. Constructing a secure attribute-based encryption for all polynomial-time predicates was posed as a central challenge by Boneh et al. [2011]. We resolve this problem affirmatively in this work.

2. OUR CONTRIBUTIONS

We construct attribute-based encryption schemes for circuits of every a-priori bounded depth, based on the learning with errors (LWE) assumption. In the course of our construction, we present a new framework for constructing attribute-based encryption schemes, based on a primitive that we call “two-to-one recoding” (TOR). Our methodology departs significantly from the current line of work on attribute-based encryption [Goyal et al. 2006; Lewko et al. 2010] and instead builds upon the connection to garbled circuits developed in the context of *bounded* collusions [Stehl and Steinfeld 2010; Gorbunov et al. 2012]. Along the way, we make the first substantial progress towards the 25-year-old open problem of constructing reusable garbled circuits. In particular, we obtain the first construction of reusable garbled circuits that satisfies authenticity but not privacy. In a follow-up work, Goldwasser et al. [2013] completely resolve this open problem by constructing reusable garbled circuits with authenticity and privacy; moreover, their construction relies crucially on our ABE scheme as an intermediate building block. More details follow.

2.1. Attribute-Based Encryption

For every class of predicate circuits with depth bounded by a polynomial function $d = d(\lambda)$ (where λ is the security parameter), we construct an ABE scheme that supports this class of circuits under the learning with errors (LWE) assumption. Informally, the (decisional) LWE problem [Regev 2009] asks to distinguish between “noisy” random linear combinations of n numbers $\mathbf{s} = (s_1, \dots, s_n) \in \mathbb{Z}_q^n$ from uniformly random numbers over $\mathbb{Z}_q$.

Regev [2009] showed that solving the LWE problem on the average is as hard as (quantumly) solving several notoriously difficult lattice problems in the worst case. Since then, the LWE assumption has become a central fixture in cryptography. We now have a large body of work building cryptographic schemes under the LWE assumption, culminating in the construction of a fully homomorphic encryption scheme [Brakerski and Vaikuntanathan 2011].

The key parameter that determines the hardness of LWE is the ratio between the modulus q and the maximum absolute value of the noise B ; as such, we refer to q/B as the hardness factor of LWE. The problem becomes easier as this ratio grows but is believed to be hard for 2^n -time algorithms when $q/B = 2^{O(n^\epsilon)}$, where $0 < \epsilon < 1/2$. Our results will hold as long as the latter holds for *some constant* ϵ . Similar LWE parameters were used in the context of leakage resilient cryptography and fully-homomorphic encryption [Akavia et al. 2009; Brakerski and Vaikuntanathan 2011].

In particular, our main result constructing selectively secure attribute-based encryption (where the adversary must commit to the challenge public attribute vector ind before seeing the public parameters) can be summarized as follows.

THEOREM 2.1 (INFORMAL). *Assume that there is a constant $0 < \epsilon < 1$ for which the LWE problem is hard for $\text{aexp}(n^\epsilon)$ factor in dimension n , for all large enough n . Then, for any polynomial d , there is a selectively secure attribute encryption scheme for general circuits of depth d .*

Moreover, our scheme has succinct ciphertexts in the sense that the ciphertext size depends polynomially on the depth d and the length ℓ of the attribute vector ind but not on the size of the circuits in the class. The construction as stated achieves the weaker notion of selective security, but we can easily obtain a fully secure scheme following Boneh and Boyen [2004] (but using subexponential hardness in a crucial way).

COROLLARY 2.2. *Assume that there is a constant $0 < \epsilon < 1/2$ such that the LWE problem with a factor of $\text{exp}(n^\epsilon)$ is hard in dimension n for $\text{exp}(n^\epsilon)$ -time algorithms. Then, for any polynomial d , there is a fully secure attribute-based encryption scheme for general circuits of depth d .*

We also obtain a new ABE scheme for branching programs (which correspond to the complexity class LOGSPACE) under the weaker quasi-polynomial hardness of LWE.

THEOREM 2.3 (INFORMAL). *There exist attribute-based encryption schemes for the class of branching programs under either (1) the hardness of the LWE problem with an $n^{\omega(1)}$ factor, or (2) the bilinear decisional Diffie-Hellman assumption.*

Here, there is no a-priori bound on the size or the depth of the branching program. In addition, we achieve succinct ciphertexts of size $O(\ell)$, where ℓ is the number of bits in the index. Prior to this work, under the same $n^{\omega(1)}$ -hardness assumption of LWE that we use, the state-of-the-art constructions were limited to IBE and inner product encryption [Cash et al. 2012; Agrawal et al. 2010a, 2011] (i.e., special cases of branching programs). However, under the same bilinear maps assumption, Goyal et al. [2006] achieved a similar result to ours using secret sharing for general access structures. Our bilinear construction is a different way to achieve the same result but exploits a combinatorial property of branching programs. The construction is inspired by a pairings-based scheme for regular languages [Waters 2012].

We now move on to provide a technical roadmap of our construction: first, we define a new primitive that we call a *two-to-one recoding* (TOR) scheme; we then show how TOR gives us an attribute-based encryption scheme for circuits and how to construct a TOR scheme from the LWE assumption.

2.2. New Framework: TOR

A *two-to-one recoding* (TOR) scheme is a family of (probabilistic) functions $\{\text{Encode}(\text{pk}, \cdot)\}$ indexed by pk , together with a “two-to-one” recoding mechanism. The basic computational security guarantee for $\text{Encode}(\text{pk}, \cdot)$ is that of (correlated)

pseudorandomness [Rosen and Segev 2010]: $\text{Encode}(\text{pk}, s)$ should be pseudorandom given $\text{Encode}(\text{pk}_i, s)$ for polynomially many pk_i 's, where s is a uniformly random “seed.”

The recoding mechanism guarantees that given any triple of public keys $(\text{pk}_0, \text{pk}_1, \text{pk}_{\text{tgt}})$, there is a recoding key rk that allows us to perform the transformation

$$(\text{Encode}(\text{pk}_0, s), \text{Encode}(\text{pk}_1, s)) \mapsto \text{Encode}(\text{pk}_{\text{tgt}}, s).$$

Such a recoding key rk can be generated using either of the two secret keys sk_0 or sk_1 . Furthermore, the recoding mechanism must satisfy a natural simulation requirement: namely, we can generate rk given just pk_0, pk_1 (and neither of the two secret keys), if we are allowed to “program” pk_{tgt} . That is, there are three ways of generating the pair $(\text{pk}_{\text{tgt}}, \text{rk})$ that are (statistically) indistinguishable: (1) given pk_{tgt} , generate rk using the secret key sk_0 ; (2) given pk_{tgt} , generate rk using the secret key sk_1 ; and (3) generate rk without either secret key by “programming” the output public key pk_{tgt} .

This requirement demonstrates the *intuitive guarantee* that we expect from a two-to-one recoding mechanism: namely, the recoding key is “useless” given only one encoding, but not both encodings. For example, it is easy to see that given $\text{Encode}(\text{pk}_0, s)$ and rk (but not $\text{Encode}(\text{pk}_1, s)$), the output $\text{Encode}(\text{pk}_{\text{tgt}}, s)$ is pseudorandom. Indeed, this is because rk could as well have been “simulated” using sk_1 , in which case it is of no help in the distinguishing task.

The simulation requirement also rules out the trivial construction from trapdoor functions, where rk is a trapdoor for inverting $\text{Encode}(\text{pk}_0, \cdot)$ or $\text{Encode}(\text{pk}_1, \cdot)$.

From TOR to Garbled Circuits. We start from the observation that our TOR primitive implies a form of *reusable* garbled circuits with no input or circuit privacy, but instead with a form of authenticity guarantee. As we will see, this leads directly into our attribute-based encryption scheme.

Consider a two-input boolean gate with input wires u, v and output wire w , computing a function $G : \{0, 1\} \times \{0, 1\} \rightarrow \{0, 1\}$. In Yao’s garbled circuit construction, we associate each wire u with a pair of strings (called “labels”) $L_{u,0}, L_{u,1}$. In addition, there is an associated translation table comprising of four values $v_{b,c}$ for $b, c \in \{0, 1\}$, where $v_{b,c}$ allows us to perform the transformation.

$$L_{u,b}, L_{v,c} \mapsto L_{w,G(b,c)}.$$

The garbled circuits construction guarantees that given the translation table and labels L_{u,b^*} and L_{v,c^*} on input wires u, v for specific input bits b^* and c^* , we can obtain $L_{w,G(b^*,c^*)}$ (i.e., the label for wire w for output bit $G(b^*, c^*)$); however, the other label at the output, namely, $L_{w,1-G(b^*,c^*)}$, remains hidden.¹

In our setting, we replace labels with public keys so that each wire is associated with a pair of public keys. As before, we also provide a translation table comprising four values $\text{rk}_{b,c}$, where the recoding key $\text{rk}_{b,c}$ allows us to perform the transformation

$$\text{Encode}(\text{pk}_{u,b}, s), \text{Encode}(\text{pk}_{v,c}, s) \mapsto \text{Encode}(\text{pk}_{w,G(b,c)}, s).$$

The security properties of the TOR scheme then give us the following guarantee: Given the translation table and encodings of s corresponding to b^*, c^* , we clearly compute the encoding of s corresponding to $G(b^*, c^*)$. However, the encoding corresponding to $1 - G(b^*, c^*)$ remains pseudorandom.

¹In the standard instantiation of Yao’s garbled circuits, each label is a secret key of a semantically secure encryption. Moreover, each translation value $v_{b,c}$ is an encryption of key $L_{w,G(b,c)}$ under a pair of keys $L_{u,b}, L_{v,c}$. Given two keys L_{u,b^*}, L_{v,c^*} corresponding to input bits b^*, c^* , it is possible to decrypt and learn key $L_{w,G(b,c)}$ in clear, whereas the key $L_{w,1-G(b,c)}$ remains hidden.

Moreover, crucially, the translation table is independent of s , so we can now “reuse” the translation table by providing fresh encodings with different choices of s . In a sentence, replacing strings by functions gives us the power of reusability.

In the garbled circuits construction, the four entries of the table are permuted and thus one can perform the translation even without knowing what the input bits b^* and c^* are. This is possible because there is an efficient way to verify when the “correct” translation key is being used. In contrast, in the preceding, reusable construction, one has to know exactly which of the recoding keys to use. This is part of the reason why we are unable to provide circuit or input privacy, but instead, only guarantee authenticity, namely, that an adversary can obtain only one of the two possible encodings at the output wire.

This construction forms the cornerstone of the subsequent work of Goldwasser et al. [2013] who construct reusable garbled circuits with input and circuit privacy by additionally leveraging the power of fully homomorphic encryption [Gentry 2009; Brakerski and Vaikuntanathan 2011].

From TOR to Attribute-Based Encryption. How is all this related to attribute-based encryption? In our attribute-based encryption scheme for circuits, the encodings of s are provided in the ciphertext, and the translation tables are provided in the secret key. More precisely, each wire is associated with two TOR public keys, and the encryption of a message m under an index ind is obtained by computing $\text{Encode}(\text{pk}_{i,\text{ind}}, s)$ for every input wire i . The output encoding $\text{Encode}(\text{pk}_{\text{out}}, s)$ is then used to mask the message. We obtain the secret key corresponding to a circuit C by “stitching” multiple translation tables together, where the public keys for the input and output wires are provided in the public parameters, and we pick fresh public keys for the internal wires during key generation. In a nutshell, this gives us the guarantee that given a secret key sk_C and an encryption $\text{Enc}(\text{ind}, m)$ such that $C(\text{ind}) = 1$, we can compute $\text{Encode}(\text{pk}_{\text{out}}, s)$ and thus recover the message. On the other hand, this value looks pseudorandom if $C(\text{ind}) = 0$.

In our outline of reusable garbled circuits with authenticity, we want to reuse the garbled circuit $G(C)$ across multiple encryptions with indices $\text{ind}_1, \text{ind}_2, \dots$ on which C always evaluates to 0. In attribute-based encryption, we also want reusability across multiple circuits $C_1, C_2, \dots$, all of which evaluate to 0 on a fixed index ind (in addition to multiple indices). Fortunately, the strong security properties of the TOR primitive provide us with this guarantee.

To obtain attribute-based encryption for branching programs, we are able to support a different notion of translation tables, which we can realize using a slightly weaker notion of TOR. In branching programs, the transition function depends on an input variable and the current state. The fact that one of these two values is always an input variable makes things simpler; in circuits, both of the input values to a gate could be internal wires.

TOR from LWE. We show how to instantiate TOR from LWE, building upon previous lattice-based IBE techniques [Gentry et al. 2008; Cash et al. 2012; Agrawal et al. 2010a; 2010b]. The public key is given by a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, and

$$\text{Encode}(\mathbf{A}, \mathbf{s}) = \mathbf{A}^T \mathbf{s} + \mathbf{e},$$

where $\mathbf{s} \in \mathbb{Z}_q^n$, $\mathbf{e} \in \mathbb{Z}_q^m$ is an error vector, and $\mathbf{A}^T$ denotes the transpose of the matrix $\mathbf{A}$. (Correlated) pseudorandomness follows directly from the LWE assumption. Given $\mathbf{A}_0, \mathbf{A}_1, \mathbf{A}_{\text{tgt}} \in \mathbb{Z}_q^{n \times m}$, the recoding key rk is given by a low-norm matrix $\mathbf{R} \in \mathbb{Z}_q^{2m \times m}$ such that

$$[\mathbf{A}_0 \parallel \mathbf{A}_1] \mathbf{R} = \mathbf{A}_{\text{tgt}}.$$

Note that

$$\mathbf{R}^T \begin{bmatrix} \mathbf{A}_0^T \mathbf{s} + \mathbf{e}_0 \\ \mathbf{A}_1^T \mathbf{s} + \mathbf{e}_1 \end{bmatrix} \approx \mathbf{A}_{\text{tgt}}^T \mathbf{s},$$

which gives us the recoding mechanism. There are three ways of generating the public key $\mathbf{A}_{\text{tgt}}$ together with the recoding key $\mathbf{R}$: (1) using the trapdoor for $\mathbf{A}_0$, (2) using the trapdoor for $\mathbf{A}_1$, or (3) first generating $\mathbf{R}$ and then “programming” $\mathbf{A}_{\text{tgt}} := [\mathbf{A}_0 || \mathbf{A}_1] \mathbf{R}$. These three ways are statistically indistinguishable by the “bonsai trick” of Cash et al. [2012]. In fact, our recoding mechanism is very similar to the lattice delegation mechanism introduced in Agrawal et al. [2010b], which also uses random low-norm matrices to move from one lattice to another.

The multiplicative mechanism for recoding means that the noise grows exponentially with the number of sequential recodings. This, in turn, limits the depth of the circuits we can handle. In particular, the noise grows by a multiplicative $\text{poly}(n)$ factor on each recoding, which means that after depth d , it becomes $n^{O(d)}$. Since $n^{O(d)} < q/4 < 2^{t^\epsilon}$, we can handle circuits of depth $\tilde{O}(n^\epsilon)$ (here, the first inequality is for correctness and the second for security). Viewed differently, setting the LWE dimension $n = d^{1/\epsilon}$ lets us handle circuits of maximum depth $d = d(\ell)$.

Our weak TOR for branching programs uses an additive mechanism, namely, the recoding key is given by a low-norm matrix $\mathbf{R} \in \mathbb{Z}_q^{m \times m}$ such that $\mathbf{A}_0 \mathbf{R} = \mathbf{A}_{\text{tgt}} - \mathbf{A}_1$. Note that $\mathbf{R}^T (\mathbf{A}_0^T \mathbf{s} + \mathbf{e}_0) + (\mathbf{A}_1^T \mathbf{s} + \mathbf{e}_1) \approx \mathbf{A}_{\text{tgt}}^T \mathbf{s}$, which gives us our recoding mechanism. Since in our branching program construction, $\mathbf{A}_0^T \mathbf{s} + \mathbf{e}_0$ will always be a fresh encoding provided in the ciphertext, the noise accumulation is additive rather than multiplicative.

2.3. Applications

Let us now explain the application of our result to the problem of publicly verifiable delegation of computation without input privacy.

A verifiable delegation scheme allows a computationally weak client to delegate expensive computations to the cloud with the assurance that a malicious cloud cannot convince the client to accept an incorrect computation [Micali 2000; Goldwasser et al. 2008; Gennaro et al. 2010; Chung et al. 2010; Applebaum et al. 2010]. Recent work by Parno et al. [2012] showed that any attribute-based encryption scheme for a class of circuits with encryption time at most linear in the length of the index immediately yields a two-message delegation scheme for the class in the preprocessing model. Namely, there is an initial preprocessing phase which fixes the circuit C the client wishes to compute, produces a circuit key, and sends it to the server. Afterwards, to delegate computation on an input x , the client only needs to send a single message. Moreover, the ensuing delegation scheme satisfies public delegatability, namely, anyone can delegate computations to the cloud as well as public verifiability, namely, anyone can check the cloud’s work (given a “verification” key published by the client). The previous delegation schemes that satisfy both these properties (secure in the standard model) supported the class NC^1 [Parno et al. 2012; Goyal et al. 2006; Lewko and Waters 2012]. Our attribute-based encryption schemes for circuits gives us a verifiable delegation scheme for all circuits, where the computation time of the client in the online phase is polynomial in the length of its input and the depth of the circuit but is otherwise independent of the circuit size. We note that this scheme does not guarantee privacy of the input. Building on this work, Goldwasser et al. [2013] show how to achieve a publicly verifiable delegation scheme with input privacy.

2.4. Related Work

Prior to this work, the state of the art for lattice-based predicate encryption was threshold and inner product predicates [Agrawal et al. 2012, 2011]; realizing Boolean formula was itself an open problem. A different line of work considers definitional issues in the more general realm of functional encryption [Boneh et al. 2011; O’Neill 2010], for which general feasibility results are known for the restricted setting of a-priori bounded collusions developed from classical “one-time” garbled circuits [Sahai and Seyalioglu 2010; Gorbunov et al. 2012] (the ciphertext size grows with both the circuit size and the collusion bound). Our methodology takes a fresh perspective on how to achieve reusability of garbled circuits with respect to authenticity. Our primitive (TOR) can be thought of as a generalization of the notion of proxy re-encryption [Blaze et al. 1998; Ateniese et al. 2006; Hohenberger et al. 2011] which can be thought of as a one-to-one re-encryption mechanism.

Note that in attribute-based encryption, we do not hide the set of attributes. However, if the attributes were hidden, then this would lead to a general solution for keyword-based search over encrypted data [Boneh et al. 2004; Ostrovsky and Skeith 2005].² In particular, we can view the attributes as keywords and secret keys for circuit as the search queries.

Independent Work. Boyen [2013] gave a construction of an ABE scheme for Boolean formulas based on LWE; our result for LWE-based branching program subsumes the result, since Boolean formulas are a subclass of branching programs. Garg et al. [2013] gave a construction of attribute-based encryption for general circuits under a DBDH-like assumption in multilinear groups; the construction extends to so-called graded encodings for which we have candidates under nonstandard assumptions in ideal lattices [Garg et al. 2012; Coron et al. 2013]. The public parameters in the construction also grow with the depth of the circuit.

Subsequent Work. Our attribute-based encryption scheme has been used as the crucial component in the subsequent work of Goldwasser et al. [2013] to construct a (private index) functional encryption scheme with succinct ciphertexts. They also show a number of applications of their construction, including reusable garbled circuits with input and circuit privacy. Also, subsequently, Boneh et al. [2014] gave asymptotic improvements on the sizes of secret keys and ciphertexts in two different constructions, respectively. Their main construction which is built from a new fully key-homomorphic encryption, reduces the size of the secret key for a predicate P from $|P| \times \text{poly}(\lambda, d)$, shown in this work, to $|P| + \text{poly}(\lambda, d)$, where λ is the security parameter and d is the circuit depth. Moreover, their construction supports arithmetic circuits and key delegation.

Organization. We present our TOR framework and its instantiation in Sections 4 and 5. We present our ABE scheme in Section 6. We present the scheme for branching programs in Section 7.

3. PRELIMINARIES

Notation. Let PPT denote probabilistic polynomial time. For any integer $q \geq 2$, we let $\mathbb{Z}_q$ denote the ring of integers modulo q and we represent $\mathbb{Z}_q$ as integers in $(-q/2, q/2]$. We let $\mathbb{Z}_q^{n \times m}$ denote the set of $n \times m$ matrices with entries in $\mathbb{Z}_q$. We use bold capital letters (e.g., $\mathbf{A}$) to denote matrices and bold lowercase letters (e.g., $\mathbf{x}$) to denote vectors. The notation $\mathbf{A}^T$ denotes the transpose of the matrix $\mathbf{A}$.

²This primitive, where the attributes are hidden, is also known as predicate-encryption [Katz et al. 2008].

If $\mathbf{A}_1$ is an $n \times m$ matrix and $\mathbf{A}_2$ is an $n \times m'$ matrix, then $[\mathbf{A}_1 \parallel \mathbf{A}_2]$ denotes the $n \times (m + m')$ matrix formed by concatenating $\mathbf{A}_1$ and $\mathbf{A}_2$. A similar notation applies to vectors. When doing matrix-vector multiplication, we always view vectors as column vectors.

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 $c > 0$, and we use $\text{poly}(n)$ to denote a polynomial function of n . We say an event occurs with *overwhelming probability* if its probability is $1 - \text{negl}(n)$. The function $\lg x$ is the base 2 logarithm of x . The notation $\lceil x \rceil$ denotes the nearest integer to x , rounding towards 0 for half-integers.

3.1. Attribute-Based Encryption

We define attribute-based encryption (ABE), following Goyal et al. [2006]. An ABE scheme ABE for a class of predicate circuits $\mathcal{C}$ (namely, circuits with a single bit output) consists of four PPT algorithms (Setup, Enc, KeyGen, Dec).

- Setup($1^\lambda, 1^\ell$) $\rightarrow$ (pp, mpk, msk). The setup algorithm gets as input the security parameter λ , the length ℓ of the index, and outputs the public parameter (pp, mpk), and the master key msk. All the other algorithms get pp as part of their inputs.
- Enc(mpk, ind, m) $\rightarrow$ ct_{ind} . The encryption algorithm gets as input mpk, an index $\text{ind} \in \{0, 1\}^\ell$, and a message $m \in \mathcal{M}$. It outputs a ciphertext ct_{ind} . Note that ind is public given ct_{ind} .
- KeyGen(msk, C) $\rightarrow$ sk_C . The key generation algorithm gets as input msk and a predicate specified by $C \in \mathcal{C}$. It outputs a secret key sk_C (where C is also public).
- Dec($\text{sk}_C, \text{ct}_{\text{ind}}$) $\rightarrow$ m . The decryption algorithm gets as input sk_C and ct_{ind} and outputs either $\perp$ or a message $m \in \mathcal{M}$.

Correctness. We define and realize perfect correctness of the ABE scheme. Namely, for all (ind, C) such that $C(\text{ind}) = 1$, all $m \in \mathcal{M}$ and $\text{ct}_{\text{ind}} \leftarrow \text{Enc}(\text{mpk}, \text{ind}, m)$, $\text{Dec}(\text{sk}_C, \text{ct}_{\text{ind}}) = m$.

Security Definition. For a stateful adversary $\mathcal{A}$ (that can maintain a global state information throughout the execution of the experiment), we define the advantage function $\text{Adv}_{\mathcal{A}}^{\text{ABE}}(\lambda)$ to be

$$\Pr \left[b = b' : \begin{array}{l} \text{ind} \leftarrow \mathcal{A}(1^\lambda, 1^\ell); \\ (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^\ell); \\ (m_0, m_1) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{mpk}), |m_0| = |m_1|; \\ b \xleftarrow{\$} \{0, 1\}; \\ \text{ct}_{\text{ind}} \leftarrow \text{Enc}(\text{mpk}, \text{ind}, m_b); \\ b' \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{ct}_{\text{ind}}) \end{array} \right] - \frac{1}{2},$$

with the restriction that all queries C that $\mathcal{A}$ makes to $\text{KeyGen}(\text{msk}, \cdot)$ satisfies $C(\text{ind}) = 0$ (i.e., sk_C does not decrypt ct_{ind}). An attribute-based encryption scheme is *selectively secure* if for all PPT adversaries $\mathcal{A}$, the advantage $\text{Adv}_{\mathcal{A}}^{\text{ABE}}(\lambda)$ is a negligible function in λ . We call an attribute-based encryption scheme *fully secure* if the adversary $\mathcal{A}$ is allowed to choose the challenge index ind after seeing secret keys, namely, along with choosing (m_0, m_1) .

3.2. Gaussian Distribution and Learning with Errors (LWE) Assumption

Gaussian Distributions. For any positive parameter $\sigma \in \mathbb{R}_{>0}$, let $\rho_\sigma(\mathbf{x}) := \exp(-\pi \|\mathbf{x}\|^2 / \sigma^2)$ be the Gaussian function on $\mathbb{R}^m$ with parameter σ (and center $\mathbf{0}$). Let $\rho_\sigma(\mathbb{Z}^m) := \sum_{\mathbf{x} \in \mathbb{Z}^m} \rho_\sigma(\mathbf{x})$ be the discrete integral of ρ_σ over $\mathbb{Z}^m$ (which always converges). Let $D_{\mathbb{Z}^m, \sigma}$ be the truncated discrete Gaussian distribution over $\mathbb{Z}^m$ with parameter σ .

Specifically, for all $\mathbf{y} \in \mathbb{Z}^m$, we have $D_{\mathbb{Z}^m, \sigma}(\mathbf{y}) = \frac{\rho_\sigma(\mathbf{y})}{\rho_\sigma(\mathbb{Z}^m)}$. Moreover, when we sample from $D_{\mathbb{Z}^m, \sigma}$ and $\|\cdot\|_\infty$ norm exceeds $\sqrt{m} \cdot \sigma$, we replace the output by $\mathbf{0}$. Note that $D_{\mathbb{Z}^m, \sigma}$ is $\sqrt{m} \cdot \sigma$ -bounded.

The LWE problem was introduced by Regev [2009], who showed that solving it on the average is as hard as solving (quantumly) several standard lattice problems, such as GapSVP and SIVP, in the worst case.

Definition 3.1 (dLWE-Assumption). For an integer $q = q(n) \geq 2$ and a (truncated) discrete Gaussian error distribution $\chi = \chi(n)$ over $\mathbb{Z}_q$, the learning with errors assumption dLWE $_{n,m,q,\chi}$ states that it is computationally hard to distinguish between the following pairs of distributions:

$$\{\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{x}\} \quad \text{and} \quad \{\mathbf{A}, \mathbf{u}\},$$

where $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}$, $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$, $\mathbf{x} \xleftarrow{\$} \chi^m$, $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$.

Connection to Lattices. Let $B = B(n) \in \mathbb{N}$. A family of distributions $\chi = \{\chi_n\}_{n \in \mathbb{N}}$ is called B -bounded if

$$\Pr[\chi \in \{-B, \dots, B-1, B\}] = 1.$$

There are known quantum [Regev 2009] and classical [Peikert 2009] reductions between dLWE $_{n,m,q,\chi}$ and approximating short vector problems in lattices in the worst case, where χ is a B -bounded (truncated) discretized Gaussian for some appropriate B . The state-of-the-art algorithms for these lattice problems run in time nearly exponential in the dimension n [Ajtai et al. 2001; Micciancio and Voulgaris 2010]; more generally, we can get a 2^k -approximation in time $2^{\tilde{O}(n/k)}$. Combined with the connection to LWE, this means that the dLWE $_{n,m,q,\chi}$ assumption is quite plausible for a poly(n)-bounded distribution χ and q as large as 2^{n^ϵ} (for any constant $0 < \epsilon < 1$). Throughout this article, the parameter $m = \text{poly}(n)$, in which case we will shorten the notation slightly to LWE $_{n,q,\chi}$.

3.3. Trapdoors for Lattices

The following lemma summarizes results about lattice trapdoors and algorithms needed for our construction [Ajtai 1999; Gentry et al. 2008; Micciancio and Peikert 2012].

LEMMA 3.2 (LATTICE TRAPDOORS). *There is an efficient randomized algorithm $\text{TrapSamp}(1^n, 1^m, q)$ that, given any integers $n \geq 1$, $q \geq 2$ and sufficiently large $m = \Omega(n \log q)$, outputs a parity check matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and a ‘trapdoor’ matrix $\mathbf{T} \in \mathbb{Z}^{m \times m}$ such that the distribution of $\mathbf{A}$ is $\text{negl}(n)$ -close to uniform.*

Moreover, there is an efficient algorithm SampleD that with overwhelming probability over all random choices, does the following: For any $\mathbf{u} \in \mathbb{Z}_q^n$, and large enough $s = \Omega(\sqrt{n \log q})$, the randomized algorithm $\text{SampleD}(\mathbf{A}, \mathbf{T}, \mathbf{u}, s)$ outputs a vector $\mathbf{r} \in \mathbb{Z}^m$ with norm $\|\mathbf{r}\|_\infty \leq \|\mathbf{r}\|_2 \leq s\sqrt{n}$ (with probability 1). Furthermore, the following distributions of the tuple $(\mathbf{A}, \mathbf{T}, \mathbf{U}, \mathbf{R})$ are within $\text{negl}(n)$ statistical distance of each other for any polynomial $k \in \mathbb{N}$.

— $(\mathbf{A}, \mathbf{T}) \leftarrow \text{TrapSamp}(1^n, 1^m, q)$; $\mathbf{U} \leftarrow \mathbb{Z}_q^{n \times k}$; $\mathbf{R} \leftarrow \text{SampleD}(\mathbf{A}, \mathbf{T}, \mathbf{U}, s)$.

— $(\mathbf{A}, \mathbf{T}) \leftarrow \text{TrapSamp}(1^n, 1^m, q)$; $\mathbf{R} \leftarrow (D_{\mathbb{Z}^m, s})^k$; $\mathbf{U} := \mathbf{A}\mathbf{R} \pmod{q}$.

4. TWO-TO-ONE RECODING SCHEMES

An overview of TOR is provided in Section 2.2. We first informally point out some properties of the encryption/decryption algorithms that are a part of TOR. Then, in Section 4.1, we provide a formal definition of all TOR algorithms and properties.

Symmetric Encryption. In our construction, we will use $\text{Encode}(\text{pk}, s)$ as a one-time key for a symmetric-key encryption scheme (E, D) . If Encode is deterministic, then we could simply use a one-time pad. However, since Encode is probabilistic, the one-time pad will not guarantee correctness. Instead, we require (E, D) to satisfy a stronger correctness guarantee, namely, for all messages m and for all ψ, ψ' in the support of $\text{Encode}(\text{pk}, s)$, $D(\psi', E(\psi, m)) = m$.

Allowing Degradation. With each recoding operation, the “quality” of encoding potentially degrades. In order to formalize this, we allow the initial global public parameters to depend on $d_{\max}$, an a-prior upper bound on the number of nested recoding operations. We then require that given any encodings ψ and ψ' that are a result of at most $d_{\max}$ nested recodings, $D(\psi', E(\psi, m)) = m$. We stress that we allow $d_{\max}$ to be superpolynomial, and in fact, provide such instantiations for a relaxed notion of TOR.

4.1. Definition of TOR

Formally, a TOR scheme over the input space $\mathcal{S} = \{S_\lambda\}$ and output space $\mathcal{K} = \{\mathcal{K}_\lambda\}$ consists of six polynomial-time algorithms (Params , Keygen , Encode , ReKeyGen , SimReKeyGen , Recode) and a symmetric-key encryption scheme (E, D) with the following properties.

- $\text{Params}(1^\lambda, d_{\max})$ is a probabilistic algorithm that takes as input the security parameter λ and an upper bound $d_{\max}$ on the number of nested recoding operations (written in binary), and then outputs “global” public parameters pp .
- $\text{Keygen}(\text{pp})$ is a probabilistic algorithm that outputs a public/secret key pair (pk, sk) .
- $\text{Encode}(\text{pk}, s)$ is a probabilistic algorithm that takes pk and an input $s \in \mathcal{S}$ and, then outputs an encoding ψ .

In addition, there is a recoding mechanism together with two ways to generate recoding keys: given one of the two secret keys or by programming the output public key.

- $\text{ReKeyGen}(\text{pk}_0, \text{pk}_1, \text{sk}_0, \text{pk}_{\text{tgt}})$ is a probabilistic algorithm that takes a key pair $(\text{pk}_0, \text{sk}_0)$, another public key pk_1 , a “target” public key pk_{tgt} , and outputs a recoding key rk .
- $\text{SimReKeyGen}(\text{pk}_0, \text{pk}_1)$ is a probabilistic algorithm that takes two public keys pk_0, pk_1 and outputs a recoding key rk together with a “target” public key pk_{tgt} .
- $\text{Recode}(\text{rk}, \psi_0, \psi_1)$ is a deterministic algorithm that takes the recoding key rk , two encodings ψ_0 and ψ_1 , and outputs an encoding ψ_{tgt} .

Remark 4.1. For our instantiation from lattices, we can in fact invert $\text{Encode}(\text{pk}, s)$ to recover s using the corresponding sk . However, we will not require this property in our generic constructions from TOR. Indeed, realizing this property over bilinear groups would be hard, since s is typically encoded in the exponent.

Correctness. Correctness of a TOR scheme requires two things. First, for sufficiently large λ , for every pk and $s \in \mathcal{S}$, there exists a family of sets $\Psi_{\text{pk}, s, j}$, $j = 0, 1, \dots, d_{\max}$.

- $\Pr[\text{Encode}(\text{pk}, s) \in \Psi_{\text{pk}, s, 0}] = 1$, where the probability is taken over the coin tosses of Encode ;
- $\Psi_{\text{pk}, s, 0} \subseteq \Psi_{\text{pk}, s, 1} \subseteq \dots \subseteq \Psi_{\text{pk}, s, d_{\max}}$.
- For all $\psi, \psi' \in \Psi_{\text{pk}, s, d_{\max}}$ and all $m \in \mathcal{M}$, $D(\psi', E(\psi, m)) = m$.

Note that these properties hold trivially if Encode is deterministic and (E, D) is the one-time pad. Second, the correctness of recoding requires that for any triple of key pairs $(pk_0, sk_0), (pk_1, sk_1), (pk_{tgt}, sk_{tgt})$ and any encodings $\psi_0 \in \Psi_{pk_0, s, j_0}$ and $\psi_1 \in \Psi_{pk_1, s, j_1}$,

$$\text{Recode}(rk, \psi_0, \psi_1) \in \Psi_{pk_{tgt}, s, \max(j_0, j_1)+1}.$$

Statistical Security Properties. Note that we have three ways of sampling recoding keys: using ReKeyGen along with one of two secret keys sk_0 or sk_1 ; using SimReKeyGen while programming pk_{tgt} . We require that for all λ , all three ways lead to the same distribution of recoding keys, up to some statistical error.

(Key Indistinguishability). Let $(pk_b, sk_b) \leftarrow \text{Keygen}(pp)$ for $b = 0, 1$ and $(pk_{tgt}, sk_{tgt}) \leftarrow \text{Keygen}(pp)$. Then, for all λ , the following two ensembles must be statistically indistinguishable:

$$\begin{aligned} & \left[\text{Aux}, \text{ReKeyGen}(pk_0, pk_1, \boxed{sk_0}, pk_{tgt}) \right] \stackrel{s}{\approx}, \\ & \left[\text{Aux}, \text{ReKeyGen}(pk_1, pk_0, \boxed{sk_1}, pk_{tgt}) \right], \end{aligned}$$

where $\text{Aux} = ((pk_0, sk_0), (pk_1, sk_1), (pk_{tgt}, sk_{tgt}))$. Informally, this says that sampling recoding keys using sk_0 or sk_1 yields the same distribution.

(Recoding Simulation). Let $(pk_b, sk_b) \leftarrow \text{Keygen}(pp)$ for $b = 0, 1$. Then, for all λ , the following two ways of sampling the tuple $[(pk_0, sk_0), (pk_1, sk_1), pk_{tgt}, rk]$ must be statistically indistinguishable:

$$\begin{aligned} & [(pk_0, sk_0), (pk_1, sk_1), pk_{tgt}, rk : (pk_{tgt}, sk_{tgt}) \leftarrow \text{Keygen}(pp); rk \\ & \quad \leftarrow \text{ReKeyGen}(pk_0, pk_1, sk_0, pk_{tgt})] \stackrel{s}{\approx}, \\ & [(pk_0, sk_0), (pk_1, sk_1), pk_{tgt}, rk : (pk_{tgt}, rk) \leftarrow \text{SimReKeyGen}(pk_0, pk_1)]. \end{aligned}$$

In addition, we require one-time semantic security for (E, D) .

(One-Time Semantic Security). For all $m_0, m_1 \in \mathcal{M}$ and all λ , the following two distributions must be statistically indistinguishable:

$$\left[E(\psi, m_0) : \psi \xleftarrow{\$} \mathcal{K} \right] \stackrel{s}{\approx} \left[E(\psi, m_1) : \psi \xleftarrow{\$} \mathcal{K} \right].$$

For all three properties, computational indistinguishability is sufficient for our applications, but we will achieve the stronger statistical indistinguishability in our instantiations.

Computational Security Property. We require that for all λ , given the encoding of a random s on $\ell = \text{poly}(\lambda)$ keys, the evaluation at a fresh key is pseudorandom.

(Correlated Pseudorandomness). For every polynomial $\ell = \ell(\lambda)$, let $(pk_i, sk_i) \leftarrow \text{Keygen}(pp)$ for $i \in [\ell + 1]$. Let $s \xleftarrow{\$} \mathcal{S}$, and let $\psi_i \leftarrow \text{Encode}(pk_i, s)$ for $i \in [\ell + 1]$. Then, the following two ensembles must be computationally indistinguishable:

$$\begin{aligned} & \left[(pk_i, \psi_i)_{i \in [\ell]}, pk_{\ell+1}, \boxed{\psi_{\ell+1}} \right] \stackrel{c}{\approx}, \\ & \left[(pk_i, \psi_i)_{i \in [\ell]}, pk_{\ell+1}, \boxed{\psi} : \psi \xleftarrow{\$} \mathcal{K} \right]. \end{aligned}$$

That is, we define the advantage function $\text{Adv}_{\mathcal{A}}^{\text{CP}}(\lambda)$ to be

$$\Pr \left[\begin{array}{l} \text{pp} \leftarrow \text{Setup}(1^\lambda); s \leftarrow \mathcal{S}; \\ (\text{pk}_i, \text{sk}_i) \leftarrow \text{Keygen}(\text{pp}), \\ \psi_i \leftarrow \text{Encode}(\text{pk}_i, s), i = 1, \dots, \ell; \\ \psi'_0 \leftarrow \text{Encode}(\text{pk}_{\ell+1}, s); \\ b \xleftarrow{\$} \{0, 1\}; \psi'_1 \xleftarrow{\$} \mathcal{K} \\ b' \leftarrow \mathcal{A}(\text{pk}_1, \dots, \text{pk}_{\ell+1}, \psi_1, \dots, \psi_\ell, \psi'_b) \end{array} \right] - \frac{1}{2},$$

and we require that for all PPT $\mathcal{A}$, the advantage function $\text{Adv}_{\mathcal{A}}^{\text{CP}}(\lambda)$ is a negligible function in λ .

4.2. Simple Applications of TOR

In this section, we provide high-level sketches of how to use TOR to construct an analogue of Yao's garbled circuits (without privacy) and identity-based encryption (IBE).

First Example. We revisit the example from Section 2.2. Consider a two-input boolean gate g with input wires u, v and output wire w , computing a function $G : \{0, 1\} \times \{0, 1\} \rightarrow \{0, 1\}$. Analogous to Yao's garbled circuit, we provide a translation table Γ comprising four values

$$\Gamma := (\text{rk}_{b,c} : b, c \in \{0, 1\}),$$

where $\text{rk}_{b,c}$ allows us to perform the transformation.

$$\text{Encode}(\text{pk}_{u,b}, s), \text{Encode}(\text{pk}_{v,c}, s) \mapsto \text{Encode}(\text{pk}_{w,G(b,c)}, s).$$

Now, fix b^*, c^* and $d^* := G(b^*, c^*)$. Given an encoding of s corresponding to b^* and c^* , we can compute that for d^* using the recoding key rk_{b^*,c^*} ; in addition, we claim that the encoding corresponding to $1 - d^*$ remains pseudorandom. To prove this, it suffices to simulate Γ given $\text{pk}_{u,b^*}, \text{pk}_{v,c^*}, \text{pk}_{w,1-d^*}$ as follows:

- we sample $(\text{pk}_{w,d^*}, \text{rk}_{b^*,c^*})$ using SimReKeyGen ;
- we sample $\text{pk}_{u,1-b^*}$ and $\text{pk}_{v,1-c^*}$ along with the corresponding secret keys; using these secret keys, we can sample the other three recoding keys $\text{rk}_{1-b^*,c^*}, \text{rk}_{b^*,1-c^*}, \text{rk}_{1-b^*,1-c^*}$.

IBE from TOR. As a warm-up, we show how to build a selectively secure IBE for identity space $\{0, 1\}^\ell$.

$$\text{mpk} := \begin{pmatrix} \text{pk}_{1,0} & \text{pk}_{2,0} & \dots & \text{pk}_{\ell,0} & \text{pk}_{\text{start}} \\ \text{pk}_{1,1} & \text{pk}_{2,1} & \dots & \text{pk}_{\ell,1} & \text{pk}_{\text{out}} \end{pmatrix}.$$

The ciphertext for identity ind and message m is given by.

$$(\text{Encode}(\text{pk}_{1,\text{ind}_1}, s), \dots, \text{Encode}(\text{pk}_{\ell,\text{ind}_\ell}, s), \text{Encode}(\text{pk}_{\text{start}}, s), \text{E}(\text{Encode}(\text{pk}_{\text{out}}, s), m)).$$

The secret key for identity ind is given by $(\text{rk}_1, \dots, \text{rk}_\ell)$, where we first sample

$$(\text{pk}'_1, \text{sk}'_1), \dots, (\text{pk}'_{\ell-1}, \text{sk}'_{\ell-1}) \leftarrow \text{Keygen}(\text{pp}),$$

and then sample

$$\begin{aligned} \text{rk}_1 &\leftarrow \text{ReKeyGen}(\text{pk}_{\text{start}}, \text{pk}_{1,\text{ind}_1}, \text{sk}_{\text{start}}, \text{pk}'_1) \\ \text{rk}_2 &\leftarrow \text{ReKeyGen}(\text{pk}'_1, \text{pk}_{2,\text{ind}_2}, \text{sk}'_1, \text{pk}'_2) \\ &\vdots \\ \text{rk}_\ell &\leftarrow \text{ReKeyGen}(\text{pk}'_{\ell-1}, \text{pk}_{\ell,\text{ind}_\ell}, \text{sk}'_{\ell-1}, \text{pk}_{\text{out}}). \end{aligned}$$

To prove selective security, we need to generate secret keys for any $\text{ind} \neq \text{ind}^*$, given $\text{sk}_{1,1-\text{ind}_1^*}, \dots, \text{sk}_{\ell,1-\text{ind}_\ell^*}$ but not sk_{start} or sk_{out} . We can achieve this as follows. Pick an i for which $\text{ind}_i \neq \text{ind}_i^*$:

- pick $(\text{rk}_1, \text{pk}'_1), \dots, (\text{rk}_{i-1}, \text{pk}'_{i-1})$ using SimReKeyGen ;
- pick $(\text{pk}'_i, \text{sk}'_i), \dots, (\text{pk}'_{\ell-1}, \text{sk}'_{\ell-1})$ using Keygen ;
- pick $\text{rk}_i, \text{rk}_{i+1}, \dots, \text{rk}_\ell$ using ReKeyGen with secret keys $\text{sk}_{1-\text{ind}_1^*}, \text{sk}'_i, \dots, \text{sk}'_{\ell-1}$, respectively.

We note that our IBE construction from TOR seems incomparable to existing IBE constructions from lattices [Cash et al. 2012; Agrawal et al. 2010a], as we follow more of a “sequential” binding to the identity ind .

5. TOR FROM LWE

In this section, we present an instantiation of TOR from LWE, building upon ideas previously introduced in [Gentry et al. 2008; Cash et al. 2012; Agrawal et al. 2010a, 2010b].

LEMMA 5.1. *Assuming $\text{dLWE}_{n,q,\chi}$ there is a TOR scheme that is correct up to $d_{\max}$ levels, where $n = n(\lambda)$, $\chi = D_{\mathbb{Z}, \sqrt{n}}$, $q = n^{\Theta(d_{\max})}$, $m = O(n \log q)$, and, $q = n^{\Theta(d_{\max})}$.*

- $\text{Params}(1^\lambda, d_{\max})$. The parameters n, m, χ, q , are given from the assumption. Set the error bound $B = B(n) = O(n)$ and the Gaussian parameter $s = s(n) = O(\sqrt{n \log q})$. Output the global public parameters $\text{pp} = (n, \chi, B, q, m, s)$.
- $\text{Keygen}(\text{pp})$. Run the trapdoor generation algorithm $\text{TrapGen}(1^n, 1^m, q)$ to obtain a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ together with the trapdoor matrix $\mathbf{T} \in \mathbb{Z}^{m \times m}$. Output $\text{pk} := \mathbf{A}$ and $\text{sk} := \mathbf{T}$.
- $\text{Encode}(\text{pk}, \mathbf{s})$. Sample an error vector $\mathbf{e} \leftarrow \chi^m$ and output the encoding $\psi := \mathbf{A}^T \mathbf{s} + \mathbf{e} \in \mathbb{Z}_q^m$.

The decoding algorithms work as follows.

- $\text{ReKeyGen}(\text{pk}_0, \text{pk}_1, \text{sk}_b, \text{pk}_{\text{tgt}})$. Let $\text{pk}_0 = \mathbf{A}_0$, $\text{pk}_1 = \mathbf{A}_1$, $\text{sk}_b = \mathbf{T}_b$, and $\text{pk}_{\text{tgt}} = \mathbf{A}_{\text{tgt}}$. Compute the matrix $\mathbf{R} \in \mathbb{Z}^{2m \times m}$ in the following way.
 - (i) Choose a discrete Gaussian matrix $\mathbf{R}_{1-b} \leftarrow (D_{\mathbb{Z},s})^{m \times m}$. Namely, each entry of the matrix is an independent sample from the discrete Gaussian distribution $D_{\mathbb{Z},s}$.
 - (ii) Compute $\mathbf{U} := \mathbf{A}_{\text{tgt}} - \mathbf{A}_{1-b} \mathbf{R}_{1-b} \in \mathbb{Z}_q^{n \times m}$.
 - (iii) Compute the matrix $\mathbf{R}_b$ by running the algorithm SampleD to compute a matrix $\mathbf{R}_b \in \mathbb{Z}^{m \times m}$ as follows:

$$\mathbf{R}_b \leftarrow \text{SampleD}(\mathbf{A}_b, \mathbf{T}_b, \mathbf{U}).$$

Output

$$\text{rk}_{0,1}^{\text{tgt}} := \begin{bmatrix} \mathbf{R}_0 \\ \mathbf{R}_1 \end{bmatrix} \in \mathbb{Z}^{2m \times m}.$$

(We remark that $\mathbf{A}_b \mathbf{R}_b = \mathbf{U} = \mathbf{A}_{\text{tgt}} - \mathbf{A}_{1-b} \mathbf{R}_{1-b}$, and thus, $\mathbf{A}_0 \mathbf{R}_0 + \mathbf{A}_1 \mathbf{R}_1 = \mathbf{A}_{\text{tgt}}$).

- $\text{SimReKeyGen}(\text{pk}_0, \text{pk}_1)$. Let $\text{pk}_0 = \mathbf{A}_0$ and $\text{pk}_1 = \mathbf{A}_1$.
 - (i) Sample a matrix $\mathbf{R} \leftarrow (D_{\mathbb{Z},s})^{2m \times m}$ by sampling each entry from the discrete Gaussian distribution $D_{\mathbb{Z},s}$.
 - (ii) Define

$$\mathbf{A}_{\text{tgt}} := [\mathbf{A}_0 \parallel \mathbf{A}_1] \mathbf{R} \in \mathbb{Z}_q^{n \times m}.$$

Output the pair $(\text{pk}_{\text{tgt}} := \mathbf{A}_{\text{tgt}}, \text{rk}_{0,1}^{\text{tgt}} := \mathbf{R})$.

—Recode($\text{rk}_{0,1}^{\text{tgt}}, \psi_0, \psi_1$). Let $\text{rk}_{0,1}^{\text{tgt}} = \mathbf{R}$. Compute the recoded ciphertext

$$\psi_{\text{tgt}} = [\psi_0^T \parallel \psi_1^T] \mathbf{R}.$$

We also need a one-time symmetric encryption scheme $(\mathbf{E}, \mathbf{D})$, which we will instantiate as an *error-tolerant* version of the one-time pad with $\mathcal{K} = \mathbb{Z}_q^m$, $\mathcal{M} = \{0, 1\}^m$, as follows.

— $\mathbf{E}(\psi, \mu)$ takes as input a vector $\psi \in \mathbb{Z}_q^m$ and a bit string $\mu \in \mathcal{M}^m$ and outputs the encryption

$$\gamma := \psi + \lceil q/2 \rceil \mu \pmod{q}.$$

— $\mathbf{D}(\psi', \gamma)$ takes as input a vector $\psi' = (\psi'_1, \dots, \psi'_m) \in \mathbb{Z}_q^m$, an encryption $\gamma = (\gamma_1, \dots, \gamma_n) \in \mathbb{Z}_q^m$, and does the following. Define a function $\text{Round}(x)$, where $x \in [-(q-1)/2, \dots, (q-1)/2]$ as

$$\text{Round}(x) = \begin{cases} 0 & \text{if } |x| < q/4, \\ 1 & \text{otherwise.} \end{cases}$$

The decryption algorithm outputs a vector $\mu = (\text{Round}(\gamma_1 - \psi'_1), \dots, \text{Round}(\gamma_n - \psi'_n))$.

The scheme is information-theoretically secure for a one-time use. As for correctness, we will show that for every two ψ and ψ' that are “close”, $\mathbf{D}(\psi', \mathbf{E}(\psi, \mu)) = \mu$. More precisely, call ψ and ψ' close if for every $i \in [n]$, $|\psi_i - \psi'_i| < q/4$.

$$\begin{aligned} \mathbf{D}(\psi', \mathbf{E}(\psi, \mu)) &= [\text{Round}(\gamma_1 - \psi'_1), \dots, \text{Round}(\gamma_n - \psi'_n)] \\ &= [\text{Round}(\mu_1 \lceil q/2 \rceil + \psi_1 - \psi'_1), \dots, \text{Round}(\mu_n \lceil q/2 \rceil + \psi_n - \psi'_n)]. \end{aligned}$$

Observe that if $|\psi_i - \psi'_i| < q/4$, then $\text{Round}(\mu_i \lceil q/2 \rceil + \psi_i - \psi'_i) = \mu_i$. This completes the proof of correctness.

5.1. Analysis

Correctness. We define the sets $\Psi_{\mathbf{A}, \mathbf{s}, j}$ for $\text{pk} := \mathbf{A} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{s} \in \mathbb{Z}_q^n$, and $j \in [1 \dots d_{\max}]$ as follows:

$$\Psi_{\mathbf{A}, \mathbf{s}, j} = \{\mathbf{A}^T \mathbf{s} + \mathbf{e} : \|\mathbf{e}\|_\infty \leq B \cdot (2sm\sqrt{m})^j\}.$$

Given this definition, observe the following.

- When $\mathbf{e} \leftarrow \chi^m$, $\|\mathbf{e}\|_\infty \leq B$ by the definition of χ and B . $\Pr[\text{Encode}(\mathbf{A}, \mathbf{s}) \in \Psi_{\mathbf{A}, \mathbf{s}, 0}] = 1$.
- $\Psi_{\mathbf{A}, \mathbf{s}, 0} \subseteq \Psi_{\mathbf{A}, \mathbf{s}, 1} \subseteq \dots \subseteq \Psi_{\mathbf{A}, \mathbf{s}, d_{\max}}$, by definition of the preceding sets.
- For any two encodings $\psi = \mathbf{A}^T \mathbf{s} + \mathbf{e}$, $\psi' = \mathbf{A}^T \mathbf{s} + \mathbf{e}' \in \Psi_{\mathbf{A}, \mathbf{s}, d_{\max}}$,

$$\|\psi - \psi'\|_\infty = \|\mathbf{e} - \mathbf{e}'\|_\infty \leq 2 \cdot B \cdot (2sm\sqrt{m})^{d_{\max}} < q/4,$$

which holds as long as $n \cdot O(n^2 \log q)^{d_{\max}} < q/4$. Thus, ψ and ψ' are “close”, and by the correctness property of the symmetric encryption scheme $(\mathbf{E}, \mathbf{D})$ previously described, $\mathbf{D}(\psi', \mathbf{E}(\psi, \mu)) = \mu$ for any $\mu \in \{0, 1\}^n$.

- Consider two encodings $\psi_0 \in \Psi_{\mathbf{A}_0, \mathbf{s}, j_0}$ and $\psi_1 \in \Psi_{\mathbf{A}_1, \mathbf{s}, j_1}$ for any $j_0, j_1 \in \mathbb{N}$, any $\mathbf{A}_0, \mathbf{A}_1 \in \mathbb{Z}_q^{n \times m}$, and $\mathbf{s} \in \mathbb{Z}_q^n$. Then, $\psi_0 = \mathbf{A}_0^T \mathbf{s} + \mathbf{e}_0$ and $\psi_1 = \mathbf{A}_1^T \mathbf{s} + \mathbf{e}_1$, where $\|\mathbf{e}_0\|_\infty \leq B \cdot (2sm\sqrt{m})^{j_0}$ and $\|\mathbf{e}_1\|_\infty \leq B \cdot (2sm\sqrt{m})^{j_1}$.

Then, the recoded ciphertext ψ_{tgt} is computed as follows.

$$\begin{aligned}\psi_{\text{tgt}}^T &:= [\psi_0^T \parallel \psi_1^T] \mathbf{R}_{0,1}^{\text{tgt}} \\ &= [\mathbf{s}^T \mathbf{A}_0 + \mathbf{e}_0^T \parallel \mathbf{s}^T \mathbf{A}_1 + \mathbf{e}_1^T] \mathbf{R}_{0,1}^{\text{tgt}} \\ &= \mathbf{s}^T [\mathbf{A}_0 \parallel \mathbf{A}_1] \mathbf{R}_{0,1}^{\text{tgt}} + [\mathbf{e}_0^T \parallel \mathbf{e}_1^T] \mathbf{R}_{0,1}^{\text{tgt}} \\ &= \mathbf{s}^T \mathbf{A}_{\text{tgt}} + \mathbf{e}_{\text{tgt}}^T,\end{aligned}$$

where the last equation is because $\mathbf{A}_{\text{tgt}} = [\mathbf{A}_0 \parallel \mathbf{A}_1] \mathbf{R}_{0,1}^{\text{tgt}}$, and we define $\mathbf{e}_{\text{tgt}} := [\mathbf{e}_0^T \parallel \mathbf{e}_1^T] \mathbf{R}_{0,1}^{\text{tgt}}$. Thus,

$$\begin{aligned}\|\mathbf{e}_{\text{tgt}}\|_\infty &\leq m \cdot \|\mathbf{R}_{0,1}^{\text{tgt}}\|_\infty \cdot (\|\mathbf{e}_0\|_\infty + \|\mathbf{e}_1\|_\infty) \\ &\leq m \cdot s\sqrt{m} \cdot (B \cdot (2sm\sqrt{m})^{j_0} + B \cdot (2sm\sqrt{m})^{j_1}) \\ &\leq B \cdot (2sm\sqrt{m})^{\max(j_0, j_1)+1},\end{aligned}$$

exactly as required. Here, the second inequality is because $\|\mathbf{R}_{0,1}^{\text{tgt}}\|_\infty \leq s\sqrt{m}$ by Lemma 3.2. This finishes our proof of correctness.

Key Indistinguishability. Recall that in ReKeyGen, we are given samplings $(\mathbf{R}_0, \mathbf{R}_1)$ satisfying $\mathbf{A}_0 \mathbf{R}_0 + \mathbf{A}_1 \mathbf{R}_1 = \mathbf{A}_{\text{tgt}}$. In key indistinguishability, we must argue that the distributions produced by sampling using a trapdoor for $\mathbf{A}_0$ or that for $\mathbf{A}_1$ are indistinguishable. Indeed, this follows directly from the following statement in Cash et al. [2012] and Gentry et al. [2008] (see also [Cash et al. 2009, Theorem 3.4]): for every $(\mathbf{A}_0, \mathbf{T}_0), (\mathbf{A}_1, \mathbf{T}_1)$ generated by $\text{TrapSamp}(1^n, 1^m, q)$, every matrix $\mathbf{A}_{\text{tgt}} \in \mathbb{Z}_q^{n \times m}$, and any $s = \Omega(\sqrt{n \log q})$, the following two experiments generate distributions with $\text{negl}(n)$ statistical distance.

- Sample $\mathbf{R}_0 \leftarrow (D_{\mathbb{Z}^m, s})^m$, compute $\mathbf{U} := \mathbf{A}_{\text{tgt}} - \mathbf{A}_0 \mathbf{R}_0 \in \mathbb{Z}_q^{n \times m}$ and $\mathbf{R}_1 \leftarrow \text{SampleD}(\mathbf{A}_1, \mathbf{T}_1, \mathbf{U}, s)$. Output $(\mathbf{A}_0, \mathbf{T}_0, \mathbf{A}_1, \mathbf{T}_1, \mathbf{A}_{\text{tgt}}, \mathbf{R}_0, \mathbf{R}_1)$.
- Sample $\mathbf{R}_1 \leftarrow (D_{\mathbb{Z}^m, s})^m$, compute $\mathbf{U} := \mathbf{A}_{\text{tgt}} - \mathbf{A}_1 \mathbf{R}_1 \in \mathbb{Z}_q^{n \times m}$ and $\mathbf{R}_0 \leftarrow \text{SampleD}(\mathbf{A}_0, \mathbf{T}_0, \mathbf{U}, s)$. Output $(\mathbf{A}_0, \mathbf{T}_0, \mathbf{A}_1, \mathbf{T}_1, \mathbf{A}_{\text{tgt}}, \mathbf{R}_0, \mathbf{R}_1)$.

Note that in both distributions, it clearly holds that $\mathbf{A}_0 \mathbf{R}_0 + \mathbf{A}_1 \mathbf{R}_1 = \mathbf{A}_{\text{tgt}}$, as required.

The recoding simulation property follows readily from Lemma 3.2, as is done in Cash et al. [2012]. In particular, given $[(\mathbf{A}_0, \mathbf{T}_0), (\mathbf{A}_1, \mathbf{T}_1), \mathbf{A}_{\text{tgt}}, (\mathbf{R}_0, \mathbf{R}_1)]$, from Lemma 3.2, it follows that if we are given $\mathbf{A}_{\text{tgt}}$, then we may sample $(\mathbf{R}_0, \mathbf{R}_1)$ using the trapdoor $\mathbf{T}_0$. Alternatively, sampling $(\mathbf{R}_0, \mathbf{R}_1)$ first and setting $\mathbf{A}_{\text{tgt}} = \mathbf{A}_0 \mathbf{R}_0 + \mathbf{A}_1 \mathbf{R}_1$ produces the same distribution.

Correlated pseudorandomness directly from the decisional LWE assumption $\text{dLWE}_{n, (\ell+1)m, q, \chi}$, where $q = n^{\Theta(d_{\max})}$. In particular, by the dLWE, given $\mathbf{A}_1, \dots, \mathbf{A}_{\ell+1}, \psi_1, \dots, \psi_\ell, \psi^*$, no adversary can distinguish between the cases when ψ^* is a valid LWE sample under key $\mathbf{A}_{\ell+1}$ or a randomly chosen value from $\mathbb{Z}_q^m$.

6. ATTRIBUTE-BASED ENCRYPTION FOR CIRCUITS

In this section, we show how to construct attribute-based encryption for circuits from any TOR scheme.³ Let TOR be the scheme consisting of algorithms (Params, Keygen, Encode) with the “two-to-one” recoding mechanism (Recode, ReKeyGen,

³We point out that our construction and the proof of security generalizes the sample IBE construction presented in Section 4.2.

SimReKeyGen) with input space S . For every $d_{\max}$, let $d_{\max}$ -TOR denote a secure “two-to-one” recoding scheme that is correct for $d_{\max}$ recoding levels.

THEOREM 6.1. *For every ℓ and polynomial $d_{\max} = d_{\max}(\lambda)$, let $\mathcal{C}_{\ell, d_{\max}}$ denote a family of polynomial-size circuits of depth at most $d_{\max}$ that take ℓ bits of input. Assuming the existence of a $d_{\max}$ -TOR scheme, there exists a selectively secure attribute-based encryption scheme $\mathcal{ABE}$ for $\mathcal{C}_{\ell, d_{\max}}$.*

Combining Theorem 6.1 and Lemma 5.1, we obtain a selectively secure attribute-based encryption scheme from LWE. Furthermore, invoking an argument from Boneh and Boyen [2004, Theorem 7.1], and using subexponential hardness of LWE, we obtain a fully secure scheme.

COROLLARY 6.2. *For all ℓ and polynomial $d_{\max} = d_{\max}(\ell)$, there exists a selectively secure attribute-based encryption scheme $\mathcal{ABE}$ for any family of polynomial-size circuits with ℓ inputs and depth at most $d_{\max}$, assuming the hardness of $\text{dLWE}_{n, q, \chi}$ for sufficiently large $n = \text{poly}(\lambda, d_{\max})$, $q = n^{O(d_{\max})}$, and some $\text{poly}(n)$ -bounded error distribution χ .*

Moreover, assuming $2^{O(\ell)}$ -hardness of $\text{dLWE}_{n, q, \chi}$ for parameters $n = \text{poly}(\lambda, d_{\max}, \ell)$ and q and χ as previously, the attribute-based encryption scheme $\mathcal{ABE}$ is fully secure.

The reader is referred to the text after the construction for further explanation of how to choose the LWE parameters. It remains an intriguing open problem to construct fully secure without complexity leveraging (whereas a generic transformation from selective to fully secure functional encryption (more powerful primitive than ABE) is already known [Ananth et al. 2014]).

Observe that if we start with a TOR scheme that supports $d_{\max} = \ell^{\omega(1)}$, then our construction immediately yields an attribute-based encryption scheme for arbitrary polynomial-size circuit families (without any restriction on the depth). This can be achieved if, for example, we had an LWE-based TOR scheme, where q grows polynomially instead of exponentially in $d_{\max}$, as in our LWE-based weak TOR.

We now prove Theorem 6.1.

Circuit Representation. Let $\mathcal{C}_\lambda$ be a collection of circuits each having $\ell = \ell(\lambda)$ input wires and one output wire. Define a collection $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$. For each $C \in \mathcal{C}_\lambda$, we index the wires of C in the following way. The input wires are indexed 1 to ℓ , the internal wires have indices $\ell + 1, \ell + 2, \dots, |C| - 1$, and the output wire has index $|C|$, which also denotes the size of the circuit. We assume that the circuit is composed of arbitrary two-to-one gates. Each gate g is indexed as a tuple (u, v, w) , where u and v are the incoming wire indices and $w > \max\{u, v\}$ is the outgoing wire index. The gate computes the function $g_w : \{0, 1\} \times \{0, 1\} \rightarrow \{0, 1\}$. The “fan-out wires” in the circuit are given a single number. That is, if the outgoing wire of a gate feeds into the input of multiple gates, then all these wires are indexed the same. (See, e.g., [Bellare et al. 2012, Fig. 4].)

6.1. Construction from TOR

The ABE scheme $\mathcal{ABE} = (\text{Setup}, \text{Enc}, \text{KeyGen}, \text{Dec})$ is defined as follows.

Setup($1^\lambda, 1^\ell, d_{\max}$) For each of the ℓ input wires, generate two public/secret key pairs. Also, generate an additional public/secret key pair:

$$\begin{aligned} (\text{pk}_{i,b}, \text{sk}_{i,b}) &\leftarrow \text{Keygen}(\text{pp}) \quad \text{for } i \in [\ell], b \in \{0, 1\}, \\ (\text{pk}_{\text{out}}, \text{sk}_{\text{out}}) &\leftarrow \text{Keygen}(\text{pp}). \end{aligned}$$

Output

$$\text{mpk} := \begin{pmatrix} \text{pk}_{1,0} & \text{pk}_{2,0} & \dots & \text{pk}_{\ell,0} \\ \text{pk}_{1,1} & \text{pk}_{2,1} & \dots & \text{pk}_{\ell,1} & \text{pk}_{\text{out}} \end{pmatrix} \quad \text{msk} := \begin{pmatrix} \text{sk}_{1,0} & \text{sk}_{2,0} & \dots & \text{sk}_{\ell,0} \\ \text{sk}_{1,1} & \text{sk}_{2,1} & \dots & \text{sk}_{\ell,1} \end{pmatrix}.$$

$\text{Enc}(\text{mpk}, \text{ind}, m)$. For $\text{ind} \in \{0, 1\}^\ell$, choose a uniformly random $s \xleftarrow{\$} S$ and encode it under the public keys specified by the index bits:

$$\psi_i \leftarrow \text{Encode}(\text{pk}_{i, \text{ind}_i}, s) \text{ for all } i \in [\ell].$$

Encrypt the message m .

$$\tau \leftarrow E(\text{Encode}(\text{pk}_{\text{out}}, s), m).$$

Output the ciphertext

$$\text{ct}_{\text{ind}} := (\psi_1, \psi_2, \dots, \psi_\ell, \tau, \text{ind}).$$

$\text{KeyGen}(\text{msk}, C)$.

- (1) For every non-input wire $w = \ell + 1, \dots, |C|$ of the circuit C and every $b \in \{0, 1\}$, generate public/secret key pairs.

$$(\text{pk}_{w,b}, \text{sk}_{w,b}) \leftarrow \text{Keygen}(\text{pp}) \text{ if } w < |C| \text{ or } b = 0,$$

and set $\text{pk}_{|C|,1} := \text{pk}_{\text{out}}$.

- (2) For the gate $g = (u, v, w)$ with outgoing wire w , compute the four recoding keys $\text{rk}_{b,c}^w$ (for $b, c \in \{0, 1\}$):

$$\text{rk}_{b,c}^w \leftarrow \text{ReKeyGen}(\text{pk}_{u,b}, \text{pk}_{v,c}, \text{sk}_{u,b}, \text{pk}_{w,g_w(b,c)}).$$

Output the secret key which is a collection of $4(|C| - \ell)$ recoding keys:

$$\text{sk}_C := (\text{rk}_{b,c}^w : w \in [\ell + 1, |C|], b, c \in \{0, 1\})$$

$\text{Dec}(\text{sk}_C, \text{ct}_{\text{ind}})$. For $w = \ell + 1, \dots, |C|$, let $g = (u, v, w)$ denote the gate with outgoing wire w . Suppose wires u and v carry the values b^* and c^* when C is evaluated on ind , so that wire w carries the value $d^* := g_w(b^*, c^*)$. Compute

$$\psi_{w,d^*} \leftarrow \text{Recode}(\text{rk}_{b^*,c^*}^w, \psi_{u,b^*}, \psi_{v,c^*}).$$

If $C(\text{ind}) = 1$, then we would have computed $\psi_{|C|,1}$. Output the message

$$m \leftarrow D(\psi_{|C|,1}, \tau).$$

If $C(\text{ind}) = 0$, output $\perp$.

LWE Parameters. Fix $\ell = \ell(\lambda)$ and $d_{\max} = d_{\max}(\ell)$, and suppose the $\text{dLWE}_{n,m,q,\chi}$ assumption holds for $q = 2^{n^\epsilon}$ for some $0 < \epsilon < 1$. Then, in our LWE-based TOR, we will set:

$$n = \tilde{\Theta}(d_{\max}^{1/\epsilon}) \quad \text{and} \quad q = n^{\Theta(d_{\max})}.$$

By Corollary 6.2, we get security under 2^{n^ϵ} -LWE.

6.2. Correctness

LEMMA 6.3 (CORRECTNESS). *Let $\mathcal{C} = \{C_\lambda\}_{\lambda \in \mathbb{N}}$ be a family, where each C_λ is a finite collection of polynomial-size circuits each of depth at most $d_{\max}$. Let TOR be a correct two-to-one recoding scheme for $d_{\max}$ levels. Then, the construction presented previously is a correct attribute-based encryption scheme.*

PROOF. Fix a circuit C of depth at most $d_{\max}$ and an input ind such that $C(\text{ind}) = 1$. Informally, we rely on recoding correctness for $d_{\max}$ recodings to show that $w = 1, \dots, |C|$; we have

$$\psi_{w,d^*} = \text{Encode}(\text{pk}_{w,d^*}, s),$$

where d^* is the value carried by the wire w and ψ_{w,d^*} is computed as in Dec. Formally, we proceed via induction on w to show that

$$\psi_{w,d^*} \in \Psi_{\text{pk}_{w,d^*,s,j}}.$$

where j is the depth of wire w . The base case $w = 1, \dots, \ell$ follows immediately from correctness of Encode. For the inductive step, consider a wire w at depth j for some gate $g = (u, v, w)$, where $u, v < w$. By the induction hypothesis,

$$\psi_{u,b^*} \in \Psi_{\text{pk}_{u,b^*,s,j_0}}, \quad \psi_{v,c^*} \in \Psi_{\text{pk}_{v,c^*,s,j_1}},$$

where $j_0, j_1 < j$ denote the depths of wires u and v , respectively. It follows immediately from the correctness of Recode that

$$\psi_{w,d^*} \in \Psi_{\text{pk}_{w,d^*,s,\max(i_0,i_1)+1}} \subseteq \Psi_{\text{pk}_{w,d^*,s,j}},$$

which completes the inductive proof. Since $C(\text{ind}) = 1$ and $\text{pk}_{|C|,1} = \text{pk}_{\text{out}}$, we have $\psi_{|C|,1} \in \Psi_{\text{pk}_{\text{out},s,d_{\max}}}$. Finally, by the correctness of (E, D), $D(\psi_{|C|,1}, \tau) = m$. $\square$

6.3. Security

LEMMA 6.4 (SELECTIVE SECURITY). *For any adversary $\mathcal{A}$ against selective security of the attribute-based encryption scheme, there exists an adversary $\mathcal{B}$ against correlated pseudorandomness of TOR whose running time is essentially the same as that of $\mathcal{A}$, such that*

$$\text{Adv}_{\mathcal{A}}^{\text{ABE}}(\lambda) \leq \text{Adv}_{\mathcal{B}}^{\text{CP}}(\lambda) + \text{negl}(\lambda),$$

where $\text{negl}(\lambda)$ captures the statistical security terms in TOR.

We begin by describing alternative algorithms, which would be useful later for constructing the adversary $\mathcal{B}$ for the correlated pseudorandomness security game.

Alternative Algorithms. Fix the selective challenge ind . We get from the “outside” the challenge $\text{pp}, (\text{pk}_i, \psi_i)_{i \in [\ell+1]}$ for correlated pseudorandomness. The main challenge is to design an alternative algorithm KeyGen^* for answering secret key queries without knowing $\text{sk}_{1,\text{ind}_1}, \dots, \text{sk}_{\ell,\text{ind}_\ell}$ or sk_{out} . The algorithm KeyGen^* will maintain the following invariant: on input C with $C(\text{ind}) = 0$, for every non-output wire $w = 1, \dots, |C| - 1$ carrying the value b^* , we will know $\text{sk}_{w,1-b^*}$, but not sk_{w,b^*} .

Moreover, we do not know $\text{sk}_{|C|,0}$ or $\text{sk}_{|C|,1} = \text{sk}_{\text{out}}$.

Setup $^*(\text{ind}, 1^\lambda, 1^\ell, d_{\max})$. Let

$$\begin{aligned} (\text{pk}_{i,1-\text{ind}_i}, \text{sk}_{i,1-\text{ind}_i}) &\leftarrow \text{Keygen}(\text{pp}) \text{ for } i \in [\ell] \\ \text{pk}_{\text{out}} &:= \text{pk}_{\ell+1} \\ \text{pk}_{i,\text{ind}_i} &:= \text{pk}_i \text{ for } i \in [\ell]. \end{aligned}$$

$$\text{Output mpk} = \begin{pmatrix} \text{pk}_{1,0} & \text{pk}_{2,0} & \dots & \text{pk}_{\ell,0} \\ \text{pk}_{1,1} & \text{pk}_{2,1} & \dots & \text{pk}_{\ell,1} & \text{pk}_{\text{out}} \end{pmatrix}.$$

Enc $^*(\text{mpk}, \text{ind}, m)$. Set $\tau \leftarrow E(\psi_{\ell+1}, m)$ and output the ciphertext

$$\text{ct}_{\text{ind}} = (\psi_1, \psi_2, \dots, \psi_\ell, \tau, \text{ind},)$$

where $\psi_1, \dots, \psi_{\ell+1}$ are provided in the challenge.

KeyGen $^*(\text{ind}, \text{msk}, C)$. Where $C(\text{ind}) = 0$, proceed as follows.

- (1) For each internal wire $w \in [\ell + 1, |C| - 1]$ of the circuit C carrying the value b^* when evaluating C on input ind , generate public/secret key pairs:

$$(\text{pk}_{w,1-b^*}, \text{sk}_{w,1-b^*}) \leftarrow \text{Keygen}(\text{pp}).$$

We will generate pk_{w,b^*} using SimReKeyGen as described next.

- (2) For $w = \ell + 1, \dots, |C|$, let $g = (u, v, w)$ denote the gate for which w is the outgoing wire. Suppose wires u and v carry the values b^* and c^* when C is evaluated on ind so that wire w carries the value $d^* := g_w(b^*, c^*)$. By this invariant, we know $\text{sk}_{u,1-b^*}$ and $\text{sk}_{v,1-c^*}$ but not sk_{u,b^*} and sk_{v,c^*} . We start by generating

$$(\text{pk}_{w,d^*}, \text{rk}_{b^*,c^*}^w) \leftarrow \text{SimReKeyGen}(\text{pk}_{u,b^*}, \text{pk}_{v,c^*}).$$

We generate the other three recoding keys using ReKeyGen as follows:

$$\begin{aligned} \text{rk}_{1-b^*,c^*}^w &\leftarrow \text{ReKeyGen}(\text{pk}_{u,1-b^*}, \text{pk}_{v,c^*}, \text{sk}_{u,1-b^*}, \text{pk}_{w,g_w(1-b^*,c^*)}), \\ \text{rk}_{b^*,1-c^*}^w &\leftarrow \text{ReKeyGen}(\text{pk}_{v,1-c^*}, \text{pk}_{u,b^*}, \text{sk}_{v,1-c^*}, \text{pk}_{w,g_w(b^*,1-c^*)}), \\ \text{rk}_{1-b^*,1-c^*}^w &\leftarrow \text{ReKeyGen}(\text{pk}_{u,1-b^*}, \text{pk}_{v,1-c^*}, \text{sk}_{u,1-b^*}, \text{pk}_{w,g_w(1-b^*,1-c^*)}). \end{aligned}$$

Note that $\text{rk}_{1-b^*,c^*}^w, \text{rk}_{1-b^*,1-c^*}^w$ are generated the same way in both KeyGen and KeyGen^* using $\text{sk}_{u,1-b^*}$.

Output the secret key.

$$\text{sk}_C := (\text{rk}_{b,c}^w : w \in [\ell + 1, |C|], b, c \in \{0, 1\}).$$

Informally, the recoding key $\text{rk}_{b^*,1-c^*}^w$ looks the same as in KeyGen because of key indistinguishability, and rk_{b^*,c^*}^w (together with the simulated pk_{w,d^*}) looks the same as in KeyGen because of the recoding simulation property.

Game Sequence. Next, consider the following sequence of games. We use $\text{Adv}_0, \text{Adv}_1, \dots$ to denote the advantage of the adversary $\mathcal{A}$ in Games 0, 1, etc. Game 0 is the real experiment. We then show that the games are either statistically or computationally indistinguishable.

Game i for $i = 1, 2, \dots, q$. As in Game 0, except the challenger answers the first i key queries using KeyGen^* and the remaining $q - i$ key queries using KeyGen . For the i th key query C_i , we consider sub-Games $i.w$ as follows.

Game $i.w$, for $w = \ell + 1, \dots, |C_i|$. The challenger switches $(\text{rk}_{b,c}^w : b, c \in \{0, 1\})$ from KeyGen to KeyGen^* . More precisely,

- we switch $(\text{pk}_{w,d^*}, \text{rk}_{b^*,c^*}^w)$ from KeyGen to KeyGen^* ;
- we switch $\text{rk}_{b^*,1-c^*}^w$ from KeyGen to KeyGen^* ;
- the other two keys $\text{rk}_{1-b^*,c^*}^w, \text{rk}_{1-b^*,1-c^*}^w$ are generated the same way in both KeyGen and KeyGen^* .

From Lemma 6.5, we have

$$|\text{Adv}_i - \text{Adv}_{i+1}| \leq \text{negl}(\lambda) \text{ for all } i.$$

Note that in Game q , the challenger runs Setup^* and answers all key queries using KeyGen^* with the selective challenge ind and generates the challenge ciphertext using Enc .

Game $q + 1$. Same as Game q , except the challenger generates the challenge ciphertext using Enc^* with $\psi_{\ell+1} = \text{Encode}(\text{pk}_{\ell+1}, s)$. Clearly,

$$\text{Adv}_{q+1} = \text{Adv}_q.$$

Game $q + 2$. Same as Game $q + 1$, except $\psi_{\ell+1} \xleftarrow{\$} \mathcal{K}$. From Lemma 6.6, there is an adversary $\mathcal{B}$ such that

$$|\text{Adv}_{q+1} - \text{Adv}_{q+2}| \leq \text{Adv}_{\mathcal{B}}^{\text{CP}}(\lambda).$$

Finally, by the one-time semantic security of (E, D) for all $m_0, m_1 \in \mathcal{M}$, $\text{E}(\psi_{\ell+1}, m_0)$ is indistinguishable from $\text{E}(\psi_{\ell+1}, m_1)$ (for $\psi_{\ell+1} \xleftarrow{\$} \mathcal{K}$). Therefore, it follows that $\text{Adv}_{q+2} \leq \text{negl}(\lambda)$, concluding the proof of security.

LEMMA 6.5. *Games i and $i + 1$ are (statistically) indistinguishable for all $i = 0, \dots, q - 1$.*

PROOF. In Game $i + 1$, first $i + 1$ key queries are answered using KeyGen^* and the remaining $q - (i + 1)$ key queries using KeyGen . Consider a key query C_{i+1} and for all $w = \ell, \dots, |C_i|$, sub-Games $(i + 1).w$ defined previously. We argue that for all $w = \ell, \dots, |C_i| - 1$, sub-Games $(i + 1).w$ and $(i + 1).(w + 1)$ are statistically indistinguishable. Clearly, sub-Game $(i + 1).\ell$ corresponds to Game i , where the $(i + 1)$ th key query is sampled using KeyGen and sub-Game $(i + 1).|C_i|$ this corresponds to Game $i + 1$, where $(i + 1)$ th query is the sampled using KeyGen^* (the remaining queries are sampled identically). Now, sub-Games $(i + 1).w$ and $(i + 1).(w + 1)$ differ in how recoding keys $(\text{rk}_{b,c}^{w+1} : b, c \in \{0, 1\})$ are sampled for query C_{i+1} . The statistical indistinguishability of keys $\text{rk}_{b^*,c^*}^{w+1}$ follows from recoding simulation $\text{rk}_{b^*,1-c^*}^{w+1}$ on indistinguishability with respect to sampling using keys sk_{b^*} and sk_{1-c^*} , and two other keys $\text{rk}_{1-b^*,c^*}^{w+1}, \text{rk}_{1-b^*,1-c^*}^{w+1}$ are generated identically in two experiments. It follows that the two sub-Games $(i + 1).w$ and $(i + 1).(w + 1)$ are statistically indistinguishable, concluding the lemma. $\square$

LEMMA 6.6. *Games $q + 1$ and $q + 2$ are computationally indistinguishable.*

PROOF. Suppose there exists an adversary B^* that distinguishes between the Games $q + 1$ and $q + 2$. We construct an adversary B that breaks the correlated pseudorandomness property. Adversary B gets as input $(\text{pk}_i, \psi_i)_{i \in [\ell]}, \text{pk}_{\ell+1}, \psi^*$, where ψ^* is either a valid encoding using key $\text{pk}_{\ell+1}$ or a randomly chosen element. Adversary B uses these values to invoke algorithms $\text{Setup}^*, \text{Enc}^*, \text{KeyGen}^*$, where it uses $\psi_{\ell+1} := \psi^*$. It invokes adversary B^* and the same thing. Clearly, if $\psi^* = \text{Encode}(\text{pk}_{\ell+1}, s)$ then this experiment corresponds to Game $q + 1$. On the other hand, if $\psi^* \xleftarrow{\$} \mathcal{K}$, then this experiment corresponds to Game $q + 2$. This concludes the proof of the lemma. $\square$

7. ATTRIBUTE-BASED ENCRYPTION FOR BRANCHING PROGRAMS

In this section, we present weak TOR and attribute-based encryption for branching programs, which capture the complexity class log-space. As noted in Section 2.2, we exploit the fact that in branching programs, the transition function depends on an input variable and the current state; this means that one of the two input encodings during recoding is always a “depth 0” encoding.

Branching Programs. Recall that a branching program Γ is a directed acyclic graph in which every nonterminal node has exactly two outgoing edges labeled $(i, 0)$ and $(i, 1)$ for some $i \in [\ell]$. Moreover, there is a distinguished terminal accept node. Every input $x \in \{0, 1\}^\ell$ naturally induces a subgraph Γ_x containing exactly those edges labeled (i, x_i) . We say that Γ accepts x if and only if there is a path from the start node to the accept node in Γ_x . At the cost of possibly doubling the number of edges and vertices, we may assume that there is at most one edge connecting any two nodes in Γ .

7.1. Weak TOR

A weak two-to-one encoding (wTOR) scheme consists of the same algorithms as TOR, except that $\text{Keygen}(\text{pp}, j)$ takes an additional input $j \in \{0, 1\}$. That is, Keygen may produce different distributions of public/secret key pairs depending on j . Moreover, in ReKeyGen , the first public key is always generated using $\text{Keygen}(\text{pp}, 0)$ and the second using $\text{Keygen}(\text{pp}, 1)$; similarly, in Recode , the first encoding is always generated with respect to a public key from $\text{Keygen}(\text{pp}, 0)$ and the second from $\text{Keygen}(\text{pp}, 1)$. Similarly, the correctness and statistical security properties are relaxed.

Correctness. First, for every pk and $s \in \mathcal{S}$, there exists a family of sets $\Psi_{\text{pk},s,j}$, $j = 0, 1, \dots, d_{\max}$.

— $\Psi_{\text{pk},s,1} \subseteq \dots \subseteq \Psi_{\text{pk},s,d_{\max}}$.

— For all $\psi, \psi' \in \Psi_{\text{pk},s,d_{\max}}$ and all $m \in \mathcal{M}$,

$$D(\psi', E(\psi, m)) = m.$$

Second, the correctness of recoding requires that for any triple of key pairs $(\text{pk}_0, \text{sk}_0)$, $(\text{pk}_1, \text{sk}_1)$, $(\text{pk}_{\text{tgt}}, \text{sk}_{\text{tgt}})$, respectively, in the support of $\text{Keygen}(\text{pp}, 0)$, $\text{Keygen}(\text{pp}, 1)$, $\text{Keygen}(\text{pp}, 1)$ and any encodings $\psi_0 \in \text{Encode}(\text{pk}_0, s)$ and $\psi_1 \in \Psi_{\text{pk}_1,s,j_1}$, where $0 < j_1$,

$$\text{Recode}(\text{rk}, \psi_0, \psi_1) \in \Psi_{\text{pk}_{\text{tgt}},s,j_1+1}.$$

Statistical Security Properties. We require recoding simulation as before, but not key indistinguishability. However, we require the following additional property.

(Back-Tracking). For all $(\text{pk}_0, \text{sk}_0) \leftarrow \text{Keygen}(\text{pp}, 0)$ and all $(\text{pk}_1, \text{sk}_1)$, $(\text{pk}_{\text{tgt}}, \text{sk}_{\text{tgt}}) \leftarrow \text{Keygen}(\text{pp}, 1)$, the following distributions are identical.

$$\text{ReKeyGen}(\text{pk}_0, \text{pk}_1, \text{sk}_0, \text{pk}_{\text{tgt}}) \equiv -\text{ReKeyGen}(\text{pk}_0, \text{pk}_{\text{tgt}}, \text{sk}_0, \text{pk}_1).$$

Informally, this says that switching the order of pk_1 and pk_{tgt} as inputs to ReKeyGen is the same as switching the “sign” of the output. In our instantiations, the output of ReKeyGen lies in a group, so negating the output simply refers to applying the group inverse operation.

Computational Security Property. We define the advantage function $\text{Adv}_{\mathcal{A}}^{\text{CP}}(\lambda)$ (modified to account for the additional input to Keygen) to be the absolute value of

$$\Pr \left[\begin{array}{l} \text{pp} \leftarrow \text{Setup}(1^\lambda); s \leftarrow \mathcal{S}; \\ (\text{pk}_i, \text{sk}_i) \leftarrow \text{Keygen}(\text{pp}, 0), \\ \psi_i \leftarrow \text{Encode}(\text{pk}_i, s), i = 1, \dots, \ell; \\ (\text{pk}_{\ell+1}, \text{sk}_{\ell+1}) \leftarrow \text{Keygen}(\text{pp}, 1); \\ \psi'_0 \leftarrow \text{Encode}(\text{pk}_{\ell+1}, s); \\ b \xleftarrow{\$} \{0, 1\}; \psi'_1 \xleftarrow{\$} \mathcal{K} \\ b' \leftarrow \mathcal{A}(\text{pk}_1, \dots, \text{pk}_{\ell+1}, \psi_1, \dots, \psi_\ell, \psi'_b) \end{array} \right] - \frac{1}{2},$$

and we require that for all PPT $\mathcal{A}$, the advantage function $\text{Adv}_{\mathcal{A}}^{\text{CP}}(\lambda)$ is a negligible function in λ .

Remark 7.1. Due to the additional back-tracking property, it is not the case that a TOR implies a weak TOR. However, we are able to instantiate weak TOR under weaker and larger classes of assumptions than TOR.

7.2. Weak TOR from LWE

We provide an instantiation of weak TOR from LWE. The main advantage over our construction of TOR in Section 5 is that the dependency of q on $d_{\max}$ is linear in $d_{\max}$ instead of exponential. Therefore, if q is quasi-polynomial, we can handle any polynomial $d_{\max}$, as opposed to an a-prior bounded $d_{\max}$.

LEMMA 7.2. *Assuming $\text{dLWE}_{n,(\ell+2)m,q,\chi}$, where $q = O(d_{\max} n^3 \log n)$, there is a weak TOR scheme that is correct up to $d_{\max}$ levels, where $n = n(\lambda)$, the error distribution $\chi = \chi(n) = D_{\mathbb{Z}, \sqrt{n}}$ the modulus $q = q(n) = d_{\max} \cdot O(n^3 \log n)$, and the number of samples $m = m(n) = O(n \log q)$.*

Note that the parameters here are better than in Lemma 5.1. The construction of weak TOR from learning with errors follows.

- Params($1^\lambda, d_{\max}$). We are given parameters n, m, χ, q , and we set the error bound $B = B(n) = O(n)$; and the Gaussian parameter $s = s(n) = O(\sqrt{n \log q})$; output the global public parameters $\text{pp} = (n, \chi, B, q, m, s)$; and define the domain $\mathcal{S}$ of the encoding scheme to be $\mathbb{Z}_q^n$.
- Keygen(pp, j). Run the trapdoor generation algorithm $\text{TrapGen}(1^n, 1^m, q)$ to obtain a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ together with the trapdoor $\mathbf{T}$. Outputs

$$\text{pk} = \mathbf{A}; \quad \text{sk} = \mathbf{T}.$$

- Encode($\mathbf{A}, \mathbf{s}$). Sample an error vector $\mathbf{e} \leftarrow \chi^m$ and output the encoding $\psi := \mathbf{A}^T \mathbf{s} + \mathbf{e} \in \mathbb{Z}_q^m$.
- ReKeyGen($\mathbf{A}_0, \mathbf{A}_1, \mathbf{A}_{\text{tgt}}, \mathbf{T}$). Outputs a low-norm matrix $\mathbf{R}$ such that $\mathbf{A}_0 \mathbf{R} = \mathbf{A}_{\text{tgt}} - \mathbf{A}_1$. In particular,

$$\mathbf{R} \leftarrow \text{SampleD}(\mathbf{A}_0, \mathbf{T}_0, \mathbf{A}_{\text{tgt}} - \mathbf{A}_1, s).$$

- SimReKeyGen($\mathbf{A}_0, \mathbf{A}_1$). Sample a matrix $\mathbf{R} \leftarrow (D_{\mathbb{Z}, s})^{m \times m}$ by sampling each entry from the discrete Gaussian distribution $D_{\mathbb{Z}, s}$. Outputs

$$\text{rk} := \mathbf{R}; \quad \mathbf{A}_{\text{tgt}} := \mathbf{A}_0 \mathbf{R} + \mathbf{A}_1.$$

- Recode($\text{rk}, \psi_0, \psi_1$). Outputs $\text{rk}^T \psi_0 + \psi_1$.

Correctness. We define the sets $\Psi_{\mathbf{A}, \mathbf{s}, j}$ for $\text{pk} := \mathbf{A} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{s} \in \mathbb{Z}_q^n$, and $j \in [1 \dots d_{\max}]$ as follows.

$$\Psi_{\mathbf{A}, \mathbf{s}, j} = \{\mathbf{A}^T \mathbf{s} + \mathbf{e} : \|\mathbf{e}\|_\infty \leq B \cdot j \cdot (sm\sqrt{m})\}.$$

The analysis is similar to that in the previous section. In particular, we observe the following right away.

- $\Psi_{\mathbf{A}, \mathbf{s}, 1} \subseteq \Psi_{\mathbf{A}, \mathbf{s}, 1} \subseteq \dots \subseteq \Psi_{\mathbf{A}, \mathbf{s}, d_{\max}}$.
- For any two encodings $\psi, \psi' \in \Psi_{\mathbf{A}, \mathbf{s}, d_{\max}}$ and $\mu \in \{0, 1\}^n$, $D(\psi', E(\psi, \mu)) = \mu$, as long as

$$B \cdot d_{\max} \cdot (sm\sqrt{m}) \leq q/4.$$

- Consider two encodings $\mathbf{A}^T \mathbf{s} + \mathbf{e} \in \text{Encode}(\mathbf{A}, \mathbf{s})$ and $\psi_1 \in \Psi_{\mathbf{A}_1, \mathbf{s}, j_1}$ for any $j_1 \in \mathbb{N}$. Then, $\psi_0 = \mathbf{A}_0^T \mathbf{s} + \mathbf{e}_0$ and $\psi_1 := \mathbf{A}_1^T \mathbf{s} + \mathbf{e}_1$, where $\|\mathbf{e}_0\|_\infty \leq B$ and $\|\mathbf{e}_1\|_\infty \leq j_1 \cdot B \cdot (sm\sqrt{m})$. Then, the recoded ciphertext ψ_{tgt} is computed as follows:

$$\begin{aligned} \psi_{\text{tgt}} &:= \mathbf{R}^T \psi_0 + \psi_1 \\ &= \mathbf{R}^T (\mathbf{A}_0^T \mathbf{s} + \mathbf{e}_0) + (\mathbf{A}_1^T \mathbf{s} + \mathbf{e}_1) \\ &= \mathbf{A}_{\text{tgt}}^T \mathbf{s} + \mathbf{e}_{\text{tgt}}, \end{aligned}$$

where the last equation is because $\mathbf{A}_{\text{tgt}} = \mathbf{A}_0 \mathbf{R} + \mathbf{A}_1$, and we define $\mathbf{e}_{\text{tgt}} := \mathbf{R}^T \mathbf{e}_0 + \mathbf{e}_1$. Thus,

$$\begin{aligned} \|\mathbf{e}_{\text{tgt}}\|_\infty &\leq m \cdot \|\mathbf{R}\|_\infty \|\mathbf{e}_0\|_\infty + \|\mathbf{e}_1\|_\infty \\ &\leq m \cdot s\sqrt{m} \cdot B + B \cdot j_1 \cdot (sm\sqrt{m}) \\ &= (j_1 + 1) \cdot B \cdot (sm\sqrt{m}), \end{aligned}$$

exactly as required. Here, the second inequality is because $\|\mathbf{R}\|_\infty \leq s\sqrt{m}$ by Lemma 3.2. This finishes our proof of correctness.

Security. Correlated pseudorandomness follows from $\text{dLWE}_{n,(\ell+2)m,q,\chi}$, where $q = n \cdot d_{\max}$. In particular, by the dLWE, given $\mathbf{A}_1, \dots, \mathbf{A}_{\ell+1}, \psi_1, \dots, \psi_\ell, \psi^*$, no adversary can distinguish between the cases when ψ^* is a valid LWE sample under key $\mathbf{A}_{\ell+1}$ or a randomly chosen value from $\mathbb{Z}_q^m$. Recoding simulation follows readily from Lemma 3.2 by an argument identical to the one for the construction of TOR in Section 5. For back-tracking, negation of a recoding matrix $\mathbf{R}$ is simply the additive inverse over $\mathbb{Z}_q^m$. That is, the output of $\text{ReKeyGen}(\text{pk}_0, \text{pk}_1, \text{sk}_0, \text{pk}_{\text{tgt}})$ is a matrix $\mathbf{R}$ such that $\mathbf{A}_0 \mathbf{R} = \mathbf{A}_{\text{tgt}} - \mathbf{A}_1$, and the output of $\text{ReKeyGen}(\text{pk}_0, \text{pk}_{\text{tgt}}, \text{sk}_0, \text{pk}_1)$ is also a matrix $\mathbf{R}'$ such that $\mathbf{A}_0 \mathbf{R}' = \mathbf{A}_1 - \mathbf{A}_{\text{tgt}}$. It is easy to see that the output $\mathbf{R}$ is distributed identically to $-\mathbf{R}'$.

7.3. Weak TOR from Bilinear Maps

We use asymmetric groups for maximal generality and for conceptual clarity. Let $G_1, G_2, G_T \leftarrow \text{GroupGen}(1^\lambda)$ be descriptions of cyclic groups (with corresponding generators) of prime order q and $e : G_1 \times G_2 \rightarrow G_T$ is a nondegenerate bilinear map. We require that the group operations in G and G_T as well the bilinear map e are computable in deterministic polynomial time with respect to λ . Let g_1, g_2 denote random generators of G_1, G_2 , respectively. The DBDH assumption says that the following two distributions are computationally indistinguishable,

$$[g_1, g_2, g_1^a, g_2^a, g_2^b, g_1^s, e(g_1, g_2)^{abs}] \stackrel{c}{\approx} [g_1, g_2, g_1^a, g_2^a, g_2^b, g_1^s, g_T^c]$$

where a, b, s , and c are randomly chosen from $\mathbb{Z}_q$.

—Params($1^\lambda, d_{\max}$). Outputs $\text{pp} := (g_1, g_2, g_1^a, g_2^a)$.

—Keygen(pp, j).

(i) If $j = 0$, then samples $t \xleftarrow{\$} \mathbb{Z}_q$ and outputs

$$(\text{pk}, \text{sk}) := ((g_1^{a/t}, g_2^{a/t}), t).$$

(ii) If $j \geq 1$, output $\text{pk} \xleftarrow{\$} G_2$.

—Encode(pk, s).

(i) If $\text{pk} = (g_1^{a/t}, g_2^{a/t}) \in G_1 \times G_2$, output $(g_1^{a/t})^s$.

(ii) If $\text{pk} \in G_2$, output $e(g_1^a, \text{pk})^s$.

—Recode(rk, c_0, c_1). Outputs $e(c_0, \text{rk}) \cdot c_1$.

—ReKeyGen($(g_1^{a/t}, g_2^{a/t}), \text{pk}_1, \text{pk}_{\text{tgt}}, t$). Outputs $\text{rk} := (\text{pk}_{\text{tgt}} \cdot \text{pk}_1^{-1})^t \in G_2$.

—SimReKeyGen($(g_1^{a/t}, g_2^{a/t}), \text{pk}_1$). Picks $z \xleftarrow{\$} \mathbb{Z}_q$ and outputs

$$\text{rk} := (g_2^{a/t})^z, \quad \text{pk}_{\text{tgt}} := \text{pk}_1 \cdot (g_2^a)^z.$$

Correctness. Define $\Psi_{\text{pk},s,j} := \{\text{Encode}(\text{pk}, s)\}$. For recoding, observe that

$$\begin{aligned} & \text{Recode}((\text{pk}_{\text{tgt}} \cdot \text{pk}_1^{-1})^t, g_1^{as/t}, e(g_1^a, \text{pk}_1)^s) \\ &= e(g_1^{as/t}, (\text{pk}_{\text{tgt}} \cdot \text{pk}_1^{-1})^t) \cdot e(g_1^a, \text{pk}_1)^s \\ &= e(g_1^a, (\text{pk}_{\text{tgt}} \cdot \text{pk}_1^{-1})^s) \cdot e(g_1^a, \text{pk}_1)^s \\ &= e(g_1^a, \text{pk}_{\text{tgt}})^s = \text{Encode}(\text{pk}_{\text{tgt}}, s). \end{aligned}$$

Security. For back-tracking, negation is simply the multiplicative inverse over G_q . Correlation pseudorandomness follows readily from the DBDH assumption. In particular, suppose there is an adversary $\mathcal{D}$ that can distinguish between

$$(g_1, g_1, g_1^a, g_2^a, (g_1^{a/t}, g_2^{a/t}, (g_1^{a/t})^s), (g_2^{b_1}, e(g_1^a, g_2^{b_1})^s), \dots, g_2^{b_\ell}, e(g_1^a, g_2^{b_\ell})^s, g_2^{b_{\ell+1}}, e(g_1^a, g_2^{b_{\ell+1}})^s)$$

and

$$(g_1, g_1, g_1^a, g_2^a, (g_1^{a/t}, g_2^{a/t}, (g_1^{a/t})^s), (g_2^{b_1}, e(g_1^a, g_2^{b_1})^s), \dots, g_2^{b_\ell}, e(g_1^a, g_2^{b_\ell})^s, g_2^{b_{\ell+1}}, g_T^c)$$

for a randomly chosen c from $\mathbb{Z}_q$. Then, we can construct an adversary $\mathcal{A}$ that breaks DBDH assumption. In particular, $\mathcal{A}$ is given $g_1, g_2, g_1^a, g_2^a, g_2^b, g_1^s, \alpha$, where α is either g_T^{abs} or g_T^c for a randomly chosen c . Then, the adversary $\mathcal{A}$ can simulate components $(g_1^{a/t}, g_2^{a/t}, (g_1^{a/t})^s)$ by sampling t' and setting $t = a/t'$. Then, $g_1^{a/t} = g_1^{t'}$, $g_2^{a/t} = g_2^{t'}$ and $(g_1^{a/t})^s = g_1^{t's}$ (which can be computed given g_1^s and t'). It can also simulate $g_2^{b_i}, e(g_1^a, g_2^{b_i})^s$ for all $i \leq \ell$, by sampling b_i in clear and then using g_1^s, g_2^a , and b_i to compute $e(g_1^s, g_2^a)^{b_i} = e(g_1^a, g_2^{b_i})^s$. Finally, it sets $b_{\ell+1} := b$ from the challenge together with α and calls the adversary $\mathcal{D}$ to distinguish.

7.4. Attribute-Based Encryption from Weak TOR

Setup($1^\lambda, 1^\ell, d_{\max}$). For each one of ℓ input bits, generate two public/secret key pairs. Also, generate a public/secret key pair for the start and accept states:

$$\begin{aligned} (\text{pk}_{i,b}, \text{sk}_{i,b}) &\leftarrow \text{Keygen}(\text{pp}, 0) \quad \text{for } i \in [\ell], b \in \{0, 1\}, \\ (\text{pk}_{\text{start}}, \text{sk}_{\text{start}}) &\leftarrow \text{Keygen}(\text{pp}, 1), \\ (\text{pk}_{\text{accept}}, \text{sk}_{\text{accept}}) &\leftarrow \text{Keygen}(\text{pp}, 1). \end{aligned}$$

Output

$$\begin{aligned} \text{mpk} &:= \begin{pmatrix} \text{pk}_{1,0} & \text{pk}_{2,0} & \dots & \text{pk}_{\ell,0} & \text{pk}_{\text{start}} \\ \text{pk}_{1,1} & \text{pk}_{2,1} & \dots & \text{pk}_{\ell,1} & \text{pk}_{\text{accept}} \end{pmatrix}, \\ \text{msk} &:= \begin{pmatrix} \text{sk}_{1,0} & \text{sk}_{2,0} & \dots & \text{sk}_{\ell,0} & \text{sk}_{\text{start}} \\ \text{sk}_{1,1} & \text{sk}_{2,1} & \dots & \text{sk}_{\ell,1} & \text{sk}_{\text{accept}} \end{pmatrix}. \end{aligned}$$

Enc($\text{mpk}, \text{ind}, m$). For $\text{ind} \in \{0, 1\}^\ell$, choose a uniformly random $s \xleftarrow{\$} S$ and encode it under the public keys specified by the index bits and the start state:

$$\begin{aligned} \psi_i &\leftarrow \text{Encode}(\text{pk}_{i,\text{ind}_i}, s), \quad \text{for all } i \in [\ell], \\ \psi_{\text{start}} &\leftarrow \text{Encode}(\text{pk}_{\text{start}}, s). \end{aligned}$$

Encrypt the message

$$\tau \leftarrow E(\text{Encode}(\text{pk}_{\text{accept}}, s), m).$$

Output the ciphertext

$$\text{ct}_{\text{ind}} = (\psi_1, \psi_2, \dots, \psi_\ell, \psi_{\text{start}}, \tau).$$

KeyGen(msk, Γ): $\Gamma : \{0, 1\}^\ell \rightarrow \{0, 1\}$ is a branching program that takes a ℓ -bit input and outputs a single bit.

—For every node u , except the start and accept nodes, sample public/secret key pair:

$$(\text{pk}_u, \text{sk}_u) \leftarrow \text{Keygen}(\text{pp}, 1).$$

—For every edge (u, v) labeled (i, b) in Γ , sample a recoding key $\text{rk}_{u,v}$ as follows:

$$\text{rk}_{u,v} \leftarrow \text{ReKeyGen}(\text{pk}_{i,b}, \text{pk}_u, \text{sk}_{i,b}, \text{pk}_v).$$

The secret key sk_Γ is the collection of all the recoding keys $\text{rk}_{u,v}$ for every edge (u, v) in Γ .

$\text{Dec}(\text{sk}_\Gamma, \text{ct}_{\text{ind}})$. Suppose $\Gamma(\text{ind}) = 1$; output $\perp$ otherwise. Let Π denote the (directed) path from the start node to the accept node in Γ_{ind} . For every edge (u, v) labeled (i, ind_i) in Π , apply the recoding algorithm on the two encodings ψ_i, ψ_u and the recoding key $\text{rk}_{u,v}$:

$$\psi_v \leftarrow \text{Recode}(\text{rk}_{u,v}, \psi_i, \psi_u)$$

If $\Gamma(\text{ind}) = 1$, we obtain ψ_{accept} . Decrypt and output the message

$$m \leftarrow \text{D}(\psi_{\text{accept}}, \tau).$$

7.4.1. Correctness

LEMMA 7.3 (CORRECTNESS). *Let $\mathcal{G} = \{\Gamma\}_\lambda$ be a collection of polynomial-size branching programs of depth at most d_{max} , and let wTOR be a weak two-to-one recoding scheme for d_{max} levels. Then, the construction previously presented is a correct attribute-based encryption scheme for $\mathcal{G}$.*

PROOF. Let Π denote the directed path from the start to the accept nodes in Γ_{ind} . We shows, via induction on nodes v along the path Π , that

$$\psi_v \in \Psi_{\text{pk}_v, s, j},$$

where j is the depth of node v along the path. The base case for $v := \text{start node}$ follows immediately from correctness of Encode . For the inductive step, consider a node v along the path Π at depth j for some edge (u, v) labeled (i, ind_i) . By the induction hypothesis,

$$\psi_u \in \Psi_{\text{pk}_u, s, j_0},$$

where $j_0 < j$ denote the depths of node u . Also by the correctness of the Encode algorithm, for all $i \in [\ell]$,

$$\psi_i \in \Psi_{\text{pk}_{i, \text{ind}_i}, s, 0}.$$

It follows immediately from the correctness of Recode that

$$\psi_v \in \Psi_{\text{pk}_v, s, j_0+1} \subseteq \Psi_{\text{pk}_v, s, j},$$

which completes the inductive proof. Since $C(\text{ind}) = 1$, we have

$$\psi_{\text{accept}} \in \Psi_{\text{pk}_{\text{accept}}, s, d_{\text{max}}}.$$

Recall that $\tau \leftarrow \text{E}(\text{Encode}(\text{pk}_{\text{accept}}, s), m)$. Finally, by the correctness of (E, D) ,

$$\text{D}(\psi_{\text{accept}}, \tau) = m. \quad \square$$

7.4.2. Selective Security

LEMMA 7.4 (SELECTIVE SECURITY). *For any adversary $\mathcal{A}$ against selective security of the attribute-based encryption scheme for branching programs, there exists an adversary $\mathcal{B}$ against correlated pseudorandomness of wTOR whose running time is essentially the same as that of $\mathcal{A}$, such that*

$$\text{Adv}_{\mathcal{A}}^{\text{ABE}}(\lambda) \leq \text{Adv}_{\mathcal{B}}^{\text{CP}}(\lambda) + \text{negl}(\lambda),$$

where $\text{negl}(\lambda)$ captures the statistical security terms in TOR .

In the proof of security, we will rely crucially on the following combinatorial property of branching programs: for any input x , the graph Γ_x does not contain any cycles as an undirected graph.

Alternative Algorithms. Fix the selective challenge ind . We also get a collection of public keys with corresponding encodings from the “outside”: $(\text{pk}_i, \psi_i)_{i \in [\ell+2]}$, where the challenge is to decide whether $\psi_{\ell+1}$ is $\text{Encode}(\text{pk}_{\ell+2}, s)$ or random. The main challenge is designing an alternative algorithm KeyGen^* for answering secret key queries without knowing $\text{sk}_{1, \text{ind}_1}, \dots, \text{sk}_{\ell, \text{ind}_\ell}$ or $\text{sk}_{\text{start}}, \text{sk}_{\text{accept}}$. We consider the following “alternative” algorithms.

$\text{Setup}^*(1^\lambda, 1^\ell, d_{\max})$. Let

$$\begin{aligned} (\text{pk}_{i, 1-\text{ind}_i}, \text{sk}_{i, 1-\text{ind}_i}) &\leftarrow \text{Keygen}(\text{pp}, 0) \text{ for } i \in [\ell] \\ \text{pk}_{i, \text{ind}_i} &:= \text{pk}_i \text{ for } i \in [\ell] \\ \text{pk}_{\text{start}} &:= \text{pk}_{\ell+1} \\ \text{pk}_{\text{accept}} &:= \text{pk}_{\ell+2}. \end{aligned}$$

Define and output the master public key as follows.

$$\text{mpk} = \begin{pmatrix} \text{pk}_{1,0} & \text{pk}_{2,0} & \dots & \text{pk}_{\ell,0} & \text{pk}_{\text{start}} \\ \text{pk}_{1,1} & \text{pk}_{2,1} & \dots & \text{pk}_{\ell,1} & \text{pk}_{\text{accept}} \end{pmatrix}.$$

$\text{Enc}^*(\text{mpk}, \text{ind}, m)$. Define

$$\begin{aligned} \psi_{i, \text{ind}_i} &:= \psi_i \quad \text{for all } i \in [\ell] \\ \psi_{\text{start}} &:= \psi_{\ell+1} \\ \psi_{\text{accept}} &:= \psi_{\ell+2}. \end{aligned}$$

Encrypt the message m .

$$\tau \leftarrow \text{E}(\psi_{\text{accept}}, b).$$

Output the simulated ciphertext.

$$\text{ct}_{\text{ind}} = (\psi_1, \psi_2, \dots, \psi_\ell, \psi_{\text{start}}, \tau).$$

$\text{KeyGen}^*(\text{msk}, \Gamma)$. Let Γ'_{ind} denote the undirected graph obtained from Γ_{ind} by treating every directed edge as an undirected edge (while keeping the edge label). Observe that Γ'_{ind} satisfies the following properties.

- Γ'_{ind} contains no cycles. This is because Γ_{ind} is acyclic and every nonterminal node contains exactly one outgoing edge.
- The start node and the accept node lie in different connected components in Γ'_{ind} , since $\Gamma(\text{ind}) = 0$.

Simulation Invariant. For each “active” edge labeled (i, ind_i) from node u to node v , simulate the recoding key. Choose our own public/secret key pair for each “inactive” edges $(i, 1 - \text{ind}_i)$ and generate the recoding key honestly.

- Run a DFS in Γ'_{ind} starting from the start node. Whenever we visit a new node v from a node u along an edge labeled (i, ind_i) , we set

$$\begin{aligned} (\text{pk}_v, \text{rk}_{u,v}) &\leftarrow \text{SimReKeyGen}(\text{pk}_{i, \text{ind}}, \text{pk}_u) \text{ if } (u, v) \text{ is a directed edge in } \Gamma, \\ (\text{pk}_v, -\text{rk}_{v,u}) &\leftarrow \text{SimReKeyGen}(\text{pk}_{i, \text{ind}}, \text{pk}_u) \text{ if } (v, u) \text{ is a directed edge in } \Gamma. \end{aligned}$$

Here, we exploit the back-tracking property in wTOR .

Note that since $\Gamma(\text{ind}) = 0$, then the accept node is not assigned a public key by this process.

- For all nodes u without an assignment, run $(\text{pk}_u, \text{sk}_u) \leftarrow \text{Keygen}(\text{pp}, 1)$.
- For every remaining edge (u, v) labeled $(i, 1 - \text{ind}_i)$ in Γ , sample a recoding key $\text{rk}_{u,v}$ as in KeyGen using $\text{sk}_{i, 1-\text{ind}}$ as follows.

$$\text{rk}_{u,v} \leftarrow \text{ReKeyGen}(\text{pk}_{i, 1-\text{ind}}, \text{pk}_u, \text{sk}_{i, 1-\text{ind}}, \text{pk}_v).$$

The secret key sk_Γ is simply the collection of all the recoding keys $rk_{u,v}$ for every edge (u, v) in Γ .

Game Sequence. Next, consider the following sequence of games. We use $Adv_0, Adv_1, \dots$ to denote the advantage of the adversary $\mathcal{A}$ in Games 0, 1, etc. Let n denote the number of edges in a branching program Γ labeled (i, ind_i) for some i , and for all $j \in [n]$, let e_j denote the actual edge.

Game 0. Real experiment.

Game i for $i = 1, 2, \dots, q$. As in Game 0, except the challenger answers the first i key queries using $KeyGen^*$ and the remaining $q - i$ key queries using $KeyGen$. For the i th key query Γ_i , we also consider sub-Games, as follows.

—Game $i.j$, for $j = 1, \dots, n$. For edge $e_j = (u, v)$ labeled (i, ind_i) , the challenger switches the simulated recoding key $rk_{u,v}$ from $KeyGen$ to $KeyGen^*$. From Lemma 7.5, it follows that

$$|Adv_i - Adv_{i+1}| \leq \text{negl}(\lambda) \text{ for all } i.$$

Note that in Game q , the challenger runs $Setup^*$ and answers all key queries using $KeyGen^*$ with the selective challenge ind and generates the challenge ciphertext using Enc .

—Game $q+1$. Same as Game q , except the challenger generates the challenge ciphertext using Enc^* with $\psi_{\ell+2} \leftarrow \text{Encode}(pk_{\ell+2}, s)$.

$$Adv_{q+1} = Adv_q.$$

—Game $q+2$ Same as Game $q+1$, except $\psi_{\ell+2} \xleftarrow{\$} \mathcal{K}$. It is straight forward to construct an adversary $\mathcal{B}$ such that

$$|Adv_{q+1} - Adv_{q+2}| \leq Adv_B^{\text{CP}}(\lambda).$$

Finally, by the one-time semantic security of (E, D) for all $m_0, m_1 \in \mathcal{M}$, $E(\psi_{\ell+2}, m_0)$ is indistinguishable from $E(\psi_{\ell+2}, m_1)$ (for $\psi_{\ell+2} \xleftarrow{\$} \mathcal{K}$). Therefore, it follows that $Adv_{q+2} \leq \text{negl}(\lambda)$, concluding the proof of security.

LEMMA 7.5. *Games i and $i+1$ are (statistically) indistinguishable for all $i = 0, \dots, q-1$.*

PROOF. The difference in two games is in how the $(i+1)$ st key query is answered. In Game i , it is answered using $KeyGen$, whereas in Game $i+1$, it is answered using $KeyGen^*$. Now, consider sub-Games $(i+1).j$ and $(i+1).(j+1)$. The difference in these two sub-Games is in the recoding key $rk_{u,v}$. Now, by the recoding simulation and backtracking, it follows that the keys sampled statistically indistinguishable, concluding the lemma. $\square$

LEMMA 7.6. *Games $q+1$ and $q+2$ are computationally indistinguishable.*

PROOF. Suppose there exists an adversary $\mathcal{B}^*$ that distinguishes between the Games $q+1$ and $q+2$. We construct an adversary $\mathcal{B}$ that breaks correlated pseudorandomness property. Adversary $\mathcal{B}$ gets as input $(pk_i, \psi_i)_{i \in [\ell+1]}, pk_{\ell+2}, \psi^*$, where ψ^* is either a valid encoding using key $pk_{\ell+2}$ or a randomly chosen element. Adversary $\mathcal{B}$ uses these values to invoke algorithms $Setup^*, Enc^*, KeyGen^*$, where it uses $\psi_{\ell+2} := \psi^*$. It invokes adversary $\mathcal{B}^*$ likewise. Clearly, if $\psi^* = \text{Encode}(pk_{\ell+2}, s)$, then this experiment corresponds to Game $q+1$. On the other hand, if $\psi^* \xleftarrow{\$} \mathcal{K}$ then this experiment corresponds to Game $q+2$. This concludes the proof of the lemma. $\square$

APPENDIXES

A. OUTSOURCING DECRYPTION

In this section, we show how to modify our main construction of attribute-based encryption to support outsourcing of decryption circuits, similar to Green et al. [2011]. Syntactically, ABE with outsourcing decryption consists of algorithms (Setup, KeyGen, Eval, Enc, Dec), satisfying the same semantics as standard ABE, except as follow.

- (1) We require that the Keygen algorithm on input (msk, C) returns two keys.
 - The evaluation key ek_C , given to a computationally powerful proxy.
 - And a decryption key dk given to the client.
- (2) Given a ciphertext ct_{ind} , the proxy must perform the “bulk” of the computation by running $\text{Eval}(\text{ek}_C, \text{ct}_{\text{ind}}) \rightarrow \text{ct}'_{\text{ind}}$. The client holding the decryption key dk can then compute $\text{Dec}(\text{dk}, \text{ct}'_{\text{ind}}) \rightarrow m$ iff $C(\text{ind}) = 1$.

Formally, we require that the runtime of Dec is $\text{poly}(\lambda, d_{\max})$, where λ is the security parameter and $d_{\max}$ is the maximum circuit depth (in particular, it is independent on the size of C). The other correctness properties are the same as in standard ABE.

We envision a semi-honest proxy server, where it performs all algorithms correctly but tries to learn the message payload. That is, the security ensures that an adversary (that may play proxy’s role) should learn nothing about the message, conditioned on that it queries for decryption keys dk ’s for predicates that are not satisfied by the challenge index (note, the adversary can query for evaluation keys separately for predicates that are satisfied). More formally, for a stateful adversary $\mathcal{A}$ (that can maintain a global state information throughout the execution of the experiment), we consider the following game.

- Phase 1.* The adversary specifies the challenge index ind , gives it to the challenger who runs the Setup algorithm, and returns the master public key mpk to the adversary.
- Phase 2.* The adversary may issue oracle queries of the form C to KeyGen algorithm and specify whether to obtain ek_C only or both (ek_C, dk) . It outputs two challenge messages m_0, m_1 .
- Phase 3.* The challenger runs $\text{ct}_{\text{ind}} \leftarrow \text{Enc}(\text{mpk}, \text{ind}, m_b)$ for a randomly chosen bit b and gives ct_{ind} to the adversary.
- Phase 4.* Same as Phase 2, where eventually the adversary outputs a guess bit b' .

The ABE with outsourcing decryption capabilities is said to be secure if no adversary can guess the bit b with probability significantly better than $1/2$, conditioned on that for all decryption keys dk obtained for predicates C , $C(\text{ind}) = 0$.

We now give a high-level overview of the changes applied to our main construction and then provide a formal construction. As before, the key-generation algorithm assigns two keys for each circuit wire. The evaluation key consists of all the recoding keys for the circuit. In addition, the output wire has another key pk_{out} which now plays a special role. The recoding key from $\text{pk}_{|C|,1}$ to pk_{out} is only given to the client as the decryption key. If $C(\text{ind}) = 1$, the the proxy computes an encoding under the $\text{pk}_{|C|,1}$ and forwards it to the client. The client applies the transformation and decrypts the message. For technical reasons, since we are using two-to-one recoding mechanism, we need to introduce an auxiliary public key pk_{in} and a corresponding encoding.

Setup($1^\lambda, 1^\ell, d_{\max}$). For each of the ℓ input wires, generate two public/secret key pairs. Also, generate an additional public/secret key pair

$$\begin{aligned} (\text{pk}_{i,b}, \text{sk}_{i,b}) &\leftarrow \text{Keygen}(\text{pp}) \quad \text{for } i \in [\ell], b \in \{0, 1\}, \\ (\text{pk}_{\text{out}}, \text{sk}_{\text{out}}) &\leftarrow \text{Keygen}(\text{pp}), \\ (\text{pk}_{\text{in}}, \text{sk}_{\text{in}}) &\leftarrow \text{Keygen}(\text{pp}). \end{aligned}$$

Output

$$\text{mpk} := \begin{pmatrix} \text{pk}_{1,0} & \text{pk}_{2,0} & \dots & \text{pk}_{\ell,0} & \text{pk}_{\text{in}} \\ \text{pk}_{1,1} & \text{pk}_{2,1} & \dots & \text{pk}_{\ell,1} & \text{pk}_{\text{out}} \end{pmatrix} \quad \text{msk} := \begin{pmatrix} \text{sk}_{1,0} & \text{sk}_{2,0} & \dots & \text{sk}_{\ell,0} & \text{sk}_{\text{in}} \\ \text{sk}_{1,1} & \text{sk}_{2,1} & \dots & \text{sk}_{\ell,1} & \text{sk}_{\text{out}} \end{pmatrix}.$$

Enc($\text{mpk}, \text{ind}, m$). For $\text{ind} \in \{0, 1\}^\ell$, choose a uniformly random $s \xleftarrow{\$} \mathcal{S}$ and encode it under the public keys specified by the following index bits.

$$\psi_i \leftarrow \text{Encode}(\text{pk}_{i,\text{ind}_i}, s) \text{ for all } i \in [\ell].$$

Encode s under the input public key.

$$\psi_{\text{in}} \leftarrow \text{Encode}(\text{pk}_{\text{in}}, s).$$

Encrypt the message m

$$\tau \leftarrow \text{E}(\text{Encode}(\text{pk}_{\text{out}}, s), m).$$

Output the ciphertext

$$\text{ct}_{\text{ind}} := (\psi_1, \psi_2, \dots, \psi_\ell, \psi_{\text{in}}, \tau).$$

KeyGen(msk, C).

- (1) For every non-input wire $w = \ell + 1, \dots, |C|$ of the circuit C , and every $b \in \{0, 1\}$, generate public/secret key pairs.

$$(\text{pk}_{w,b}, \text{sk}_{w,b}) \leftarrow \text{Keygen}(\text{pp}).$$

- (2) For the gate $g = (u, v, w)$ with output wire w , compute the four recoding keys $\text{rk}_{b,c}^w$ (for $b, c \in \{0, 1\}$).

$$\text{rk}_{b,c}^w \leftarrow \text{ReKeyGen}(\text{pk}_{u,b}, \text{pk}_{v,c}, \text{sk}_{u,b}, \text{pk}_{w,g_w(b,c)}).$$

- (3) Also, compute the following recoding key.

$$\text{rk}_{\text{out}}^{\text{out}} \leftarrow \text{ReKeyGen}(\text{pk}_{|C|,1}, \text{pk}_{\text{in}}, \text{sk}_{|C|,1}, \text{pk}_{\text{out}}).$$

Output the evaluation key which is a collection of $4(|C| - \ell)$ recoding keys.

$$\text{ek}_C := (\text{rk}_{b,c}^w : w \in [\ell + 1, |C|], b, c \in \{0, 1\}).$$

Output the decryption key $\text{dk} := \text{rk}_{\text{out}}^{\text{out}}$.

Eval($\text{ek}_C, \text{ct}_{\text{ind}}$). We tacitly assume that ct_{ind} contains the index ind . For $w = \ell + 1, \dots, |C|$, let $g = (u, v, w)$ denote the gate with output wire w . Suppose wires u and v carry the values b^* and c^* so that wire w carries the value $d^* := g_w(b^*, c^*)$. Compute

$$\psi_{w,d^*} \leftarrow \text{Recode}(\text{rk}_{b^*,c^*}^w, \psi_{u,b^*}, \psi_{v,c^*}).$$

If $C(\text{ind}) = 1$, then we would have computed $\psi_{|C|,1}$. **Output**

$$\text{ct}'_{\text{ind}} := (\psi_{|C|,1}, \psi_{\text{in}}, \tau).$$

If $C(\text{ind}) = 0$, output $\perp$.

Dec($\text{dk}, \text{ct}'_{\text{ind}}$). Apply the transformation

$$\psi_{\text{out}} \leftarrow \text{Recode}(\text{rk}_{\text{out}}^{\text{out}}, \psi_{\text{in}}, \psi_{|C|,1}).$$

and output the message

$$m \leftarrow D(\psi_{\text{out}}, \tau).$$

Security. We informally state how to modify the simulator in the proof of security in Section 6.4. The simulator gets $\{\text{pk}_i, \psi_i\}_{i \in [\ell+2]}$ from the “outside.” It assigns $\text{pk}_1, \dots, \text{pk}_\ell$ as the public keys specified by the bits of ind and $\text{pk}_{\text{in}} := \text{pk}_{\ell+1}, \text{pk}_{\text{out}} := \text{pk}_{\ell+2}$. It is easy to see how to simulate the ciphertext: all the input encodings become a part of it, as well as an encryption of the message using $\psi_{\text{out}} := \psi_{\ell+2}$. Now, the evaluation key ek is simulated by applying the TOR simulator.

- For query C such that $C(\text{ind}) = 0$, the simulator can choose $(\text{pk}_{|C|,1}, \text{sk}_{|C|,1})$ by itself (the public key $\text{pk}_{|C|,0}$ is “fixed” by the TOR simulator). Hence, the decryption key dk can be computed using $\text{sk}_{|C|,1}$.
- On the other hand, for query C such that $C(\text{ind}) = 1$, the adversary is not allowed to obtain the decryption key dk , hence there is not need to simulate it.

B. EXTENDING SECRET KEYS

Consider the following problem: a users holds two (or more) secret keys sk_{C_1} and sk_{C_2} . C_1 allows decrypting all ciphertexts addressed to *human resources* department and C_2 allows to decrypt ciphertexts addressed to *share holders*. The user wishes to create (and delegate) another secret key sk_{C^*} that allows to decrypt ciphertexts addressed to *human resources* and *share holders*. The question that we study is whether it is possible to allow the user to compute sk_{C^*} without calling the authority holding the master secret key msk .

More formally, given $\{\text{sk}_{C_i}\}_{i \in [q]}$, we require an additional algorithm $\text{Extend}(\{\text{sk}_{C_i}\}_{i \in [q]}, C^*)$ that returns a secret key sk_{C^*} for any circuit C^* that is a black-box monotone composition of C_i ’s. The correctness properties are that of the standard ABE in addition to that sk_{C^*} can be used to decrypt ciphertexts ct_{ind} if $C^*(\text{ind}) = 1$. The security properties also remain similar to standard ABE: no adversary should be able to learn anything from the ciphertext ct_{ind} about the message m given many secret keys sk_C for queries C such that $C(\text{ind}) = 0$, where secret key sk_C may be generated via KeyGen or Extend algorithms. Note that only monotone compositions are realizable, since otherwise, a user holding a secret keys sk_C , where $C(\text{ind}) = 0$, could come up with a secret key for $\bar{C}$ and hence break the security game.

To suppose monotone extensions, it is enough to show how to obtain (1) $\text{sk}_{C_1 \text{ AND } C_2}$ given $\text{sk}_{C_1}, \text{sk}_{C_2}$, and (2) $\text{sk}_{C_1 \text{ OR } C_2}$ given $\text{sk}_{C_1}, \text{sk}_{C_2}$. We start from the construction presented in Section 1 and explain how to modify it.

- First, instead of fixing $\text{pk}_{|C_i|,1} = \text{pk}_{\text{out}}$ in this construction, we sample $(\text{pk}_{|C_i|,1}, \text{sk}_{|C_i|,1}) \leftarrow \text{Keygen}(\text{pp})$ for every query C_i . In the secret key sk_{C_i} , in addition to all recoding keys for the internal gates, we publish

$$\text{rk}_{\text{aux}} \leftarrow \text{ReKeyGen}(\text{pk}_{|C_i|,1}, \text{pk}_{|C_i|,1}, \text{sk}_{|C_i|,1}, \text{pk}_{\text{out}})$$

and also include the secret key $\text{sk}_{|C_i|,1}$. Clearly, this does not affect the correctness, because given an encoding under $\text{pk}_{|C_i|,1}$, the user can compute the encoding under pk_{out} and record the message if $C_i(\text{ind}) = 1$. Moreover, this does not affect the security since our simulation proceeds by sampling $(\text{pk}_{|C_i|,1}, \text{sk}_{|C_i|,1})$ honestly using Keygen algorithm and the fact that the adversary is restricted to queries C_i such that $C_i(\text{ind}) = 0$.

- Now, suppose the user holds sk_{C_1} and sk_{C_2} , let $C^* = C_1 \text{ AND } C_2$ (binary OR-gate is handled analogously). Then, we define $\text{Extend}(\text{sk}_{C_1}, \text{sk}_{C_2}, C^* = C_1 \text{ AND } C_2)$ as follows.

Sample keys associated with the active value (i.e., 1) at the output wire of C^* .

$$(pk_{|C^*|,1}, sk_{|C^*|,1}) \leftarrow \text{Keygen}(pp).$$

Also, sample four recoding keys $rk_{b,c}^{C^*}$ (for $b, c \in \{0, 1\}$):

$$\begin{aligned} (pk_{|C^*|,0}, rk_{0,0}^{C^*}) &\leftarrow \text{SimReKeyGen}(pk_{|C_1|,0}, pk_{|C_2|,0}), \\ rk_{0,1}^{C^*} &\leftarrow \text{ReKeyGen}(pk_{|C_1|,0}, pk_{|C_2|,1}, sk_{|C_2|,1}, pk_{|C^*|,0}), \\ rk_{1,0}^{C^*} &\leftarrow \text{ReKeyGen}(pk_{|C_1|,1}, pk_{|C_2|,0}, sk_{|C_1|,1}, pk_{|C^*|,0}), \\ rk_{1,1}^{C^*} &\leftarrow \text{ReKeyGen}(pk_{|C_1|,1}, pk_{|C_2|,1}, sk_{|C_1|,1}, pk_{\text{out}}). \end{aligned}$$

And an auxiliary recoding key,

$$rk_{\text{aux}} \leftarrow \text{ReKeyGen}(pk_{|C^*|,1}, pk_{|C^*|,1}, sk_{|C^*|,1}, pk_{\text{out}}).$$

The secret key sk_{C^*} consists of all recoding keys for the internal gates of C_1 and C_1 , recoding keys $rk_{b,c}^{C^*}$ for $b, c \in \{0, 1\}$, rk_{aux} and secret key $sk_{|C^*|,1}$.

It is easy to see that the correctness is easy to satisfy as previously. In particular, given encoding under output public key $pk_{|C^*|,1}$ the user can obtain encoding under pk_{aux} , and therefore recover the message if $C^*(\text{ind}) = 1$. Also, the security holds similarly by satisfying the following invariant: for all wires w carrying value b^* when evaluating $C^*(\text{ind})$, the simulator knows the secret key $sk_{w,1-b^*}$. Hence, it knows the secret key $sk_{|C^*|,1}$ for all C^* such that $C^*(\text{ind}) = 0$ and can use it to compute rk_{aux} by running.

$$rk_{\text{aux}} \leftarrow \text{ReKeyGen}(pk_{|C^*|,1}, pk_{|C^*|,1}, sk_{|C^*|,1}, pk_{\text{out}}).$$

The remaining internal recoding keys (and the ciphertext) are sampled by the simulator identically as in the security proof in Section 6.3.

ACKNOWLEDGMENTS

We thank Dan Boneh and Shafi Goldwasser for their helpful comments and insightful conversations.

REFERENCES

- Shweta Agrawal, Dan Boneh, and Xavier Boyen. 2010a. Efficient lattice (H)IBE in the standard model. In *Proceedings of EUROCRYPT*. 553–572.
- Shweta Agrawal, Dan Boneh, and Xavier Boyen. 2010b. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In *Proceedings of CRYPTO*. 98–115.
- Shweta Agrawal, Xavier Boyen, Vinod Vaikuntanathan, Panagiotis Voulgaris, and Hoeteck Wee. 2012. Functional encryption for threshold functions (or, fuzzy IBE) from lattices. In *Proceedings of the 15th International Conference on Practice and Theory in Public Key Cryptography*. 280–297.
- Shweta Agrawal, David Mandell Freeman, and Vinod Vaikuntanathan. 2011. Functional encryption for inner product predicates from learning with errors. In *Proceedings of ASIACRYPT*. 21–40.
- Miklós Ajtai. 1999. Generating hard instances of the short basis problem. In *Proceedings of ICALP*. 1–9.
- Miklós Ajtai, Ravi Kumar, and D. Sivakumar. 2001. A sieve algorithm for the shortest lattice vector problem. In *Proceedings of STOC*. 601–610.
- Adi Akavia, Shafi Goldwasser, and Vinod Vaikuntanathan. 2009. Simultaneous hardcore bits and cryptography against memory attacks. In *Proceedings of 6th Theory of Cryptography Conference (TCC09)*. 474–495.
- Joseph A. Akinyele, Matthew W. Pagano, Matthew D. Green, Christoph U. Lehmann, Zachary N. J. Peterson, and Aviel D. Rubin. 2011. Securing electronic medical records using attribute-based encryption on mobile devices. In *Proceedings of the 1st ACM Workshop on Security and Privacy in Smartphones and Mobile Devices (SPSM'11)*. ACM, New York, NY, 75–86. DOI: <http://dx.doi.org/10.1145/2046614.2046628>
- Prabhanjan Ananth, Zvika Brakerski, Gil Segev, and Vinod Vaikuntanathan. 2014. From selective to adaptive security in functional encryption. Cryptology ePrint Archive: Report 2014/917. (2014).
- Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. 2010. From secrecy to soundness: Efficient verification via secure computation. In *Proceedings of ICALP*. 152–163.

- Giuseppe Ateniese, Kevin Fu, Matthew Green, and Susan Hohenberger. 2006. Improved proxy re-encryption schemes with applications to secure distributed storage. *ACM Trans. Inf. Syst. Secur.* 9, 1 (2006), 1–30.
- Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. 2012. Foundations of garbled circuits. In *Proceedings of the ACM CCS*. 784–796.
- Matt Blaze, Gerrit Bleumer, and Martin Strauss. 1998. Divertible protocols and atomic proxy cryptography. In *Proceedings of EUROCRYPT*. 127–144.
- Dan Boneh and Xavier Boyen. 2004. Efficient selective-ID secure identity-based encryption without random oracles. In *Proceedings of EUROCRYPT*. 223–238.
- Dan Boneh, Giovanni Di Crescenzo, Rafail Ostrovsky, and Giuseppe Persiano. 2004. Public key encryption with keyword search. In *Advances in Cryptology - Proceedings of the International Conference on the Theory and Applications of Cryptographic Techniques (Eurocrypt'04)*. 506–522.
- Dan Boneh and Matthew K. Franklin. 2001. Identity-based encryption from the Weil pairing. In *Proceedings of CRYPTO*. 213–229.
- Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. 2014. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In *Proceedings of EUROCRYPT*.
- Dan Boneh, Amit Sahai, and Brent Waters. 2011. Functional encryption: Definitions and challenges. In *Proceedings of TCC*. 253–273.
- Xavier Boyen. 2013. Attribute-based functional encryption on lattices. In *Proceedings of TCC*. 122–142.
- Zvika Brakerski and Vinod Vaikuntanathan. 2011. Efficient fully homomorphic encryption from (standard) LWE. In *Proceedings of FOCS*. 97–106.
- David Cash, Dennis Hofheinz, and Eike Kiltz. 2009. How to delegate a lattice basis. Cryptology ePrint Archive, Report 2009/351. (2009).
- David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. 2012. Bonsai trees, or how to delegate a lattice basis. *J. Cryptol.* 25, 4 (2012), 601–639.
- Kai-Min Chung, Yael Kalai, and Salil P. Vadhan. 2010. Improved delegation of computation using fully homomorphic encryption. In *Proceedings of CRYPTO*. 483–501.
- Clifford Cocks. 2001. An identity based encryption scheme based on quadratic residues. In *Proceedings of the IMA International Conference*. 360–363.
- Jean-Sbastien Coron, Tancrède Lepoint, and Mehdi Tibouchi. 2013. Practical multilinear maps over the integers. In *Advances in Cryptology - (CRYPTO'13)*, Ran Canetti and Juan A. Garay (Eds.). Lecture Notes in Computer Science, Vol. 8042. Springer, Berlin Heidelberg, 476–493. DOI: http://dx.doi.org/10.1007/978-3-642-40041-4_26
- Sanjam Garg, Craig Gentry, and Shai Halevi. 2012. Candidate multilinear maps from ideal lattices and applications. Cryptology ePrint Archive, Report 2012/610. (2012).
- Sanjam Garg, Craig Gentry, Shai Halevi, Amit Sahai, and Brent Waters. 2013. Attribute-based encryption for circuits from multilinear maps. In *Proceedings of CRYPTO*. 479–499. Also, Cryptology ePrint Archive, Report 2013/128.
- Rosario Gennaro, Craig Gentry, and Bryan Parno. 2010. Non-interactive verifiable computing: Outsourcing computation to untrusted workers. In *Proceedings of CRYPTO*. 465–482.
- Craig Gentry. 2009. Fully homomorphic encryption using ideal lattices. In *Proceedings of STOC*. 169–178.
- Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. 2008. Trapdoors for hard lattices and new cryptographic constructions. In *Proceedings of STOC*. 197–206.
- Shafi Goldwasser, Yael Kalai, Raluca Ada Popa, Vinod Vaikuntanathan, and Nikolai Zeldovich. 2013. Succinct functional encryption and its power: Reusable garbled circuits and beyond. In *Proceedings of STOC*.
- Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. 2008. Delegating computation: Interactive proofs for muggles. In *Proceedings of STOC*. 113–122.
- Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. 2012. Functional encryption with bounded collusions via multi-party computation. In *Proceedings of CRYPTO*. 162–179.
- Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. 2006. Attribute-based encryption for fine-grained access control of encrypted data. In *Proceedings of the ACM CCS*. 89–98.
- Matthew Green, Susan Hohenberger, and Brent Waters. 2011. Outsourcing the decryption of ABE ciphertexts. In *Proceedings of the 20th USENIX Conference on Security (SEC'11)*. USENIX Association. <http://dl.acm.org/citation.cfm?id=2028067.2028101>.
- Susan Hohenberger, Guy N. Rothblum, Abhi Shelat, and Vinod Vaikuntanathan. 2011. Securely obfuscating re-encryption. *J. Cryptol.* 24, 4 (2011), 694–719.

- Jonathan Katz, Amit Sahai, and Brent Waters. 2008. Predicate encryption supporting disjunctions, polynomial equations, and inner products. In *Proceedings of EUROCRYPT*. 146–162.
- Allison B. Lewko, Tatsuaki Okamoto, Amit Sahai, Katsuyuki Takashima, and Brent Waters. 2010. Fully secure functional encryption: Attribute-based encryption and (hierarchical) inner product encryption. In *Proceedings of EUROCRYPT*. 62–91.
- Allison B. Lewko and Brent Waters. 2010. New techniques for dual system encryption and fully secure HIBE with short ciphertexts. In *Proceedings of TCC*. 455–479.
- Allison B. Lewko and Brent Waters. 2012. New proof methods for attribute-based encryption: Achieving full security through selective techniques. In *Proceedings of CRYPTO*. 180–198.
- Ming Li, Shucheng Yu, Yao Zheng, Kui Ren, and Wenjing Lou. 2013. Scalable and secure sharing of personal health records in cloud computing using attribute-based encryption. *IEEE Trans. Parallel Distrib. Syst.* 24, 1 (2013), 131–143. DOI: <http://dx.doi.org/10.1109/TPDS.2012.97>
- Silvio Micali. 2000. Computationally sound proofs. *SIAM J. Comput.* 30, 4 (2000), 1253–1298.
- Daniele Micciancio and Chris Peikert. 2012. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *Proceedings of EUROCRYPT*. 700–718.
- Daniele Micciancio and Panagiotis Voulgaris. 2010. A deterministic single exponential time algorithm for most lattice problems based on voronoi cell computations. In *Proceedings of STOC*. 351–358.
- Tatsuaki Okamoto and Katsuyuki Takashima. 2010. Fully secure functional encryption with general relations from the decisional linear assumption. In *Proceedings of CRYPTO*. 191–208.
- Adam O’Neill. 2010. Definitional issues in functional encryption. Cryptology ePrint Archive, Report 2010/556. (2010).
- Rafail Ostrovsky and William E. Skeith III. 2005. Private streaming in streaming data. In *Advances in Cryptology (CRYPTO’05)*, Victor Shoup (Ed.). Lecture Notes in Computer Science, Vol. 3621. Springer Berlin Heidelberg, 223–240.
- John P. Papanis, Stavros I. Papapanagiotou, Aziz S. Mousas, Georgios V. Lioudakis, Dimitra I. Kaklamani, and Iakovos S. Venieris. 2014. On the use of attribute-based encryption for multimedia content protection over information-centric networks. *Trans. Emerg. Telecommu. Techno.* 25, 4 (2014), 422–435. DOI: <http://dx.doi.org/10.1002/ett.2722>
- Bryan Parno, Mariana Raykova, and Vinod Vaikuntanathan. 2012. How to delegate and verify in public: Verifiable computation from attribute-based encryption. In *Proceeding of TCC*. 422–439.
- Chris Peikert. 2009. Public-key cryptosystems from the worst-case shortest vector problem. In *Proceeding of STOC*. 333–342.
- Oded Regev. 2009. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM* 56, 6 (2009).
- Alon Rosen and Gil Segev. 2010. Chosen-ciphertext security via correlated products. *SIAM J. Comput.* 39, 7 (2010), 3058–3088.
- Amit Sahai and Hakan Seyalioglu. 2010. Worry-free encryption: Functional encryption with public keys. In *Proceeding of the ACM CCS*. 463–472.
- Amit Sahai and Brent Waters. 2005. Fuzzy identity-based encryption. In *Proceeding of EUROCRYPT*. 457–473.
- Adi Shamir. 1984. Identity-based cryptosystems and signature schemes. In *Proceeding of CRYPTO*. 47–53.
- Damien Stehlé and Ron Steinfeld. 2010. Faster fully homomorphic encryption. In *Proceeding of ASIACRYPT*. 377–394.
- Brent Waters. 2009. Dual system encryption: Realizing fully secure IBE and HIBE under simple assumptions. In *Proceeding of CRYPTO*. 619–636.
- Brent Waters. 2012. Functional encryption for regular languages. In *Proceeding of CRYPTO*. 218–235.

Received July 2014; revised April, July 2015; accepted September 2015