

Attribute-Based Functional Encryption on Lattices

Xavier Boyen

Abstract. We introduce a broad lattice manipulation technique for expressive cryptography, and use it to realize functional encryption for access structures from post-quantum hardness assumptions.

Specifically, we build an efficient key-policy attribute-based encryption scheme, and prove its security in the selective sense from learning-with-errors intractability in the standard model.

1 Introduction

Attribute-Based Encryption (ABE) is a very powerful notion of encryption, where ciphertexts are not decipherable according to the ownership of a specific key (as in public-key encryption), or a specific name (as in identity-based encryption), but according to the fulfillment of a functional condition expressed as a predicate that takes multiple attributes as input.

Attribute-based encryption was first coined in a paper by Goyal *et al.* [22], although the idea was already implicit in the Fuzzy IBE of Sahai and Waters [32], which for the first time permitted ciphertexts to be addressed on the basis of a condition that was strictly richer than a mere equality (of keys or identities). Since then, the notion of ABE has blossomed into an entire research program known as Functional Encryption [23,11], whereby rich functions driven by inputs from both the ciphertext and the key attempting to decrypt it, determine whether the message, or some function thereof, can be accessed. As an illustrative example of recent developments in this area, Waters very recently built a functional cryptosystem whose predicates are deterministic finite automata [33].

As impressive as these results may be, almost all of them appear to require the machinery of bilinear maps [27]—which leaves them completely vulnerable to quantum cryptanalysis, by virtue of hinging on the classically hard but quantumly easy Discrete Log problem. (Limited instances of construction from yet other techniques [16,10] do exist, but, with assumptions that hinge on Factoring, they are equally vulnerable to quantum attacks.) With quantum computers rapidly moving from a scientific to an engineering problem, it behooves us to have safe cryptographic alternatives ready before they become a reality—possibly with nary an advance warning. Lattices appear to be our best defense, for not only are they increasingly conjectured to thwart the quantum threat in a fundamental way, they also have a rich mathematical structure that makes them well suited for building “complex” and expressive cryptographic systems.

Lattices have made their apparition in cryptography with the work Ajtai [5], and have since been used to construct a vast variety of primitives, including one-way and collision-resistant hash functions [5,26], signatures [12,25], public-key

encryption [7,30,31], identity-based encryption schemes [21,15,1,2], lossy trapdoor functions [29], and even a couple instances of functional encryption for inner-product [4] and threshold [3] functions. Lattices have also been very instrumental in cracking the long-standing question of realizing fully homomorphic encryption [19,20,14].

Lattices are indeed rapidly emerging as a mathematical platform of choice for building increasingly powerful and efficient cryptographic primitives. In addition to lattice problems being generally conjectured to withstand quantum attacks, the mathematical properties of these objects make them both relatively efficient and flexible to enable the construction of powerful cryptosystems. Research in lattice-based cryptosystems that reduce from the “Learning With Errors” (LWE) hardness assumption has been particularly active, in no small part because the average-case LWE problem is itself reducible [31,28] from a slew of worst-case lattice problems, for a sound foundation.

Despite all of those incentives and successes, the reality is that functional encryption so far remains largely confined to the world of bilinear maps. In recent years, only a handful of such systems have been successfully realized using lattices, such as the already cited constructions of IBE [21,1], HIBE [15,2], IPE [4], and FuzzyIBE [3]. Further advances have remained elusive, despite the “pull” exerted by the faster pace of progress in that other world of bilinear maps. Rather disconcertingly indeed, as attempts are made to translate high-level principles of bilinear-map functional encryption into lattice analogues, serious difficulties tend to crop up in the most unexpected places when one tries to prove security. A pointed example, documented in [3], relates to the unresolved difficulties faced by those authors when trying to build ABE from LWE.

If anything, this brief history of functional encryption from lattices suggests that new ideas are in order for progress, beyond the field’s classic paradigms.

1.1 Main Motivations

“Attribute-Based Encryption using Lattices” is by many authors’ account an important research question, having been posed and left unanswered in an number of recent works including [15,1,4,3]. Perhaps the best evidence of the problem’s popularity is none other than a recent attempt by a large corporation to lay claim on its solution, in an eponymous patent application [17], even though the problem explicitly remained open to this day.¹ Why such eager enthusiasm?

First and foremost, functional encryption in general and ABE in particular are extremely powerful cryptographic constructs that would seem almost incredible—*e.g.*, by the standards of *circa* 2000. FE and ABE primarily give us unprecedented flexibility and expressiveness with which recipients can be designated in a wholesale manner. Not only do there exist direct use cases for such

¹ The US patent application [17] appears to refer to a precursor of the “Fuzzy IBE using Lattices” subsequently published in [3], wherein a superset of the authors explicitly acknowledge that it did not extend to a proper ABE. We further opine on mathematical but not legal grounds that our ABE falls outside of the claims of [17].

power (we refer to the early literature on the subject for examples), but the prospects that it opens for protocol building are highly intriguing.

As already alluded to, such rewards would be for naught if the looming threat of a catastrophic quantum cryptanalysis kept relegating it to where damage would be contained. It would be foolish to believe that because quantum registers have only grown from 5 to 7 qubits during the last decade, that their size could not suddenly become cryptographically devastating during the next one. This is where lattices come into play.

Compounding their conjectured quantum robustness, lattices also have a number of rather unique efficiency and implementation advantages. For instance, while bilinear-map cryptosystems tend to be convenient to work with on paper thanks to the availability of clean abstractions, this view hides a rather complex elliptic-curve machinery that must be securely implemented in any physical implementation. In lattice-based cryptography, the situation is reversed: schemes and proofs tend to be more complex and mirrored in details, but implementations require only small-number arithmetic and basic linear algebra.

Those are the reasons—from quantum peace of mind, to the sheer challenge of solving compelling theory with practical applications—why it is far from wasted effort to “reinvent” Attribute-Based Encryption, not from bilinear maps but from lattices. (And as a bonus, we introduce a new technique whose power likely reaches into FE far beyond mere ABE.)

1.2 Our Contributions

Our main result is the construction of a functional encryption scheme for monotone access structures, also known as (key-policy) attribute-based encryption, and reduce its security from LWE.

We achieve this result by way of a new lattice manipulation framework suited to the handling of complex access policies. Compared to earlier works on lattice-based IBE and FE, our framework has two distinguishing characteristics: the reliance on *ephemeral lattices* for all private-key extractions, and the subsequent application of a *basis splicing* technique which allows a recipient to convert an ephemeral lattice’s basis into a basis for any lattice in a given family, as needed.

We introduce our framework in relation to a number of observations we make in our attempt to shed some light on the difficulties previously faced. This leads us to a (rather informal) discussion of FE with uniform and non-uniform policies, and how the latter appeared hard to tackle based on previous lattice techniques.

Here we focus solely on introducing our framework and building “key-policy” KP-ABE from it. We defer to future work the study of “ciphertext-policy” CP-ABE and even more ambitious FE.

2 Preliminaries

We refer to the Appendix—available in the eprint version of the paper [13]—for background on lattices in cryptography.

2.1 Attribute-Based Encryption

We follow the definition of the ABE functionality as given by Goyal et al. [22], albeit for security we consider the notion of ciphertext privacy which implies both semantic security and recipient anonymity.

Definition 1 (Key-Policy Attribute-Based Encryption). A Key-Policy Attribute-Based Encryption scheme consists of the following four algorithms:

Setup(λ, ℓ) \rightarrow (Pub, Msk): This algorithm is input a security parameter λ and an attribute number ℓ . It outputs a public key Pub and a master key Msk.

Extract(Pub, Msk, Policy) \rightarrow Key: This algorithm takes a public key Pub, a master key Msk, and an access policy Policy. It outputs a decryption key Key.

Encrypt(Pub, Attrib, Msg) \rightarrow Ctx: This algorithm is input a public key Pub, a list of attributes Attrib, and a message bit Msg. It outputs a ciphertext Ctx.

Decrypt(Pub, Key, Ctx) \rightarrow b : This algorithm takes a public key Pub, a decryption key Key, and a ciphertext Ctx. It outputs the bit b if the attributes Attrib used to create Ctx satisfy the policy Policy used in the creation of Key.

Definition 2 (Selective-Model KP-ABE Security). A KP-ABE scheme is ciphertext-private in the selective-attribute model of security if all probabilistic polynomial time (PPT) adversaries have at most a negligible advantage in this game:

Target: The adversary declares the challenge attributes, Attrib^\dagger , that it wishes to be challenged upon.

Setup: The challenger runs the Setup algorithm and gives the public key to the adversary.

Queries: The adversary is allowed to issue adaptive queries for private keys corresponding to policies Policy of its choice, as long as Attrib^\dagger does not satisfy Policy.

Challenge: The adversary signals its readiness to accept a challenge, and proposes a message to encrypt. The challenger encrypts the message for the challenge attributes Attrib^\dagger , and then flips a random coin r . If $r = 1$, the ciphertext is given to the adversary; if $r = 0$, a random element of the ciphertext space is returned.

Queries: This is a continuation of the earlier query phase.

Guess: The adversary outputs a guess r' of r . The advantage of an adversary A in this game is defined as $|\Pr[r' = r] - \frac{1}{2}|$

One also defines an adaptive-attribute version of the above game, where the adversary may defer the choice of target attributes until requesting the challenge.

2.2 Linear Secret Sharing

Definition 3 (LSSS over \mathbb{Z}_q). An LSSS Π over a set of parties \mathcal{P} consists of an “index map” ρ and a “share-generating matrix” $L \in \mathbb{Z}_q^{\ell \times \theta}$ with ℓ rows and θ columns, where ℓ is the number of shares specified by Π , and θ depends on the structure of Π . For all $i = 1, \dots, \ell$, the function ρ maps the i -th row of L to its

corresponding party. The matrix \mathbf{L} maps an input θ -vector $\mathbf{v} = (s, r_2, \dots, r_\theta)$, where $s \in \mathbb{Z}_q$ is the secret to be shared, and $r_2, \dots, r_\theta \in \mathbb{Z}_q$ are random, into an output ℓ -vector $\mathbf{L}\mathbf{v} = (s_1, \dots, s_\ell)$ containing the shares of the secret s according to Π . The share $s_i = (\mathbf{L}\mathbf{v})_i$ is assigned to party $\rho(i)$.

Every LSSS according to the above definition enjoys the linear reconstruction property. This means that if Π is an LSSS for the access structure \mathbb{A} , then the following is true. Let $S \in \mathbb{A}$ be any authorized set, and let $I \subset \{1, 2, \dots, \ell\}$ be defined as $I = \{i : \rho(i) \in S\}$. Then, there exist constants $\{\kappa_i \in \mathbb{Z}_q\}$ for $i \in I$, such that, if the $\{\lambda_i = (\mathbf{L}\mathbf{v})_i\}$ are valid shares of any secret s according to Π , then $\sum_{i \in I} \kappa_i \lambda_i = s$. It was shown by Beimel [9], that these constants $\{\kappa_i\}$ can be found in time polynomial in the size of the share-generating matrix \mathbf{L} .

Vector Secrets and Reconstruction over \mathbb{Z} . For the purpose of this paper, we will need a slightly modified notion of LSSS, where secrets and shares are ℓ -dimensional integer vectors in \mathbb{Z}^ℓ , and share-generating matrices are defined over \mathbb{Z} rather than over \mathbb{Z}_q . This creates a few issues:

1. Since secrets and shares are themselves vectors, the vector \mathbf{v} of all such shares should be viewed as a tensor, and the product $(\mathbf{L} \cdot \mathbf{v})$ interpreted accordingly.
2. There is no notion of uniform share distribution over \mathbb{Z} : a benign issue here.
3. Reconstruction in \mathbb{Z} may require fractional interpolation coefficients $\kappa_i \in \mathbb{Q}$. We alleviate this difficulty by relaxing our notion of reconstruction, allowing the reconstructed vector to be a non-zero multiple of the original vector (which is non-trivial only if the vector has dimension greater than one). Such reconstruction is possible using only integer coefficients $\kappa_i \in \mathbb{Z}$.

Low-Norm Share Generation. We will use the generic construction mechanism described in Appendix G of [24, eprint] to convert a monotone access structure into a deterministic LSSS matrix. For access formulas with AND (\wedge) and OR (\vee) gates only, it has the further advantage to build share-generating matrices $\mathbf{L} \in \{0, \pm 1\}^{\ell \times \theta}$ with ternary elements in $\{0, \pm 1\}$. For such formulas, the (unrelaxed) reconstruction coefficients κ_i will be binary in $\{0, 1\}$ by construction, even when working in \mathbb{Z} , hence already integer and low-norm without further relaxation.

Duplicated Attributes. For ease of exposition, we first restrict our attention to formulas where each attribute appears exactly once. Since ρ is then the identity function, we omit it from the notation altogether—until Section 4.5 and the Example Appendix of [13] where we handle missing and duplicated attributes.

3 Framework

3.1 Functional Encryption from Lattices

The Regev Cryptosystem. Recall that the Regev PKE scheme [31] makes use of an Ajtai lattice [5], defined as $\Lambda_q^\perp(\mathbf{A}) = \{\mathbf{x} : \mathbf{A}\mathbf{x} = \mathbf{0} \pmod{q}\} \subseteq \mathbb{Z}^m$, where

$q \in \mathbb{Z}^+$ and $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ together specify the lattice (though not necessarily in a unique way). In Regev’s PKE scheme, one assumes q fixed and $m > n \log q$. The private key is a vector $\mathbf{d} \in \mathbb{Z}^m$ with low euclidean norm $\|\mathbf{d}\| \ll q \sqrt{m}$. The public key is a pair (\mathbf{A}, \mathbf{u}) such that $\mathbf{A}\mathbf{d} = \mathbf{u} \pmod{q}$. To encrypt a bit $m \in \{0, 1\}$, one selects a random ephemeral vector $\mathbf{s} \in \mathbb{Z}_q^n$, and output a pair (c_0, c_1) , where $c_0 = \mathbf{s}^\top \mathbf{u} + \lfloor q/2 \rfloor m + \nu_0$ and $c_1 = \mathbf{s}^\top \mathbf{A} + \nu_1$, and where the additive terms ν_0 and ν_1 are low-norm independent discrete gaussian noise terms. To decrypt, the private-key holder computes the difference $\Delta = c_0 - c_1 \mathbf{d}$ in \mathbb{Z}_q , and interprets it as “ $m = 1$ ” if (the smallest non-negative representative of the coset) Δ lies in $\{\lfloor q/4 \rfloor, \lfloor 3q/4 \rfloor\}$, and as “ $m = 0$ ” otherwise.

Preimage Sampling. The Regev system has served as a starting point for many “expressive” functional generalizations of public-key cryptography. The key turning point in this generalization has been the development, in [21], of a “preimage sampling” technique that, given \mathbf{A} and \mathbf{u} , allow one to obtain a preimage \mathbf{d} such that $\mathbf{A}\mathbf{d} = \mathbf{u} \pmod{q}$ and such that \mathbf{d} has the same conditional distribution given \mathbf{u} as if it had been sampled first and its image computed from it. What makes the preimage-sampling approach cryptographically interesting, is that in order to sample a preimage of good quality (where the “quality” of a sample is an inverse measure of its norm), it is (conjectured) necessary to possess a good quality or low-norm basis \mathbf{B} for the lattice $\Lambda_q^\perp(\mathbf{A})$. Furthermore, Ajtai’s original result [5] does give us an efficient way to co-generate both a uniformly random matrix \mathbf{A} and an associated short basis \mathbf{B} for the lattice it induces; whereas it is a conjectured hard problem to find even a single short vector “after the fact” for a given random \mathbf{A} . Together, these methods provide an effective way to obtain provably secure trapdoors from lattice hardness assumptions, that have been used in interesting ways to construct increasingly “expressive” functional cryptosystems: IBE [21,1], HIBE [15,2], IPE [4], FuzzyIBE [3], and now ABE.

More Expressive Predicates. The combination of the lattice/basis co-generation algorithm of [5], the basic public-key framework of [31], and the preimage sampling approach of [21], has led to the invention of several functional encryption schemes for various classes of functions, starting with the identity-based encryption scheme in the original paper [21]. A handful of other functional encryption schemes from lattices were later devised, including IBE in the standard model [15,1], hierarchical IBE [15,2], inner-product encryption [4], and fuzzy IBE [3]. At a high level, all of those schemes find their roots in the Regev PKE system, which they generalize in various ways following a common principle. The common principle is to extend Regev so that either or both the matrix \mathbf{A} and/or the syndrome \mathbf{u} depend on the functional decryption criterion, rather than being constant. In IBE, the decryption criterion is a match of identities, so we let \mathbf{A} and/or \mathbf{u} be function of the identity. In IPE and FuzzyIBE, the decryption criterion is an inner product equality or a threshold of equalities, obtained by splitting \mathbf{A} and/or \mathbf{u} into multiple shares \mathbf{A}_i and/or \mathbf{u}_i , each of which depending on one of the attributes of the decryption predicate.

3.2 Complex Policies and Non-uniformity

In our quest to understand what differentiates successes from failures in earlier lattice-based FE construction attempts, we are drawn to observe the emergence of a pattern that we shall attempt to characterize informally (based on inductive rather than deductive reasoning).

Uniform Policies. The “successes” share a crucial simplifying characteristic: all attributes taken as formal arguments in the decryption policy are of equal importance; they play symmetrical roles.

- IBE and HIBE use trivial examples of uniform policies, because the decryption predicate is a mere equality test that treats a full identity string as a single atomic input (of variable length in the case of HIBE), comparing that of the ciphertext with that of the private key.
- IPE uses uniform policies, because none of the multiple attributes taken as inputs to the decryption predicate, plays a different role or is more important than the others. Indeed, the predicate is of the form, “ $\langle \mathbf{k}, \mathbf{c} \rangle = 0 \pmod{q}$?” (where \mathbf{k} and \mathbf{c} are the key’s and the ciphertext’s attribute vectors). Now let us consider a permutation π . If we apply it to the components of \mathbf{k} and also to the components of \mathbf{c} , one obtains the new predicate, “ $\langle \pi(\mathbf{k}), \pi(\mathbf{c}) \rangle = 0 \pmod{q}$?”, which is in fact unchanged and evaluates to the same value.
- FuzzyIBE uses uniform policies by same reasoning. The only difference is that here the predicate is a θ -out-of- ℓ threshold equality test between key and ciphertext attributes.

Non-Uniform Policies. To contrast, consider the following basic ABE decryption predicate: “ $(A_k = A_c) \vee ((B_k = B_c) \wedge (C_k = C_c))$?” It falls within the scope of the ABE model; yet it is non-uniform since the atomic clause that takes attribute A as input, $(A_k = A_c)$, can by itself truthify the entire predicate, whereas neither the clause in B nor in C can do the same. The attributes are not symmetrical, since A carries more weight than either B or C . Per our earlier criterion, some permutations π of the attributes would not leave the predicate invariant.

Leakage from Non-Uniformity. The authors of [3] observe that the difficulty with extending existing lattice techniques into ABE stems from the conjunction of two risk factors: the necessity to prevent short-vector private keys from spilling a full basis; and the propensity of keys with asymmetrical components to do just that.

To be sure, there are examples of earlier “FE successes” that allow full-bases to be used as keys: all the HIBE schemes [15,1,2] fall in that category, since full bases are needed for key delegation. However, we contend that passing out full bases is not damaging in this case, because HIBE policies are trivially uniform, involving only a single attribute, so that either there is a full match or there is no match at all—no need to finesse the power of the decryption key in any way.

The other past “FE success” with multi-vector keys is the FuzzyIBE from [3]. There, a private key is a Regev key randomly secret-shared into a number of vectors function of the threshold—definitely not a full basis which would give

too much power. Such sharing finesse led to an attack when one attempted to extend the scheme to ABE with non-uniform policies, because of discrepancies in the relative importance of the private key components. *E.g.*, a key for $A \vee (B \wedge C)$ would be “heavier” at attribute A . In this situation, an adversary could, by making multiple key queries for related but distinct policies, obtain a collection of short vectors whose “heavy” coordinates together leak enough information to allow the adversary to reconstitute a “rogue” (sub-)basis. The uneven weight of the coordinates made it difficult to randomize the keys to prevent the “heavy” coordinates from leaking, without necessarily drowning the “light” coordinates in noise and render them useless.

3.3 Robust Embedding of Policies

Instead of trying to prevent the reconstitution of rogue bases from private-key vectors (which was the direction of future research envisioned in [3]), we shall make our private keys into full bases outright—albeit, bases of *ephemeral random lattices that vary with every invocation of key extraction*.

Ephemeral Lattices. Making keys from constantly changing, ephemeral lattices seems great for security—but how can such keys be useful for decryption in a Regev-like system, if the lattices used for encryption and key extraction are different? In a nutshell, the ephemeral lattices (or, rather, the Ajtai matrices defining them) will have a known structure, featuring both deterministic and randomized subcomponents. The ephemeral lattice is rather high-dimensional and its structure will encode the private-key policy attributes. The structure will allow the recipient to transform this “useless” random-lattice basis, into a basis for any target lattice, typically of a lower dimension, that belongs in a certain authorized set that corresponds to the policy encoded into the initial structure. Thus, if a private key is valid for a given ciphertext, meaning that the attributes of one satisfy the policy of the other, then the recipient is able to transform it into a basis for the lattice used in the ciphertext construction, and from there decryption à la Regev can proceed. Conversely, if a private key is invalid for a given ciphertext, the encryption lattice will be outside the authorized set, and the private key will be useless to derive a (short) basis for that lattice.

Basis Splicing. We refer as *basis splicing* the internal operations that let the recipient transform the given high-dimensional ephemeral-lattice basis, into a basis for any desired lower-dimensional lattice in the authorized set. In the case of ABE, the structure embedded in the ephemeral lattice will be obtained from an LSSS, and the basis splicing operations will amount to taking linear combinations of the basis vectors. Certain linear combinations will cause all the blinding randomness to vanish, transforming the initial *unknown* ephemeral lattice into a smaller *known* target lattice in the authorized set.

Security versus Functionality. At an intuitive level, the security benefits that we derive from our approach are twofold:

- *Private keys as full bases are more robust than single vectors.* In a system where private keys are mere vectors, there is an incentive to obtain more than one such vector, in a bid to reconstruct a rogue basis. If the key is a full basis, there is nothing to be gained in trying to obtain another, which can be generated from the first.
- *Ephemeral lattices make a very potent blinding and firewalling mechanism.* This is perhaps the most important aspect of the framework we propose: since the key-extraction mechanism involves an independently rerandomized lattice that changes upon each invocation, the private keys are in a very strong sense firewalled from one another and from the master secret.

These two properties should intuitively make it easy to construct a secure system, which should translate into easy-to-construct reductionist simulations.

4 Scheme

4.1 Intuition

Setup. The system setup is very straightforward. To each (binary) attribute Attrib_i named in the system, is associated a random Ajtai matrix A_i and a matching trapdoor B_i such that $A_i B_i = 0$ for small $\|B_i\|$. The matrices A_i form the global public key. The trapdoors B_i form the keying authority’s master key.

In KP-ABE, ciphertexts are created for sets of (binary) attributes, while private keys embed the decryption policies. To make it possible to encrypt for a set of attributes, a natural idea is, for each (binary) attribute in the system, to create an Ajtai matrix A_i and an associated trapdoor T_i . The matrices A_i will form the public key; the trapdoors T_i form the master key.

Encryption. To encrypt for an attribute set $\{\text{Attrib}_i\}$, one creates a matrix F by concatenating the public matrices A_i designated by the Attrib_i , filling the gaps with the zero matrix 0 ; one then uses F as an “encryption matrix” à la Regev.²

Key Extraction. To create a private key for a given decryption policy represented as an LSSS, the key-extraction authority starts by constructing a (high-dimensional) ephemeral matrix $M = [M_{\text{diag}} | M_{\text{lsss}}]$, where M_{diag} is a block-diagonal assembly of all the A_i , and M_{lsss} is a tensor product of the LSSS matrix and a secret ephemeral randomization matrix. Using its knowledge of the master-key bases B_i , the authority creates a short basis W for the lattice $\Lambda_q^\perp(M)$, randomizes it into a structure-less short basis K , and returns K as the private key. Notice that the basis K is that of a fresh random lattice whose defining Ajtai matrix M is not even revealed to the recipient.

² A Regev ciphertext (c_0, c_1) is created in reference to an Ajtai lattice $\Lambda_q^\perp(F)$ defined by a known matrix F . We call the matrix F , the *Regev encryption matrix*. (It is usually denoted A but we use F to emphasize that it is a function of the encryption attributes; we reserve the notation A_i for the constant matrices in the public key.)

Decryption. Given a Regev ciphertext created from some encryption matrix F , the first step is to transform the private key K into a basis T for the lattice $A_q^\perp(F)$, using the basis-splicing technique.

The transformation requires the encryption matrix F to lie in the “span” of the (undisclosed!) ephemeral matrix M , *i.e.*, that there be a linear combination of the rows of M that yields $M \hookrightarrow [F|0]$. By the structure of $M = [M_{\text{diag}}|M_{\text{lss}}]$, it follows that the i -th block-column of F is a multiple of the i -th block of M_{diag} , or, in other words, that F is the concatenation of $g_i A_i$ with computable coefficients g_i . Though K was orthogonal to M , it is not orthogonal to $[F|0]$. We can obtain orthogonality to $[F|0]$ by multiplying each row of K by an integer coefficient $\bar{g}_i \propto 1/g_i \pmod{q}$ inversely proportional modulo q to the coefficient g_i of the corresponding column of $[F|0]$ (taking $\bar{g}_i = 0$ when corresponding to the columns of 0 or those of F associated with a coefficient $g_i = 0$).

The basis K thus transformed is a matrix $[T^\top|0^\top]^\top$ where T has full rank and is orthogonal to F . The final observation is to take $\bar{g}_i = (\prod_{j:g_j \neq 0} g_j)/g_i$. Because those \bar{g}_i are already in \mathbb{Z} , no modular reduction is necessary to ensure that $\bar{g}_i \propto 1/g_i \pmod{q}$. Hence the norm $\|T\|$ remains small when the g_i are binary or small enough. This makes of T a low-norm full-rank set, convertible into a basis suitable as a trapdoor for sampling low-norm vectors in $A_q^\perp(F)$.

We see that, by properly constructing M , it is possible for the recipient to know how its trapdoor K can be transformed into the desired trapdoor T , even though M itself is not revealed. Once the trapdoor T is obtained, it can be used to decrypt the ciphertext, *e.g.*, by finding a short preimage \mathbf{d} of the encryption syndrome \mathbf{u} , *i.e.*, such that $F\mathbf{d} = \mathbf{u} \pmod{q}$, and applying Regev.³

Issues. For this approach to work, it is necessary that the norm of the reconstructed trapdoor T be small in order to apply Regev. The only operation that can cause the norm of T to grow out of hand, is the LSSS-based derivation of T from K . In general, for circuits containing “proper” threshold gates—not just \wedge nor \vee —with large fan-in, the coefficients g_i can become exponentially large, which would overwhelm the noise tolerance of the Regev decryption scheme unless the modulus q is itself chosen to be exponentially large.

The first good news is that, even in the pessimal case, the issue of the LSSS coefficients is somewhat mitigated by the fact that we only perform LSSS reconstruction “half-way”, eschewing full-fledged Lagrange interpolation. Indeed, the worst way in which LSSS coefficients intervene in T is through simple products $\prod_j g_j$ —and not as ratios of products that would further require denominator elimination as, say, in the Fuzzy IBE of [3]. Intuitively, the reason why we do not need to account for—and then eliminate—the common denominator in LSSS reconstruction, is because what needs to be reconstructed is not the secret decryption itself (such as a short pre image or basis), but merely a multiple of the (public) encryption matrix F ; only a multiple is needed because F induces the same Ajtai lattice as all its multiples relatively prime to q .

³ Because the private key is a full basis, it allows the recipient to find a preimage for any syndrome; hence the encryption syndrome \mathbf{u} may change with each ciphertext.

The second and main good news is that, as long as the only gates present are \wedge and \vee , regardless of their size or circuit complexity, the coefficients g_i can be made binary $\in \{0, 1\}$, thereby ensuring that $\|\mathbf{T}\| \leq \|\mathbf{K}\|$. This restriction is not as severe as it looks, as it should be emphasized that circuits of \wedge and \vee gates already capture most cases of practical interest for (monotone) access policies. Until now, it was not known how to realize ABE involving even the simplest non-uniform policies, *e.g.*, involving only one \wedge and one \vee gate.

4.2 Construction

We assume the existence of the following PPT algorithms for certain lattice sampling operations. See the Appendix in [13] for some background, and the rapidly evolving literature for the fastest and tightest instantiations, *e.g.*, [18].

- **TrapGen** for co-sampling a uniform Ajtai lattice and a short basis for it [5,6];
- **SampleGaussian** for discrete Gaussian sampling a point on a given Ajtai lattice;
- **SamplePreimage** for sampling a preimage of a given Ajtai syndrome, with a discrete Gaussian conditional density [21,8].
- **ExtendRight** for extending a trapdoor of an Ajtai matrix \mathbf{A} into a trapdoor of any Ajtai matrix of the form $[\mathbf{A}|\mathbf{Z}]$, as long as \mathbf{A} has full rank [15,1].

Remark. (Black-Box Sampling and Algorithm Parameters)

In the scheme description, we view all of the above sampling algorithms as (commodity, interchangeable) *black boxes*, without concern for their precise parameter requirements. For now, it suffices to know that the available sampling algorithms are both sufficiently fast and sufficiently tight, to make the entire system security reducible from the learning-with-error (LWE) hardness assumption with polynomially bounded parameters, so that is can in turn be further (quantumly [31], or for large moduli classically [28]) reduced from worst-case lattice assumptions.

The KP-ABE scheme consists of four algorithms specified as follows.

kpABE.Setup($1^\lambda, 1^\ell$): Given a security parameter λ , and an attribute bound ℓ :

1. Select a security dimension $n > \Omega(\lambda)$ and a base lattice dimension $m > 2n \log q$, together with a prime modulus $q > 2$. (See the Appendix for the constraints on q in function of the desired tightness α of LWE—the larger the modulus, the weaker the assumption.)
2. Use algorithm **TrapGen**(1^λ) to select, for each $i \in [\ell]$, a uniformly random $n \times m$ -matrix $\mathbf{A}_i \in \mathbb{Z}_q^{n \times m}$ with a full-rank m -vector set $\mathbf{B}_i \subseteq \Lambda_q^\perp(\mathbf{A}_i)$ that satisfies a low-norm condition.
3. Select a uniformly random $n \times m$ -matrix $\mathbf{A}_0 \in \mathbb{Z}_q^{n \times m}$.
4. Select a uniform random n -vector $\mathbf{u} \in \mathbb{Z}_q^n$.
5. Output the public key and master key,

$$\text{Pub} = \left(\{\mathbf{A}_i\}_{i \in [\ell]}, \mathbf{A}_0, \mathbf{u} \right) \quad ; \quad \text{Msk} = \left(\{\mathbf{B}_i\}_{i \in [\ell]} \right)$$

kpABE.Extract(Pub, Msk, Policy): On input a public key denoted Pub, a master key denoted Msk, and an access structure denoted Policy, do:

1. Convert Policy into a (low-norm, and preferably deterministic) Linear Span Program matrix $L \in \mathbb{Z}^{\ell \times (\ell + \theta)}$, assigning the i -th row of L to the binary attribute of index $i \in [\ell]$. The columns $j \in [0, \theta]$ are numbered from 0 to θ , with $\theta \leq \ell$ being a function of Policy. The linear encoding rule we adopt for L is that, for a binary attribute list represented as $\text{Attrib} \in \{0, 1\}^\ell$ or $\text{Attrib} \subseteq [\ell]$, the (monotone) access policy is satisfied iff the rows of L selected by Attrib contain in their span the row-vector $[1, 0, \dots, 0] \in \mathbb{Z}^{1+\theta}$.
2. Select θ ephemeral uniform random $n \times m$ -matrices $Z_j \in \mathbb{Z}_q^{n \times m}$ for $j \in [\theta]$.
3. Construct a “virtual encryption matrix” $M \in \mathbb{Z}_q^{\ell n \times (\ell + 1 + \theta)m}$, consisting of $\ell \times (\ell + 1 + \theta)$ blocks of $n \times m$ -“sub-matrices”, by translating the sharing matrix $L = (l_{i,j})_{i \in [\ell], j \in [1+\theta]}$ as follows,

$$M = \left[\begin{array}{c|c|c} \begin{array}{ccc} \boxed{A_1} & & \\ & \boxed{A_2} & \\ & & \ddots \\ & & & \boxed{A_\ell} \end{array} & \begin{array}{c} l_{1,0} \boxed{A_0} \\ l_{2,0} \boxed{A_0} \\ \vdots \\ l_{\ell,0} \boxed{A_0} \end{array} & \begin{array}{ccc} l_{1,1} \boxed{Z_1} \dots l_{1,\theta} \boxed{Z_\theta} \\ l_{2,1} \boxed{Z_1} \dots l_{2,\theta} \boxed{Z_\theta} \\ \vdots \\ l_{\ell,1} \boxed{Z_1} \dots l_{\ell,\theta} \boxed{Z_\theta} \end{array} \\ \hline \text{Public, constant, from Pub} & \text{From Pub} & \text{Secret, random, ephemerals} \end{array} \right] \text{ mod } q$$

Each row of L maps to a particular attribute according to the map ρ associated with the secret-sharing scheme. In this section, we are assuming for simplicity that each attribute (of index $\#i$) appears exactly once (on the i -th row), making ρ the identity function. This restriction is lifted in Section 4.5, to handle missing and duplicated attributes.

4. Build a “structureless” random trapdoor K for $\Lambda_q^\perp(M)$, thus satisfying $M \cdot K = 0 \pmod{q}$. This can be done using `ExtendRight`, based on the fact that $M = [M_{\text{trapdoor}} | M_{\text{extension}}]$, where $M_{\text{trapdoor}} = \text{Diag}(A_1, \dots, A_\ell)$ has full rank and a trivial trapdoor $\text{Diag}(B_1, \dots, B_\ell)$.

Unless `ExtendRight` is already guaranteed to produce an extended basis W whose vectors are independently and identically distributed, it is necessary to rerandomize it to achieve this condition. Let K be the resulting “structureless” trapdoor for M .

5. A redundant form of the policy-based private key may be output, as,

$$\text{Key} = \left(K, L \right)$$

However, two optimizations can be made:

- (a) If the sharing matrix L is deterministic in Policy, it may be omitted.
- (b) It is not necessary to transmit all of K since the decryptor will only ever need the upper-left quadrant of dimension $(\ell + 1)m \times (\ell + 1)m$, which we denote by $K' \in \mathbb{Z}^{(\ell+1)m \times (\ell+1)m}$.

Hence, the private key for Policy may be given in compressed form, as,

$$\text{Key} = \boxed{K'}$$

kpABE.Encrypt(Pub, Attrib, Msg): On input a public key Pub, an attribute list $\text{Attrib} \subseteq [\ell]$, and a message bit $\text{Msg} \in \{0, 1\}$, do:

1. Assemble an “encryption matrix” $F \in \mathbb{Z}_q^{n \times (\ell+1)m}$, obtained as the concatenation of, for each $i \in [\ell]$, either A_i if $i \in \text{Attrib}$, or 0 if $i \notin \text{Attrib}$, and A_0 , as follows,

$$F = \left[\begin{array}{c|c|c|c} F_1 \doteq & & F_\ell \doteq & F_0 \doteq \\ \hline \boxed{A_1} & \cdots & \boxed{A_\ell} & \boxed{A_0} \\ \hline \text{or } 0 & & \text{or } 0 & \\ \hline \underbrace{\hspace{10em}}_{A_i \text{ included iff } i \in \text{Attrib}} & & & \end{array} \right]$$

2. Select a uniform random n -vector $\mathbf{s} \in \mathbb{Z}_q^n$.
3. Select a low-norm Gaussian noise scalar $\nu_0 \in \mathbb{Z}$ according to some parametric distribution Ψ_α (see Appendix), and compute the scalar,

$$c_0 = \left(\mathbf{s}^\top \cdot \mathbf{u} + \nu_0 + \lfloor \frac{q}{2} \rfloor \cdot \text{Msg} \right) \bmod q$$

4. Select a low-norm Gaussian noise vector $\nu_1 \in \mathbb{Z}^{(\ell+1)m}$ whose components are identically and independently distributed from Ψ_α , and compute the vector,

$$\mathbf{c}_1 = \left(\mathbf{s}^\top \cdot F + \nu_1 \right) \bmod q$$

5. Output the ciphertext,

$$\text{Ctx} = \left(c_0, \mathbf{c}_1 \right)$$

(It is not necessary to transmit the components of \mathbf{c}_1 that contain only added ν_1 -noise, i.e., we only need to transmit the components of \mathbf{c}_1 at coordinates where $F_i \neq 0$.)

kpABE.Decrypt(Pub, Key, Ctx): Given a public key Pub, a policy-based key Key (for known policy Policy), and a ciphertext Ctx (for known attributes Attrib):

1. Find an as-short-as-feasible ℓ -vector $\mathbf{g} \in \mathbb{Z}^\ell$ satisfying the two conditions:

$$\mathbf{g}^\top \cdot \mathbf{L} = [d, 0, \dots, 0] \times [1, 0, \dots, 0] \quad ; \quad \forall i \in [\ell] : (g_i = 0) \vee (i \in \text{Attrib})$$

Namely, one finds a linear combination of the rows of \mathbf{L} that yields some small d -multiple of $[1, 0, \dots, 0]$ with $d \in \mathbb{Z} \setminus \{0\}$, using only rows corresponding to attributes in Attrib. This is possible iff Attrib satisfies Policy.

2. Notionally apply the linear combination \mathbf{g} to the “block-rows” of \mathbf{M} , to transform the “virtual” encryption matrix \mathbf{M} into a “real” encryption matrix \mathbf{M}' that matches the encryption matrix F of the given ciphertext (up to constant factors):

$$\mathbf{M}' = \left[\begin{array}{c|c|c|c|c|c} g_1 \boxed{A_1} & g_2 \boxed{A_2} & \cdots & g_\ell \boxed{A_\ell} & d \cdot \boxed{A_0} & \underbrace{0 \cdot Z_1 \mid \cdots \mid 0 \cdot Z_\theta}_0 \\ \hline \text{or } 0 & \text{or } 0 & & \text{or } 0 & & \end{array} \right] \bmod q$$

This is defined, even though the decryptor does not know the Z_i , for they all cancel out.

3. Let M'' be the matrix containing only the $|\text{Attrib}| + 1$ non-zero “block-columns” of M' as shown above. Let K'' be the matrix obtained by removing from K the matching rows and columns—i.e., rows and columns with the same indices as the columns removed from M' . (Dimension-wise, we obtain $M'' \in \mathbb{Z}_q^{n \times (|\text{Attrib}|+1)m}$ and $K'' \in \mathbb{Z}^{(|\text{Attrib}|+1)m \times (|\text{Attrib}|+1)m}$.) We have $M' \cdot K = 0$; therefore $M'' \cdot K'' = 0$, and K'' is a short basis of $\Lambda_q^\perp(M'')$.
4. Likewise, let F'' be the matrix retaining the $|\text{Attrib}| + 1$ non-zero “block-columns” of F ; and let \mathbf{c}_1'' be the ciphertext vector from which only the matching components of \mathbf{c}_1 remain.
5. We now build a trapdoor for the encryption matrix F , or, rather, its reduced form F'' . Let $\mathbf{1}$ be the $m \times m$ identity matrix, and define the diagonal matrices,

$$G = \begin{bmatrix} g_1 \cdot \boxed{\mathbf{1}} & & & & \\ & \ddots & & & \\ & & g_\ell \cdot \boxed{\mathbf{1}} & & \\ & & & & d \cdot \boxed{\mathbf{1}} \end{bmatrix} ; \quad G'' = \begin{bmatrix} \text{non-zero} \\ \text{diagonal} \\ \text{blocks} \\ \text{of } G \end{bmatrix} \in \mathbb{Z}^{\binom{(|\text{Attrib}|+1)m \times (|\text{Attrib}|+1)m}}{(|\text{Attrib}|+1)m}}$$

Notice $F'' \cdot G'' = M'' \pmod{q}$. Since $M'' \cdot K'' = 0 \pmod{q}$, we have $F'' \cdot G'' \cdot K'' = 0 \pmod{q}$. Compute $T'' = G'' \cdot K''$, whose norm is bounded as $\|T''\| \leq \|G''\| \|K''\| \leq \max\{g_i, d\} \|K\|$. The result T'' is our desired trapdoor for sampling short vectors in $\Lambda_q^\perp(F'')$.

6. Using `SamplePreimage` with trapdoor T'' , find a short solution \mathbf{f}'' of $F'' \cdot \mathbf{f}'' = \mathbf{u} \pmod{q}$.
7. Compute $v = c_0 - (\mathbf{f}'')^\top \cdot \mathbf{c}_1'' \pmod{q}$, and represent its coset as an integer $v \in [-\lfloor \frac{q}{2} \rfloor, \lfloor \frac{q}{2} \rfloor]$.
8. Output the decrypted message bit as,

$$b = \begin{cases} 0 & \text{if } \|v\| \leq \lfloor \frac{q}{4} \rfloor \\ 1 & \text{if } \|v\| \geq \lceil \frac{q}{4} \rceil \end{cases}$$

4.3 Correctness

Theorem 4. *For usual values of the lattice parameters in Regev-like encryption systems, the key-policy attribute-based encryption scheme of the previous section will correctly decrypt authorized ciphertexts with overwhelming probability.*

Proof. To see this, suppose that the “independent” initial bases and short vectors (namely, B_i , Y_i , $\mathbf{e}_{i,j}$, $\mathbf{d}_{i,j}$) are sampled with a suitable Gaussian parameter σ , for instance using the tools from [21,8]. Then, the norm of all “dependent” bases and vectors that are supposed to be short, will be bounded by multiples of σ to which certain “growth coefficients” will have applied. To bound those, we note that the only processes in the whole system that will induce “growth”, are:

- in **Extract**: the randomized invocation of **ExtendRight** to obtain K , which merely multiplies the norm of the master-key trapdoors by a constant factor independent of the data;
- in **Decrypt**: the calculation of the trapdoor T'' from K'' , which as we already noted multiplies the norm of K'' by a factor $\leq \max\{g_i, d\}$ that only depends on the linear-sharing reconstruction vector \mathbf{g} , itself function of the function **Policy** and its inputs **Attrib**.

Bounding $\max\{g_i, d\}$ for access-structure circuits with many gates can be tedious, but we note that $\max\{g_i, d\}$ will be dominated by the presence of large threshold gates. On the contrary, \wedge and \vee gates are essentially harmless, as shown below.

Claim. For a circuit consisting only of \wedge and \vee gates, $\max\{g_i, d\} = 1$.

Proof. There exists a deterministic construction of a linear sharing matrix L that guarantees binary reconstruction coefficients in this case (see Preliminaries). \square

We defer to the full paper the exact quantification of the various norm and noise parameters. Of course, while the growing norm of supposedly short vectors can be compensated by commensurately increasing the modulus q , this is best avoided for efficiency reasons. \square

4.4 Security

Theorem 5. *If there exists a probabilistic polynomial-time algorithm \mathcal{A} with advantage $\epsilon > 0$ in a selective-security key-policy attack against the above scheme, then there exists a probabilistic polynomial-time algorithm \mathcal{B} that decides the $(\mathbb{Z}_q, n, \bar{\Psi}_\alpha)$ -LWE problem with advantage $\epsilon/2$, where $\alpha = O(\text{poly}(n))$.*

Proof. In the LWE problem, the decision algorithm is given access to a sampling oracle, \mathcal{O} , which is either a pseudo-random sampler \mathcal{O}_s with embedded secret $s \in \mathbb{Z}_q^n$, or a truly random sampler $\mathcal{O}_\$. Our decider algorithm \mathcal{B} will simulate an attack environment for, and exploit the prowesses of \mathcal{A} , to decide which oracle it is given. The reduction proceeds as follows.$

Instance. \mathcal{B} requests from \mathcal{O} and obtains $((1 + \ell)m + 1)$ LWE samples that we denote as,

$$\begin{aligned} [(\mathbf{w}_{-1}, v_{-1})] &\in (\mathbb{Z}_q^n \times \mathbb{Z}_q) \\ [(\mathbf{w}_0^1, v_0^1), \dots, (\mathbf{w}_0^m, v_0^m)] &\in (\mathbb{Z}_q^n \times \mathbb{Z}_q)^m \\ [(\mathbf{w}_1^1, v_1^1), \dots, (\mathbf{w}_1^m, v_1^m)] &\in (\mathbb{Z}_q^n \times \mathbb{Z}_q)^m \\ &\vdots \\ [(\mathbf{w}_\ell^1, v_\ell^1), \dots, (\mathbf{w}_\ell^m, v_\ell^m)] &\in (\mathbb{Z}_q^n \times \mathbb{Z}_q)^m \end{aligned}$$

Target. \mathcal{A} announces a target attribute vector, denoted Attrib^\dagger , on which it wishes to be challenged.

Setup. \mathcal{B} constructs the public key **Pub** as follows:

1. The vector $\mathbf{u} \in \mathbb{Z}_q^n$ is constructed from the LWE samples of index -1 : simply set $\mathbf{u} = \mathbf{w}$.
2. The matrix $\mathbf{A}_0 \in \mathbb{Z}_q^{n \times m}$ is built from the LWE samples of index 0: set $\mathbf{A}_0 = [\mathbf{w}_0^1 | \dots | \mathbf{w}_0^m]$.
3. For each $i \in [\ell]$ such that attribute $i \in \text{Attrib}^\dagger$, the matrix \mathbf{A}_i is constructed from the LWE samples of index i in a similar way as above: for $i \in \text{Attrib}^\dagger$, set $\mathbf{A}_i = [\mathbf{w}_i^1 | \dots | \mathbf{w}_i^m]$.
4. For each $i \in [\ell]$ such that attribute $i \notin \text{Attrib}^\dagger$, the matrix \mathbf{A}_i is constructed as in the real scheme using TrapGen , which provides an associated low-norm full-rank matrix \mathbf{B}_i such that $\mathbf{A}_i \cdot \mathbf{B}_i = 0$. (The LWE samples of all indices $i \notin \text{Attrib}^\dagger$ will remain unused.)

The resulting public key Pub is given to \mathcal{A} .

Queries. \mathcal{A} is allowed to make adaptive queries for keys Key for policies Policy that the target attribute list Attrib^\dagger does not satisfy. \mathcal{B} constructs and returns a key Key for each query Policy , as follows.

1. As in the real scheme, derive from Policy a (low-norm) linear sharing matrix $\mathbf{L} \in \mathbb{Z}^{\ell \times (1+\theta)}$.
2. Let $\phi = |\text{Attrib}^\dagger|$. Make \mathbf{L}' from \mathbf{L} , keeping only the rows of index i such that $i \in \text{Attrib}^\dagger$. Make \mathbf{L}'' from \mathbf{L}' by dropping the leftmost column of index $j = 0$ (keeping $j = 1, \dots, \theta$).
3. W.l.o.g., suppose that $\text{Attrib}^\dagger = \{i_1, i_2, \dots, i_\phi\} = \{1, 2, \dots, \phi\}$; i.e., the first ϕ attributes, from 1 to ϕ , are arbitrarily assumed to be the attacker's targets.
4. W.l.o.g., suppose that the ϕ left-most columns of \mathbf{L}'' form a ϕ -dimensional square matrix of full rank. The columns of \mathbf{L} from which \mathbf{L}'' is derived can always be reordered to achieve this, since the order of its columns (other than that of index $j = 0$) is arbitrary. Notice that this step requires that the challenge Attrib^\dagger do *not* satisfy the query Policy . If it did, by definition some non-zero $[d, 0, \dots, 0]^\top$ would be in the span of \mathbf{L} , and thus $[0, \dots, 0]^\top$ non-trivially in that of \mathbf{L}'' ; therefore the ϕ left-most columns of \mathbf{L}'' would not be full-rank.
5. Invoking TrapGen , sample ϕ random matrices $\mathbf{Z}_i \in \mathbb{Z}_q^{n \times m}$ with short bases $\mathbf{Y}_i \in \mathbb{Z}^{m \times m}$, for all $i \in \text{Attrib}^\dagger$ (i.e., w.l.o.g., $i = 1, \dots, \phi$ are the indices of the \mathbf{Z}_i with trapdoor \mathbf{Y}_i).
6. Build a “virtual encryption matrix” \mathbf{M} exactly as in the real scheme (see below about the boxes), as,

$$\mathbf{M} = \left[\begin{array}{ccc|ccc|cccc}
 \mathbf{A}_1 & & & l_{1,0} \mathbf{A}_0 & & l_{1,1} \mathbf{Z}_1 & \dots & l_{1,\phi} \mathbf{Z}_\phi & \dots & l_{1,\theta} \mathbf{Z}_\theta \\
 & \ddots & & \vdots & & \vdots & & \vdots & & \vdots \\
 & & \mathbf{A}_\phi & l_{\phi,0} \mathbf{A}_0 & & \boxed{l_{\phi,1} \mathbf{Z}_1 \dots l_{\phi,\phi} \mathbf{Z}_\phi} & \dots & l_{\phi,\theta} \mathbf{Z}_\theta & & \\
 & & & \vdots & & \vdots & & \vdots & & \\
 & & \boxed{\phantom{\mathbf{A}_\phi}} & \vdots & & \vdots & & \vdots & & \\
 & & & l_{\ell,0} \mathbf{A}_0 & & l_{\ell,1} \mathbf{Z}_1 & \dots & l_{\ell,\phi} \mathbf{Z}_\phi & \dots & l_{\ell,\theta} \mathbf{Z}_\theta \\
 & & & & & & & & &
 \end{array} \right] \bmod q$$

7. Denote by \mathbf{Z} the $(\phi n \times \phi m)$ -submatrix of \mathbf{M} made of the blocks $l_{j,i} \mathbf{Z}_i$ whose $i, j \in [\phi]$. Per Lemma 6, we can build (from the \mathbf{Y}_i) a single trapdoor \mathbf{Y} for \mathbf{Z} as a whole.

Lemma 6. For $i = 1, \dots, \phi$, let $Z_i \in \mathbb{Z}_q^{n \times m}$ and $Y_i \in \mathbb{Z}^{m \times m}$ such that $Z_i Y_i = 0 \pmod{q}$. Suppose also that each Y_i is a basis of $\Lambda_q^\perp(Z_i)$ and has low norm $\|Y_i\| \leq \beta \in \mathbb{R}$. Define,

$$Z = \begin{bmatrix} l_{1,1} Z_1 & \cdots & l_{1,\phi} Z_\phi \\ \vdots & \ddots & \vdots \\ l_{\phi,1} Z_1 & \cdots & l_{\phi,\phi} Z_\phi \end{bmatrix} \pmod{q}$$

Then, for any full-rank integer matrix $(l_{i,j})$ with $i, j \in [\phi]$, the Ajtai lattice induced by $Z \in \mathbb{Z}_q^{\phi n \times \phi m}$ admits an efficiently computable (in fact constant) trapdoor $Y \in \mathbb{Z}^{\phi m \times \phi m}$ i.e., such that Y is a basis of $\Lambda_q^\perp(Z)$ with bounded norm $\|Y\| \leq \beta$.

Proof. Take,

$$Y = \begin{bmatrix} Y_1 & & 0 \\ & \ddots & \\ 0 & & Y_\phi \end{bmatrix}$$

We have that $Z \cdot Y = 0 \pmod{q}$, that Y is a basis for $\Lambda_q^\perp(Z)$, and that $\|Y\| \leq \max_i \|Y_i\|$. \square

8. Observe that we now have a trapdoor for every lattice defined by a submatrix of M encased in one of the boxes shown in Step 6. Let us notionally reorder the columns of M by swapping the ϕ left-most A_i -block-columns with the ϕ left-most Z_i -block-columns. We get a matrix $M' = [M'_{\text{trapdoor}} | M'_{\text{extension}}]$, where M'_{trapdoor} is full-rank, block-diagonal, and each of its blocks has an associated trapdoor. We can thus trivially build a trapdoor for all of M'_{trapdoor} . By invoking `ExtendRight`, we extend this into a trapdoor W' for all of M' . Reordering the rows of W' yields a trapdoor for the original M above: call it W .
9. Randomize W into a structure-less basis K whose norm matches that of the real scheme. (This step is only necessary if `ExtendRight` does not already produce a basis whose vectors all have the target discrete Gaussian distribution already; if they do, let $K = W$.)

This concludes the simulation of the private-key extraction. The adversary \mathcal{A} is given the resulting $\text{Key} = (K, L)$. Notice that it has exactly the same distribution as in the real scheme.

Challenge. \mathcal{A} signals that it is ready to accept a challenge, and chooses a message bit $\text{Msg}^\dagger \in \{0, 1\}$. \mathcal{B} responds with a ciphertext $\text{Ctx}^\dagger = (c_0^\dagger, c_1^\dagger)$ assembled from the LWE instance, as follows:

1. Let $c_0^\dagger = v_{-1} + \lfloor \frac{q}{2} \rfloor \cdot \text{Msg}^\dagger$.
2. Let $c_1^\dagger = \left[\underbrace{v_1^1, \dots, v_1^m}_{\text{if } 1 \in \text{Attrib}^\dagger}, \dots, \underbrace{v_\ell^1, \dots, v_\ell^m}_{\text{if } \ell \in \text{Attrib}^\dagger}, \underbrace{v_0^1, \dots, v_0^m}_{\text{always}} \right]$

Observe that when the v_i come from a genuine LWE oracle, the foregoing is a well-formed Regev-like encryption of Msg^\dagger for the encryption matrix F

indicated by the challenge Attrib^\dagger . On the contrary, when the v_i come from a random fake LWE oracle, the ciphertext is independent of the message bit since c_0^\dagger in particular is uniformly and independently distributed.

Continuation. \mathcal{A} is allowed to continue making further private-key extraction queries, after having obtained the challenge ciphertext.

Decision. \mathcal{A} eventually emits a guess, whether Ctx^\dagger was actually a valid encryption of $\text{Msg} \in \{0, 1\}$ as requested. \mathcal{B} uses the guess to decide whether the LWE oracle \mathcal{O} was genuine. If \mathcal{A} says “valid”, then \mathcal{B} says “*genuine*”; if \mathcal{A} says “invalid”, then \mathcal{B} says “*fake*”.

If the adversary succeeds in guessing Msg^\dagger with probability at least $\frac{1}{2} + \epsilon$, then our decision algorithm \mathcal{B} will correctly guess the nature of the LWE oracle with probability at least $\frac{1}{2} + \frac{\epsilon}{2}$. This concludes the proof of the security reduction. \square

4.5 Extensions

So far we have assumed, merely for simplicity of notation, that policies will only encode monotone access structures given as formulas where each attribute appears as argument exactly once. We now show how to list such limitations.

Duplicated Attributes. Arbitrary monotone policies will generally be expressed as formulas where various attributes appear zero, once, or even multiple times. Accordingly, we show how to handle policies that can comport arbitrarily many \wedge and \vee gates, and an arbitrary wiring of the attribute inputs to feed them, including duplication.⁴ The idea is very simple:

kpABE.Setup' is unchanged from the original version: to each attribute one continues to associate one Ajtai matrix A_i and its trapdoor B_i .

kpABE.Setup' also remains the same: the ciphertext is constructed as before, around a Regev encryption matrix F that either includes or excludes each submatrix A_i depending on whether or not the respective attribute $i \in \text{Attrib}$.

kpABE.Extract' must be modified to allow for duplicate occurrences of the same attribute in the Boolean expression of Policy . This is done as follows:

1. Give each occurrence of some attribute $\#i$ in Policy a unique label, say $\#i.1$ and $\#i.2$, and accordingly rewrite the policy Policy into Policy' as a function of the augmented attributes. Policy' has the same topology (structure and size) as Policy , but its input literals are now unique. Keep track of the mapping from the augmented attributes i' to the original attributes i by means of a surjective map $\rho : i' \mapsto i$.
2. Construct the sharing matrix L in the regular way from the augmented-attribute formula Policy' . For each original attribute $\#i$, there will be as many rows in L as the number of occurrences of $\#i$ in the original Policy .

⁴ We must however continue to caution on the use of t -out-of- n threshold gates \geq_t , because unless $t = 1$ or $t = n$ we cannot guarantee in general that the LSSS matrix L and the reconstruction coefficients will be small. Fortunately, as long as repeated attribute inputs are allowed, every possible monotone access structure can be expressed using only \vee and \wedge gates, in such a way that L is a binary or ternary matrix.

3. Construct the “virtual encryption matrix” M from L as before. Since the augmented attributes that emanate from the same original attribute, all refer to the same public matrix A_i , the key-extraction matrix M will thus contain multiple copies of A_i , albeit on different columns.

Once M has been constructed with possibly duplicated A_i on its left-side block-diagonal, key extraction both in the real scheme and in the simulation will proceed as usual. The only effect of the duplication is that, in the simulation, knowledge of trapdoors B_i will be linked to the presence of the original attributes—not the augmented ones—in Attrib^\dagger .

kpABE.Decrypt' requires a small adjustment to cope with duplicated attributes in the Policy encoded in the decryption key. Essentially, before applying the decryption algorithm, the decryptor needs to avail himself as many copies of the attribute as he will need. This is done by duplicating the various fragments of \mathbf{c}_1 that correspond to the attributes that need to be duplicated, before using the result in the normal decryption process.

This construction is very efficient as the ciphertext size remains unchanged in $|\text{Attrib}|$, and the private key size has the same dependency on $|\text{Policy}|$ as it did without attribute duplication (of course, $|\text{Policy}|$ can now grow arbitrarily).

5 Conclusion

In this paper, we have introduced a new cryptographic framework for performing complex lattice basis manipulations, of the kind that seemingly can unlock the construction of very powerful and expressive cryptosystems such as functional encryption. We demonstrated its power and flexibility by building the first known attribute-based cryptosystem from “learning with errors”, a (conjectured) quantum-resistant hardness assumption tied to many lattice problems.

Acknowledgments. The author would like to thank Dan Boneh for suggesting a simplification of the scheme and its proof by way of the `ExtendRight` abstraction, and to thank the TCC 2013 program committee for what appears to be a very thorough review.

References

1. Agrawal, S., Boneh, D., Boyen, X.: Efficient Lattice (H)IBE in the Standard Model. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 553–572. Springer, Heidelberg (2010)
2. Agrawal, S., Boneh, D., Boyen, X.: Lattice Basis Delegation in Fixed Dimension and Shorter-Ciphertext Hierarchical IBE. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 98–115. Springer, Heidelberg (2010)
3. Agrawal, S., Boyen, X., Vaikuntanathan, V., Voulgaris, P., Wee, H.: Functional Encryption for Threshold Functions (or Fuzzy IBE) from Lattices. In: Fischlin, M., Buchmann, J., Manulis, M. (eds.) PKC 2012. LNCS, vol. 7293, pp. 280–297. Springer, Heidelberg (2012)

4. Agrawal, S., Freeman, D.M., Vaikuntanathan, V.: Functional Encryption for Inner Product Predicates from Learning with Errors. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 21–40. Springer, Heidelberg (2011)
5. Ajtai, M.: Generating hard instances of lattice problems (extended abstract). In: STOC 1996 (1996)
6. Ajtai, M.: Generating Hard Instances of the Short Basis Problem. In: Wiedermann, J., Van Emde Boas, P., Nielsen, M. (eds.) ICALP 1999. LNCS, vol. 1644, pp. 1–9. Springer, Heidelberg (1999)
7. Ajtai, M., Dwork, C.: A public-key cryptosystem with worst-case/average-case equivalence. In: STOC 1997 (1997)
8. Alwen, J., Peikert, C.: Generating shorter bases for hard random lattices. In: STACS 2009 (2009)
9. Beimel, A.: Secure schemes for secret sharing and key distribution. PhD thesis, Department of Computer Science, Technion (1996)
10. Boneh, D., Gentry, C., Hamburg, M.: Space-efficient identity based encryption without pairings. In: FOCS 2007 (2007)
11. Boneh, D., Sahai, A., Waters, B.: Functional Encryption: Definitions and Challenges. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 253–273. Springer, Heidelberg (2011)
12. Boyen, X.: Lattice Mixing and Vanishing Trapdoors: A Framework for Fully Secure Short Signatures and More. In: Nguyen, P.Q., Pointcheval, D. (eds.) PKC 2010. LNCS, vol. 6056, pp. 499–517. Springer, Heidelberg (2010)
13. Boyen, X.: Attribute-based functional encryption on lattices. Cryptology ePrint Archive, Report 2012/??? (December 21, 2012), <http://eprint.iacr.org/>
14. Brakerski, Z., Vaikuntanathan, V.: Fully Homomorphic Encryption from Ring-LWE and Security for Key Dependent Messages. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 505–524. Springer, Heidelberg (2011)
15. Cash, D., Hofheinz, D., Kiltz, E., Peikert, C.: Bonsai Trees, or How to Delegate a Lattice Basis. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 523–552. Springer, Heidelberg (2010)
16. Cocks, C.: An Identity Based Encryption Scheme Based on Quadratic Residues. In: Honary, B. (ed.) Cryptography and Coding 2001. LNCS, vol. 2260, pp. 360–363. Springer, Heidelberg (2001)
17. Microsoft Corporation. Attribute based encryption using lattices. Application USPTO 20120155635 (December 17, 2010)
18. Ducas, L., Nguyen, P.Q.: Faster Gaussian Lattice Sampling Using Lazy Floating-Point Arithmetic. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 415–432. Springer, Heidelberg (2012)
19. Gentry, C.: Fully homomorphic encryption using ideal lattices. In: STOC 2009 (2009)
20. Gentry, C.: Toward Basing Fully Homomorphic Encryption on Worst-Case Hardness. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 116–137. Springer, Heidelberg (2010)
21. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: STOC 2008 (2008)
22. Goyal, V., Pandey, O., Sahai, A., Waters, B.: Attribute-based encryption for fine-grained access control of encrypted data. In: CCS 2006 (2006)
23. Katz, J., Sahai, A., Waters, B.: Predicate Encryption Supporting Disjunctions, Polynomial Equations, and Inner Products. In: Smart, N.P. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 146–162. Springer, Heidelberg (2008)

24. Lewko, A., Waters, B.: Decentralizing attribute-based encryption. Cryptology ePrint Archive, Report 2010/351 (2010), <http://eprint.iacr.org/>
25. Lyubashevsky, V.: Lattice Signatures without Trapdoors. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 738–755. Springer, Heidelberg (2012)
26. Micciancio, D.: Generalized compact knapsacks, cyclic lattices, and efficient one-way functions from worst-case complexity assumptions. In: FOCS 2002 (2002)
27. Miller, V.: The Weil pairing, and its efficient calculation. J. Cryptology (2004)
28. Peikert, C.: Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In: STOC 2009 (2009)
29. Peikert, C., Waters, B.: Lossy trapdoor functions and their applications. SIAM J. Computing (2011)
30. Regev, O.: New lattice-based cryptographic constructions. J. ACM (2004)
31. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: STOC 2005 (2005)
32. Sahai, A., Waters, B.: Fuzzy Identity-Based Encryption. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 457–473. Springer, Heidelberg (2005)
33. Waters, B.: Functional Encryption for Regular Languages. In: Safavi-Naini, R. (ed.) CRYPTO 2012. LNCS, vol. 7417, pp. 218–235. Springer, Heidelberg (2012)