Lattice-Inspired Broadcast Encryption and Succinct Ciphertext-Policy ABE

Zvika Brakerski* Vinod Vaikuntanathan†

Abstract

We propose a candidate ciphertext-policy attribute-based encryption (CP-ABE) scheme for circuits, where the ciphertext size depends only on the depth of the policy circuit (and not its size). This, in particular, gives us a Broadcast Encryption (BE) scheme where the size of the keys and ciphertexts have a poly-logarithmic dependence on the number of users. This goal was previously only known to be achievable assuming ideal multilinear maps (Boneh, Waters and Zhandry, Crypto 2014) or indistinguishability obfuscation (Boneh and Zhandry, Crypto 2014) and in a concurrent work from generic bilinear groups and the learning with errors (LWE) assumption (Agrawal and Yamada, Eurocrypt 2020).

Our construction relies on techniques from lattice-based (and in particular LWE-based) cryptography. We analyze some attempts at cryptanalysis, but we are unable to provide a security proof.

1 Introduction

Broadcast Encryption (BE) [FN93] is an important multi-user generalization of public-key encryption where a *broadcaster* can send the same message m to an arbitrary subset $S \subseteq \mathcal{U}$, where $\mathcal{U}$ is the universe of all the N possible users. A trivial, communication-inefficient, way of achieving this would involve the broadcaster encrypting m separately with the public keys of all users in S, resulting in a ciphertext of size $O(|S|)$ (ignoring dependence on the security parameter λ). Broadcast encryption seeks to achieve the same end goal with much better parameters, ideally ciphertexts and keys of size $O(\lambda)$ (ignoring polylogarithmic factors in $|S|$ and $|\mathcal{U}|$).[1]

The first solution to the broadcast encryption problem was proposed by Boneh, Gentry and Waters [BGW05] using bilinear maps on elliptic curves.[2] Their construction had ciphertexts of size

*Weizmann Institute of Science. Supported by the Binational Science Foundation (Grant No. 2016726), and by the European Union Horizon 2020 Research and Innovation Program via ERC Project REACT (Grant 756482) and via Project PROMETHEUS (Grant 780701).

†MIT. Research supported in part by NSF Grants CNS-1350619 and NSF/BSF grant CNS-1718161, the US/Israel Binational Science Foundation and by the Defense Advanced Research Projects Agency (DARPA) and the U.S. Army Research Office under contracts W911NF-15-C-0226 and W911NF-15-C-0236.

[1] We note that in the context of BE, it is assumed that the decryptor knows the set S, either out-of-band, or as an addendum to the ciphertext. In any case, the description of S is not counted towards the ciphertext size or the communication complexity.

[2] Prior to [BGW05], solutions were either for the case of small sets S, having $O(|S|)$ size ciphertexts, or a line of work (starting from [NNL01]) that handled very large sets, that is sets of size $N - r$ with $O(r)$ size ciphertexts. Both solutions degrade to $O(N)$ ciphertext size for "typical" size broadcast sets.

$O(1)$ and user secret keys of size $O(1)$, but public keys of size $O(N)$.[3] Subsequently, there were constructions from indistinguishability obfuscation (iO) [BZ17] and from multilinear maps [BWZ14], the latter of which achieved polylog(N) size for ciphertexts, secret keys and the public key. Very recently, in an independent work, Agrawal and Yamada [AY20] showed a construction with the same dependence on N in the generic group model for groups with bilinear maps, assuming in addition that the learning with errors (LWE) assumption holds (For a more detailed comparison, see Section 1.4.)

Broadcast encryption is a fundamental cryptographic primitive, and one of the few advanced cryptographic primitives that have been actually commercially used and recognized as such [FN16]. It is also one of the last standing "natural" primitives for which no post-quantum solution is known, despite much effort.[4] This is the case even if we require *only that the ciphertext* (but not necessarily the keys) *have constant size*, a goal that we know how to achieve from bilinear maps [BGW05]. Indeed, achieving even sublinear, that is $o(N)$, ciphertext size from LWE, regardless of the length of the keys, is a wide open question. This leads us to ask:

Can we construct non-trivial broadcast encryption systems from LWE?

While we do not provide a complete answer in this paper, we propose a novel "LWE-inspired" construction based on a natural extension of known learning with errors structures. We hope that both the construction and the assumption will prove useful, both in constructing truly LWE-based broadcast encryption and in other LWE-based constructions of advanced cryptographic objects. In fact, our construction solves the more general problem of succinct ciphertext-policy attribute-based encryption (CP-ABE) which we describe next.

A More General Problem: Succinct CP-ABE. It is possible to cast the BE problem as a special case of another cryptographic task which is by itself very interesting, namely that of constructing *succinct ciphertext-policy* attribute-based encryption schemes, as explained below.

In a ciphertext-policy attribute-based encryption scheme (CP-ABE), messages are encrypted with respect to access policies, specified as Boolean circuits f, and secret keys SK_x are generated for users with certain attributes, specified as strings x. The key SK_x is able to decrypt a ciphertext CT if and only if the attributes (encoded in the key) satisfy the policy circuit (encoded in the ciphertext), that is if and only if $f(x) = 1$. A key-policy ABE (KP-ABE) scheme is the dual notion, where the role of the inputs and circuits is reversed. That is, keys SK_f are bound to circuits f, and the ciphertexts encode attributes. We know by now how to construct *key-policy* ABE schemes from the LWE assumption [GVW13, BGG+14]. The ciphertexts in these schemes have size $|x| \cdot \mathrm{poly}(\lambda, d)$ where d is an upper bound on the depth of functions f supported by the scheme. The keys have size $|f| \cdot \mathsf{poly}(\lambda, |x|, d)$ in [GVW13] and $\mathsf{poly}(\lambda, |x|, d)$ in [BGG+14] independent of $|f|$. Looking ahead, we remark that the succinct keys feature of [BGG+14] will turn out to be crucial for our purposes.

To connect this to broadcast encryption, consider a function f_S that encodes membership in a set S. That is, the Boolean circuit f_S takes as input $x \in \{0,1\}^{\log N}$, and outputs 1 *if and only if* $x \in S$. A CP-ABE scheme where the ciphertext encodes f_S and the secret keys encode user identities x will syntactically give us a broadcast encryption scheme. In fact, using the [BGG+14]

[3] Here and in much of the overview, we ignore polynomial dependence on the security parameter.

[4] With the exception of obfuscation-based solutions, which are very involved and removed from concrete efficiency.

scheme, the keys will have quasi-optimal size $\mathsf{poly}(\lambda, \log N, d) = \mathsf{poly}(\lambda, \log N)$ since f_S can be implemented as a $O(\log N)$ depth circuit.

To satisfy the ciphertext-succinctness property of broadcast encryption, we need a *succinct CP-ABE* scheme. By this, we mean that the ciphertext size should not depend only polylogarithmically on the size of the policy circuit f_S. (Here, as in the case of broadcast encryption, the decryptor is given the policy circuit out-of-band, and it is not counted against the ciphertext size.)

A naïve first attempt at constructing a CP-ABE scheme is to use the known LWE-based KP-ABE schemes [GVW13,BGG+14] and flip the roles of the policy and the attributes using a universal circuit. That is, to generate a CP-ABE key for x, generate a KP-ABE key for the circuit $\mathcal{U}_x(\cdot)$ which takes the description of a function f as input and outputs $f(x)$. The CP-ABE ciphertext w.r.t. a policy f is then the KP-ABE ciphertext w.r.t. the description f. The key difficulty is that the [GVW13, BGG+14] schemes are *not ciphertext-succinct*. Indeed, the ciphertexts have size $|f| \cdot \mathrm{poly}(\lambda, d)$ where d is the depth of the universal circuit $\mathcal{U}_x$. This linear dependence on $|f|$ brings us back to square one.

1.1 Our Results: LWE-Inspired BE and CP-ABE Candidates

In this work, we present a candidate construction of a succinct CP-ABE scheme, and as a corollary, we obtain a quasi-optimal BE scheme. Our construction is based on a heuristic that allows to "invert" the key-succinctness of the BGG+ KP-ABE scheme. We do not have a security reduction for this heuristic, and we pose its security as an open problem.

A Historical Note. We have been circulating this candidate since 2015 and sharing it with researchers in the field, in attempt to find a proof or an attack. Some approaches for attacks were proposed by us and by others, but it so far appears that the scheme withstands those attempts. The scheme in this paper is identical to the scheme that we circulated in 2015, and we did not need to make any adjustments in order to maintain (heuristic) security. This somewhat boosts our confidence in the security of the scheme, but we were nevertheless unable to come up with a proof, or even to come up with a closed-form assumption that implies the security of our candidate. We therefore believe that at this point it may be justified to publish our candidate and utilize the collective wisdom of the cryptographic community to find a proof or an attack. We outline some of the main cryptanalysis attempts in Section 1.3 and Section 5 and explain the reasons why they do not seem to apply. A technical overview of our construction is provided in Section 1.2 below.

1.2 Technical Overview

As explained above, our scheme starts from the KP-ABE scheme of [BGG+14], which has succinct keys and non-succinct ciphertexts, and attempts to turn it on its head, to achieve CP-ABE with succinct ciphertexts (and keys). We use LWE-inspired techniques in order to (heuristically) guarantee collusion resistance. Details follow.

The BGG+ Key Policy ABE Scheme. The high-level idea of the BGG+ scheme is a method of encrypting the attribute x in a way that allows for a specific type of homomorphic evaluation. In particular, the encryption algorithm takes as input a master public key $(\mathbf{A}_0, \mathbf{A}_1, \ldots, \mathbf{A}_k)$ where each matrix lives in $\mathbb{Z}_q^{n \times m}$, and attributes $x \in \{0,1\}^k$ and encrypts it as

$$\left(\mathbf{s}\mathbf{A}_0, \mathbf{s}(\mathbf{A}_1 + x_1\mathbf{G}), \ldots, \mathbf{s}(\mathbf{A}_k + x_k\mathbf{G})\right) + \mathbf{e}$$

where $\mathbf{s} \in \mathbb{Z}_q^{1\times n}$ is a uniformly random row vector ("the LWE secret") and $\mathbf{G}$ is the "gadget matrix" from [MP12]. The vector $\mathbf{e}$ is the "LWE noise vector" which contains short entries (Gaussians, almost always); we will omit the noise vector henceforth and take care to only multiply the ciphertext by small values so that the presence of noise does not significantly change the output. We also ignore the way the encrypted message is encoded in the ciphertext. For the purpose of this exposition one can simply think about using the above to encrypt 0, and using a completely random ciphertext for 1, so that the LWE assumption guarantees that the two are indistinguishable to an adversary, but given a proper key they can be distinguished. We note that more elegant encoding methods exist.

Before describing the decryption process, we note an important *decomposability* feature of this encryption procedure that was already utilized in [BV16] (in fact, [BV16] and this work are concurrent) and recently also in [AY20]. The encryption process can be split into an "offline step" that is computed before the attribute x is known, and an "online step" whose complexity, and in particular its dependence on x, are extremely simple. Explicitly, in the offline phase, one can compute $\mathbf{c}_0 = \mathbf{s}\mathbf{A}_0$, and then for all $i \in [k]$ and $b \in \{0,1\}$ the value $\mathbf{c}_{i,b} = \mathbf{s}(\mathbf{A}_i + b\mathbf{G})$. Then, in the offline phase, simply output $\mathbf{c}_0$ and $\mathbf{c}_{i,x_i}$ for all i.

The key technique in [BGG+14] is a two-fold homomorphic evaluation procedure which works as follows. Given $(\mathbf{A}_1, \ldots, \mathbf{A}_k)$, a function f, and $x \in \{0,1\}^k$, it is possible to derive a matrix $\mathbf{H}$ with relatively small values (essentially exponential in the depth of f, which should be considered small in our setting), s.t.

$$[\mathbf{A}_1 + x_1\mathbf{G}|| \ldots ||\mathbf{A}_k + x_k\mathbf{G}] \cdot \mathbf{H} = \mathbf{A}_f + f(x)\mathbf{G} \ .$$

The matrix $\mathbf{A}_f$ itself can be efficiently derived from $\vec{\mathbf{A}}$ and f without any knowledge or dependence on x.

Using this evaluation procedure, it is possible, given the ciphertext and f, x, to derive

$$\mathbf{s}[\mathbf{A}_0||\mathbf{A}_1 + x_1\mathbf{G}|| \ldots ||\mathbf{A}_k + x_k\mathbf{G}] \cdot \begin{bmatrix} \mathbf{I} \\ \mathbf{H} \end{bmatrix} = \mathbf{s}[\mathbf{A}_0||\mathbf{A}_f + f(x)\mathbf{G}] \tag{1}$$

up to (somewhat increased but, with proper choice of parameters, still small) noise.

The secret key for f is a pair of "relatively short" vectors $(\mathbf{r}_0, \mathbf{r}_f)$ s.t. $\mathbf{A}_0\mathbf{r}_0 + \mathbf{A}_f\mathbf{r}_f = 0 \pmod q$. Such vectors can be derived given a so-called lattice trapdoor for $\mathbf{A}_0$. Decryption involves multiplying Eq. (1) by the column vector consisting of $\mathbf{r}_0$ and $\mathbf{r}_f$, and checking that the result is small. It is easy to check that decryption succeeds if $f(x) = 0$.

The resulting scheme is a KP-ABE scheme with succinct keys, since the keys contain, apart from f itself, two vectors with short entries whose dimensions are independent of f. The ciphertexts, however, are not succinct: for each attribute bit we need to include an entire vector $\mathbf{c}_{i,x_i}$.

An Approach for (Ciphertext-Policy) ABE with Succinct Ciphertexts. We would like to turn the aforementioned KP-ABE scheme on its head: if we could only invert the roles of the secret key and ciphertext, we would be in good shape, since now the ciphertexts will encode f and be succinct, and the key size depend polynomially on $|x|$ (which is fine for the BE application as $|x| = O(\log N)$). However, duality of this sort does not seem straightforward. (Indeed, we considered and dismissed a strawman attempt using the universal circuit earlier on in this introduction.) We therefore consider the following high level outline.

Whenever we wish to encrypt a message with respect to policy f, we will generate a new instance of the KP-ABE scheme. We will generate a key with respect to f and then toss the master secret key of the scheme. We then produce the offline phase of the encryption, i.e. create all ciphertext pieces $\mathbf{c}_0$ and $\{\mathbf{c}_{i,b}\}_{i\in[k],b\in\{0,1\}}$. Now, seemingly, all we need to do is find a way to encode the $\mathbf{c}_{i,b}$ pieces so that user x can only recover $\mathbf{c}_{i,x_i}$.

The first idea that comes to mind is to use an IBE scheme (as was done in [BV16]): just encrypt each $\mathbf{c}_{i,b}$ with respect to attribute (i,b) and provide user x with a key that decrypts all ciphertexts with attribute (i,x_i). Indeed, this will provide functionality, and even security against an adversary that only gets access to a single decryption key, but not *collusion resistance*: if users x, x' come together, they will be able to learn, for some i, both $\mathbf{c}_{i,0}$ and $\mathbf{c}_{i,1}$, which will render the ciphertext completely insecure.

Our key idea to solve the collusion-resistance problem is to construct a special variant of IBE which we refer to as a *randomizing IBE* scheme. During the decryption process, a randomizing IBE scheme modifies the "message", specifically the $\mathbf{s}$ part of the ciphertext, so that users x, x' will each receive a variant of $\mathbf{c}_{i,x_i}$ or $\mathbf{c}_{i,x'_i}$ respectively, but such that each of them corresponds to a different $\mathbf{s}$. If we were able to do this in such a way that the $\mathbf{s}$ values are completely uniform and independent, and such that in some way the noise values in the $\mathbf{c}_{i,x_i}$ pieces are also randomized, then we will have solved the problem. Unfortunately we cannot achieve this and we therefore resort to a heuristic.

Before describing our heuristic, let us point out that even the variant which enjoys only one-key security (and is not collusion resistant) is not fully succinct, in the sense that the ciphertext size grows with the input length $k = |x|$ (times a polynomial in the security parameter). However, this level of succinctness is sufficient for all purposes covered in this work, in particular for broadcast encryption.

Our Heuristic: Select on the Left, Decrypt on the Right. We first "beef up" the KP-ABE scheme. Instead of using a single vector $\mathbf{s}$, we will use a matrix $\mathbf{S}$, or equivalently we will generate a number of independent ciphertexts with independent $\mathbf{s}$ values. Intuitively, this will be useful since it will allow us to define, for each x, its own personal $\mathbf{s}_x$ that will be defined as (roughly) a subset sum of the rows of $\mathbf{S}$, with coefficients that will be specified for x. This modification changes very little in the semantics of the scheme. It still has the same offline/online nature, only now the pieces $\mathbf{C}_0 \approx \mathbf{SA}$ and $\mathbf{C}_{i,b} \approx \mathbf{S}(\mathbf{A}_i + b\mathbf{G})$ are matrices and not vectors. Recall that decryption only involves multiplying the respective ciphertext pieces on the *right*, first by the matrix $\mathbf{H}$ and then by the secret key vectors $(\mathbf{r}_0, \mathbf{r}_f)$. We will therefore mask the ciphertext pieces, using an encoding that is decryptable *on the left*.

Concretely, we will consider matrices $\mathbf{B}_0$ and $\{\mathbf{B}_{i,b}\}$, and set $\hat{\mathbf{C}}_{i,b} = \mathbf{B}_{i,b}\hat{\mathbf{S}}_{i,b} + \mathbf{E}_{i,b} + \mathbf{C}_{i,b}$, and likewise $\hat{\mathbf{C}}_0 = \mathbf{B}_0\mathbf{S}_0 + \mathbf{E}_0 + \mathbf{C}_0$ will encrypt $\mathbf{C}_0$. In itself, the LWE assumption implies that without any additional information on the $\mathbf{B}$ matrices, the $\hat{\mathbf{C}}$ pieces are indistinguishable from uniform, even given arbitrary side information on the $\mathbf{C}$ pieces.

Now, in order to allow decryption, we will provide user x with a short vector $\mathbf{t}_x$ which is sampled randomly (from a discrete Gaussian) subject to $\mathbf{t}_x\mathbf{B}_{i,x_i} = 0 \pmod q$ for all i, and in addition $\mathbf{t}_x\mathbf{B}_0 = 0 \pmod q$. Such a vector can be efficiently sampled if the $\mathbf{B}$ matrices are generated together with a trapdoor for the *intersection* of lattices that they represent. The details are not very important at this point, however we note that crucially, $\mathbf{t}_x\mathbf{B}_{i,1-x_i}$ have no structure, and are uniformly random.

Note that given $\mathbf{t}_x$, it is possible to derive $\mathbf{t}_x\hat{\mathbf{C}}_{i,x_i} \approx \mathbf{t}_x\mathbf{C}_{i,x_i}$, but for $\hat{\mathbf{C}}_{i,1-x_i}$ the vector $\mathbf{t}_x$ is useless (at least heuristically) and does not enable deriving any useful variant of $\mathbf{C}_{i,1-x_i}$. Since the decryption procedure is performed by right multiplication, correctness is maintained even when applying decryption to the $\mathbf{t}_x\mathbf{C}_{i,x_i}$ values. The hope is that collusion attacks are prevented since an adversary that is given, say, $\mathbf{t}_x, \mathbf{t}_{x'}$ for two attribute vectors x and x' should not be able to correlate the ciphertext pieces that it derives in the selection process.

This completes the description of our scheme. Note that we obtain succinct ciphertexts and succinct keys. We only use the "symmetric key" version of the KP-ABE scheme, and only use it to generate a single key. These are very appealing properties in terms of functionality, but we have to be very careful when arguing about its security. We will discuss this in the following section.

1.3 Security and Attempted Cryptanalysis

Our heuristic scheme is based on similar logic to the one standing behind the [BGG+14] KP-ABE scheme. We mask a set of values using LWE instances of the form $\mathbf{SA} + \mathbf{E}$, and provide a trapdoor that exposes only certain relations between the masked values. Let us try to speculate what would be the nature of a security proof for our scheme, and focus on possible aspects where problems could arise. We consider the weakest notion of security, namely fully selective security, and focus only on the broadcast encryption application rather than consider CP-ABE in full generality. Specifically, we consider a security game, where the adversary selects k which is logarithmic in the security parameter, and then specifies two *disjoint* sets $X, S \subseteq \{0,1\}^k$. The set X indicates "corrupted" x values, and the set S indicates the recipients of the broadcast. The adversary then receives the public parameters of the scheme, the secret keys for all $x \in X$, and a ciphertext encrypting a random bit with respect to the set S. The adversary is required to guess the value of the encrypted bit with non-trivial success probability.

Let us take the cryptanalyst's perspective and see what kind of potentially-sensitive information can be derived from the adversary's view of the experiment. Assume that the set X is simply a set of randomly selected identities. Then it must be the case that for every index i, roughly half $x \in X$ have $x_i = 0$ and half have $x_i = 1$. This means that the adversary obtains many short $\mathbf{t}_x$ that annihilate $\mathbf{B}_{i,0}$ (and likewise many *other* short vectors that annihilate $\mathbf{B}_{i,1}$). Therefore, the adversary can use these values to reconstruct a trapdoor for all individual matrices $\mathbf{B}_{i,b}$.

While this feature appears to be concerning, it does not seem to lead to a break of any sort. Indeed, what is needed in order to learn meaningful information is a trapdoor (or at least a short vector) that *simultaneously* annihilates new subsets of the $\mathbf{B}_{i,b}$ matrices. This perspective can be applied when notice that as a necessary condition for security, it should not be possible, given secret keys for the set X, to efficiently generate a valid key for any $x^* \notin X$. Note that this in particular means finding a short vector that annihilates a new subset of the $\mathbf{B}_{i,b}$, and doing so would immediately violate security. This again translates to a question on "mixing and matching" or "malleability" of lattice trapdoors. We show in Section 5 that learning with error trapdoors have a "non-malleability" property so that it is computationally intractable to mix-and-match lattice trapdoors in a way that will compromise the hardness of LWE on a lattice to which an explicit trapdoor has not been given. This is true even if this target lattice is an intersection of lattices for which we provide a trapdoor. This shows that such "key-recovery attacks" on our system are not possible under the LWE assumption.

Of course, so far we only considered very limited attack strategies, namely ones that only observe the set of keys of colluding adversarial parties, and not the challenge ciphertext itself.

Indeed, when trying to convert ABE schemes into the stronger notions of Functional Encryption (FE) or Witness Encryption (WE), collusion resistance is lost because even keys that are not authorized to decrypt reveal sufficient information about the *encryption randomness*, which allows to recover it and break security. In this context, we usually consider the set of linear equations that are obtained by applying the keys that the adversary has to the ciphertext. Of course none of the keys individually will decrypt, but each such attempted decryption implies a linear equation over the unknowns which are the randomness in the encryption (in our case the matrix $\mathbf{S}$ and the LWE noise values). The hope is to obtain sufficiently many such equations to recover, say, $\mathbf{S}$.

At first glance, this approach appears to be successful. It appears that given sufficiently many keys, it should be possible to generate more equations than variables, and thus supposedly break security. However, a seemingly minor detail plays an important role here. The set of equations that can be obtained in this way is linearly independent *over the integers*, but the equations themselves are only applicable *modulo* q, and it turns out (not surprisingly, as we explain below) that the equations obtained are not full rank modulo q, and thus the information obtained in this way appears to be useless to an attacker.

This is indeed a significant difference between our approach and the aforementioned approaches towards FE and WE. In our case, the adversary is unable to obtain equations over the integers, only modular ones (since decrypting with a non-certified key should not yield any functionality, as opposed to FE where all keys have correctness requirement with respect to all ciphertexts). To illustrate this issue, consider a plain LWE instance of the form $(\mathbf{A}, \mathbf{b} = \mathbf{sA} + \mathbf{e})$. Now assume that we are given a short full rank matrix $\mathbf{T}$ s.t. $\mathbf{AT} = 0$. This matrix is a trapdoor for $\mathbf{A}$ and can be used to recover the secrets $\mathbf{s}, \mathbf{e}$ via computing the value $\mathbf{bT} \pmod q$ and noticing that this is equal to the value $\mathbf{eT}$ *over the integers*. Since $\mathbf{T}$ is invertible over the integers $\mathbf{e}$ can be recovered, which leads to recovery of $\mathbf{s}$ as well. However, if $\mathbf{b} = \mathbf{sA} + \mathbf{y}$ for a non-short $\mathbf{y}$ ($\mathbf{y}$ can still be very structured, e.g. LWE instance with respect to a known matrix $\mathbf{B}$), this approach fails. This is because now we can only obtain $\mathbf{yT} \pmod q$, and we note that modulo q the matrix $\mathbf{T}$ is degenerate since it lies in the mod-q kernel of $\mathbf{A}$. It appears that all "linear" cryptanalysis attempts fall into this framework, but we are unable to show a formal statement of this form.

To conclude, it appears that attacks that were successful in other contexts are not applicable to our scheme. This is of course far from constituting a proof of security. As explained above, we hope that one of the readers of this manuscript will be able to advance it, either in the direction of a proof, or in the direction of finding new attacks.

1.4 Comparison to Agrawal-Yamada

The recent work of Agrawal and Yamada [AY20] constructs broadcast encryption and succinct ciphertext-policy ABE relying on the LWE assumption in addition to generic groups with bilinear maps. Their proof of security relies on the *generic model* for groups with bilinear maps. As explained above, our scheme predates this work. Let us point out a few points of comparison.

In terms of the final result, our work relies only on the lattice structure, allows CP-ABE for all functions, but has no proof of security. Their work relies on (standard) LWE and generic bilinear groups (and therefore immediately broken by quantum attacks). Perhaps the most significant difference is that their CP-ABE scheme only supports NC^1 functions, whereas we support all polynomial-time computable functions, with an a-priori bounded (polynomial) depth. Additionally, our scheme has a polynomial LWE modulus for NC^1 functions, whereas their modulus is necessarily exponential, due to the interaction with bilinear groups. However [AY20] has the signif-

icant advantage of being provably secure in a reasonable (yet heuristic) attack model. A followup work [AWY20] replaces the generic group model in [AY20] with a knowledge assumption, however the properties of their construction remain the same.

In terms of techniques, their scheme also starts with a KP-ABE scheme with succinct keys and attempts to "turn it on its head". They use the bilinear structure to prevent collusion by only allowing the parties to see their ciphertext in the exponent of a group generator which is distinct for each user.

Acknowledgements. We thank Shweta Agrawal, Dan Boneh, Yilei Chen, Sam Kim, Alex Lombardi, Rotem Tsabary, and Hoeteck Wee for discussions about the scheme and attempts at proofs and cryptanalysis.

2 Preliminaries

2.1 Attribute Based Encryption (ABE)

Let $\mathcal{F} = \{\mathcal{F}_\lambda\}_\lambda$ be an ensemble of function classes such that $\mathcal{F}_\lambda \subseteq \{0,1\}^* \to \{0,1\}$. We assume that the functions are represented as Boolean circuits. A ciphertext-policy attribute based encryption (CP-ABE) scheme is defined by PPT algorithms $\mathsf{ABE} = (\mathsf{ABE.Params}, \mathsf{ABE.Enc}, \mathsf{ABE.Keygen}, \mathsf{ABE.Dec})$ such that:

- The setup algorithm $\mathsf{ABE.Params}(1^\lambda)$ takes the security parameter and the input length of the supported functions as input and outputs a master secret key msk and a set of public parameters pp.
- The encryption algorithm $\mathsf{ABE.Enc}_{\mathsf{pp}}(\mu, f)$ uses the public parameters pp and takes as input a function $f \in \mathcal{F}_\lambda$ and a message μ from a message space $\mathcal{M} = \mathcal{M}_\lambda$. It outputs a ciphertext $\mathsf{ct} \in \{0,1\}^*$.
- The key generation algorithm $\mathsf{ABE.Keygen}_{\mathsf{msk}}(x)$ uses the master secret key msk and takes as input an attribute vector $x \in \{0,1\}^k$. It outputs a secret key sk_x.
- The decryption algorithm $\mathsf{ABE.Dec}_{\mathsf{pp}}(\mathsf{sk}_x, x, f, \mathsf{ct}_f)$ takes as input a function secret key sk_f, an attribute $x \in \{0,1\}^*$, a function f and a ciphertext ct_f, and outputs a message $\mu' \in \mathcal{M}$.

Remark 1. *It is often the case that a construction of ABE (whether key or ciphertext policy) is applicable to a parameterized class of functions. For example, we may be able to construct ABE for all $\mathcal{F}^{(c)} = \{\mathcal{F}^{(c)}_\lambda\}_\lambda$, where $\mathcal{F}^{(c)}_\lambda$ denotes the set of all circuits of depth λ^c, but such that the parameters of the scheme vary with c. In such cases, we sometimes consider c (or λ^c in this case), as an additional input to the setup algorithm rather than an external specification of the function class. Usually, and specifically in the context of this manuscript, when the setup algorithm takes additional inputs, we explain how these inputs refer to the function classes supported by the construction.*

Remark 2. *We note that the ABE key generation algorithm needs to know (and can run in time polynomial in) the input length of the functions, but* not the size of the circuit computing the supported functions. *This is a key feature of CP-ABE schemes which cannot be obtained generically from KP-ABE schemes, flipping the role of the function and input using a universal circuit.*

Definition 2.1 (Correctness of CP-ABE)**.** *A scheme* ABE *is* correct *if the following holds. Consider a sequence of functions* $\{f_\lambda \in \mathcal{F}_\lambda\}_\lambda$ *and a sequence of attributes* $\{x_\lambda \in \{0,1\}^*\}_\lambda$*, such that for all* λ*, the input size of* f_λ *is* exactly $|x_\lambda|$ *and* $f_\lambda(x_\lambda) = 0$.[5] *For all such sequences and for any sequence* $\{m_\lambda \in \mathcal{M}_\lambda\}_\lambda$*, it holds that*

$$\Pr[\mathsf{ABE.Dec}_{\mathsf{pp}}(\mathsf{sk}_x, x, f, \mathsf{ct}_f) \neq \mu] = \mathrm{negl}(\lambda) \ ,$$

where $(\mathsf{msk}, \mathsf{pp}) = \mathsf{ABE.Params}(1^\lambda, 1^k)$*,* $\mathsf{ct} = \mathsf{ABE.Enc}_{\mathsf{pp}}(\mu, f)$*,* $\mathsf{sk}_f = \mathsf{ABE.Keygen}_{\mathsf{msk}}(x)$*.*

Definition 2.2 (Security for CP-ABE)**.** *Let* ABE *be a CP-ABE encryption scheme as above, and consider the following game between the challenger and adversary.*

1. *The challenger generates* $(\mathsf{msk}, \mathsf{pp}) = \mathsf{ABE.Params}(1^\lambda, 1^k)$*, and sends* pp *to the adversary.*
2. *The adversary makes arbitrarily many key queries by sending attributes* x_i *to the challenger. Upon receiving such an attribute, the challenger creates* $\mathsf{sk}_i = \mathsf{ABE.Keygen}_{\mathsf{msk}}(x_i)$ *and sends* sk_i *to the adversary.*
3. *The adversary sends a function* f *and a pair of messages* μ_0, μ_1 *to the challenger. The challenger samples* $b \in \{0,1\}$ *and computes the challenge ciphertext* $\mathsf{ct}^* = \mathsf{ABE.Enc}_{\mathsf{pp}}(\mu_b, f)$*. It sends* ct^* *to the adversary.*
4. *The adversary makes arbitrarily many key queries as in Step 2 above.*
5. *The adversary outputs* $\tilde{b} \in \{0,1\}$*.*
6. *Let* legal *denote the event where all key queries of the adversary are such that* $f(x_i) = 1$*. If* legal*, the output of the game is* $b' = \tilde{b}$*, otherwise the output* b' *is a uniformly random bit.*

The advantage *of an adversary* $\mathcal{A}$ *is* $|\Pr[b' = b] - 1/2|$*, where* b, b' *are generated in the game played between the challenger and the adversary* $\mathcal{A}(1^\lambda)$*.*

The game above is called the adaptive security game *for ABE, and it has relaxed variants. In the* selective security game*, the adversary sends* f *before Step 1.*

The scheme ABE *is* adaptively/selectively secure *if any* PPT *adversary* $\mathcal{A}$ *only has negligible advantage in the adaptive/selective security game (respectively).*

Negated Policies. We allow decryption when $f(x) = 0$ and require that in the security game all queries are such that $f(x_i) = 1$. In LWE-based constructions it is often much more convenient to work with this negated version of the policy, which explains the apparent strangeness. This variant is obviously equivalent.

Succinct Ciphertext-Policy ABE. The crucial succinctness property we care about in this work is that the ciphertext size depends polynomially in the input length of the functions, namely k, and the security parameter, but is otherwise independent of the size of the functions being encrypted. As a result of a technical deficit common to all known LWE based ABE schemes, we will allow the ciphertext size to grow with the depth of the circuits (but not the size).

[5] Recall our convention that $f(x) = 0$ is the event when decryption succeeds.

2.2 Learning with Errors and Lattice Trapdoors

This section summarizes tools from previous works that are used in our construction. This includes the definition of the LWE problem and its relation to worst case lattice problems, the notion of trapdoors for lattices and operations on trapdoors, and homomorphic evaluation of matrices with special properties.

Learning with Errors (LWE). The Learning with Errors (LWE) problem was introduced by Regev [Reg05] as a generalization of "learning parity with noise" [BFKL93, Ale03]. We now define the decisional version of LWE. (Unless otherwise stated, we will treat all vectors as column vectors in this paper).

Definition 2.3 (Decisional LWE (DLWE) [Reg05])**.** *Let λ be the security parameter, $n = n(\lambda)$, $m = m(\lambda)$, and $q = q(\lambda)$ be integers and $\chi = \chi(\lambda)$ be a probability distribution over $\mathbb{Z}$. The* $\mathrm{DLWE}_{n,q,\chi}$ *problem states that for all $m = \mathrm{poly}(n)$, letting $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\mathbf{e} \leftarrow \chi^m$, and $\mathbf{u} \leftarrow \mathbb{Z}_q^m$, the following distributions are computationally indistinguishable:*

$$\left(\mathbf{A}, \mathbf{s}^T\mathbf{A} + \mathbf{e}^T\right) \overset{c}{\approx} \left(\mathbf{A}, \mathbf{u}^T\right)$$

There are known quantum (Regev [Reg05]) and classical (Peikert [Pei09]) reductions between $\mathrm{DLWE}_{n,q,\chi}$ and approximating short vector problems in lattices. Specifically, these reductions take χ to be a discrete Gaussian distribution $D_{\mathbb{Z},\alpha q}$ for some $\alpha < 1$. We write $\mathrm{DLWE}_{n,q,\alpha}$ to indicate this instantiation. We now state a corollary of the results of [Reg05, Pei09, MM11, MP12]. These results also extend to additional forms of q (see [MM11, MP12]).

Corollary 1 ([Reg05, Pei09, MM11, MP12])**.** *Let $q = q(n) \in \mathbb{N}$ be either a prime power $q = p^r$, or a product of co-prime numbers $q = \prod q_i$ such that for all i, $q_i = \mathrm{poly}(n)$, and let $\alpha \geq \sqrt{n}/q$. If there is an efficient algorithm that solves the (average-case)* $\mathrm{DLWE}_{n,q,\alpha}$ *problem, then:*

- *There is an efficient quantum algorithm that solves* $\mathsf{GapSVP}_{\tilde{O}(n/\alpha)}$ *(and* $\mathsf{SIVP}_{\tilde{O}(n/\alpha)}$*) on any n-dimensional lattice.*
- *If in addition $q \geq \tilde{O}(2^{n/2})$, there is an efficient classical algorithm for* $\mathsf{GapSVP}_{\tilde{O}(n/\alpha)}$ *on any n-dimensional lattice.*

Recall that GapSVP_γ is the (promise) problem of distinguishing, given a basis for a lattice and a parameter d, between the case where the lattice has a vector shorter than d, and the case where the lattice doesn't have any vector shorter than $\gamma \cdot d$. SIVP is the search problem of finding a set of "short" vectors. The best known algorithms for GapSVP_γ require at least $2^{\tilde{\Omega}(n/\log\gamma)}$ time [Sch87]. We refer the reader to [Reg05, Pei09] for more information.

In this work, we will only consider the case where $q \leq 2^n$. Furthermore, the underlying security parameter λ is assumed to be polynomially related to the dimension n.

Lastly, we derive the following corollary which will allow us to choose the LWE parameters for our scheme. The corollary follows immediately from the fact that the discrete Gaussian $D_{\mathbb{Z},\alpha q}$ is $(\alpha q \cdot t, 2^{-\Omega(t^2)})$-bounded for all t.

Corollary 2. *For all $\epsilon > 0$ there exist functions $q = q(n) \leq 2^n, \chi = \chi(n)$ such that χ is B-bounded for some $B = B(n)$, $q/B \geq 2^{n^\epsilon}$ and such that* $\mathrm{DLWE}_{n,q,\chi}$ *is at least as hard as the classical hardness of* GapSVP_γ *and the quantum hardness of* SIVP_γ *for $\gamma = 2^{\Omega(n^\epsilon)}$.*

The Gadget Matrix. Let $N = n \cdot \lceil \log q \rceil$ and define the "gadget matrix" $\mathbf{G} = \mathbf{g} \otimes \mathbf{I}_n \in \mathbb{Z}_q^{n \times N}$ where $\mathbf{g} = (1, 2, 4, \ldots, 2^{\lceil \log q \rceil - 1}) \in \mathbb{Z}_q^{\lceil \log q \rceil}$. We will also refer to this gadget matrix as the "powers-of-two" matrix. We define the inverse function $\mathbf{G}^{-1} : \mathbb{Z}_q^{n \times m} \to \{0,1\}^{N \times m}$ which expands each entry $a \in \mathbb{Z}_q$ of the input matrix into a column of size $\lceil \log q \rceil$ consisting of the bits of the binary representation of a. We have the property that for any matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, it holds that $\mathbf{G} \cdot \mathbf{G}^{-1}(\mathbf{A}) = \mathbf{A}$.

Trapdoors. Let $n, m, q \in \mathbb{N}$ and consider a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$. For all $\mathbf{V} \in \mathbb{Z}_q^{n \times m'}$, we let $\mathbf{A}_\tau^{-1}(\mathbf{V})$ denote the random variable whose distribution is a Gaussian $D_{\mathbb{Z}^m, \tau}^{m'}$ conditioned on $\mathbf{A} \cdot \mathbf{A}_\tau^{-1}(\mathbf{V}) = \mathbf{V}$. A τ-trapdoor for $\mathbf{A}$ is a procedure that can sample from the distribution $\mathbf{A}_\tau^{-1}(\mathbf{V})$ in time $\mathrm{poly}(n, m, m', \log q)$, for any $\mathbf{V}$. We slightly overload notation and denote a τ-trapdoor for $\mathbf{A}$ by $\mathbf{A}_\tau^{-1}$.

The following properties had been established in a long sequence of works.

Corollary 3 (Properties of Trapdoors [Ajt96, Ajt99, GPV08, ABB10a, CHKP12, ABB10b, MP12])**.** *Lattice trapdoors exhibit the following properties.*

1. *Given* $\mathbf{A}_\tau^{-1}$*, one can obtain* $\mathbf{A}_{\tau'}^{-1}$ *for* any $\tau' \geq \tau$*.*
2. *Given* $\mathbf{A}_\tau^{-1}$*, one can obtain* $[\mathbf{A} \| \mathbf{B}]_\tau^{-1}$ *and* $[\mathbf{B} \| \mathbf{A}]_\tau^{-1}$ *for* any $\mathbf{B}$*.*
3. *For all* $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ *and* $\mathbf{R} \in \mathbb{Z}^{m \times N}$*, with* $N = n \lceil \log q \rceil$*, one can obtain* $[\mathbf{AR} + \mathbf{G} \| \mathbf{A}]_\tau^{-1}$ *for* $\tau = O(m \cdot \|\mathbf{R}\|_\infty)$*.*
4. *There exists an efficient procedure* $\mathsf{TrapGen}(1^n, q)$ *that outputs* $(\mathbf{A}, \mathbf{A}_{\tau_0}^{-1})$ *where* $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ *for some* $m = O(n \log q)$ *and is* 2^{-n}*-uniform, where* $\tau_0 = O(\sqrt{n \log q \log n})$*.*

3 The BGG$^+$-Lite Single-Key KP-ABE Scheme

In this section, we recall the KP-ABE scheme of Boneh et al. [BGG+14]. This scheme will be the basis of our construction in Section 4. In fact, we present a simplified version of the scheme which works only to generate a single function key. We start by describing the notion of homomorphic evaluation that underlies their construction.

Key-Homomorphic Evaluation. Let f be a boolean circuit of depth d computing a function from $\{0,1\}^k$ to $\{0,1\}$, and assume that f contains only NAND gates. We will show a version of f that computes on matrices. In particular, for every NAND gate whose inputs are associated to matrices $\mathbf{A}_1$ and $\mathbf{A}_2$, we associate to the output the matrix

$$\mathbf{A}_{\mathsf{nand}} := \mathbf{A}_1 \mathbf{G}^{-1}(\mathbf{A}_2) - \mathbf{G} \tag{2}$$

Note that

$$\begin{aligned}[\mathbf{A}_1 + x_1\mathbf{G} \| \mathbf{A}_2 + x_2\mathbf{G}] \cdot \begin{bmatrix} \mathbf{G}^{-1}(\mathbf{A}_2) \\ -x_1\mathbf{I} \end{bmatrix} &= \mathbf{A}_{\mathsf{nand}} + (1 - x_1x_2)\mathbf{G} \\ &= \mathbf{A}_{\mathsf{nand}} + \mathsf{NAND}(x_1, x_2) \cdot \mathbf{G}\end{aligned}$$

which provides us with a mathematical framework for homomorphic evaluation.

In particular, given matrices $\vec{\mathbf{A}} := (\mathbf{A}_1, \ldots, \mathbf{A}_k)$, we can define the matrix $\mathbf{A}_f$ associated to the output wire of f (by using the NAND rules above repeatedly). This has the associated homomorphic evaluation property that

$$[\mathbf{A}_1 + x_1\mathbf{G}|| \ldots ||\mathbf{A}_k + x_k\mathbf{G}] \cdot \mathbf{H}_{f,x,\vec{\mathbf{A}}} = \mathbf{A}_f + f(x_1, \ldots, x_k) \cdot \mathbf{G} \tag{3}$$

for some matrix $\mathbf{H}_{f,x,\vec{\mathbf{A}}}$ that has norm $(n \log q)^{O(d)}$.

The KP-ABE-Lite Scheme. We are now ready to describe a simplified, single-key, version of the KP-ABE scheme of [BGG+14].

- $\mathsf{BGG.KeyGen}(1^\lambda, 1^k, f)$ generates LWE matrices $\mathbf{A}_1, \ldots, \mathbf{A}_k \in \mathbb{Z}_q^{n \times \ell}$. It generates a random short matrix $\mathbf{T}_f \in \mathbb{Z}^m$ and sets $\mathbf{A}_0 := \mathbf{A}_f\mathbf{T}_f \pmod q$ where $\mathbf{A}_f$ is the matrix described above.

  The master public key is
  $$\mathsf{BGG.MPK} = (\mathbf{A}_0, \mathbf{A}_1, \ldots, \mathbf{A}_k)$$
  and the functional key for f is
  $$\mathsf{BGG.sk}_f = \mathbf{T}_f$$

- $\mathsf{BGG.Enc}(\mathsf{BGG.MPK}, x, \mu)$ where $x \in \{0,1\}^k$ and $\mu \in \{0,1\}$ does the following. Pick a random LWE secret $\mathbf{s} \in \mathbb{Z}_q^n$ and let $\mathbf{u} = \mathbf{s}\mathbf{A}_0 + \mathbf{e}_0$ if $\mu = 0$ and a random vector if $\mu = 1$.

  Output the ciphertext
  $$\mathsf{BGG.ct} := \left(\mathbf{u}, \mathbf{s}[\mathbf{A}_1 + x_1\mathbf{G}|| \ldots ||\mathbf{A}_k + x_k\mathbf{G}] + \mathbf{e}\right)$$
  where $\mathbf{e}$ is a small Gaussian error.

  Looking ahead, we also define the $\mathsf{BGG.OfflineEnc}$ procedure that simply takes the master public key and produces "encryptions relative to all possible x". That is, it outputs
  $$\mathbf{u}, \begin{bmatrix} \mathbf{s}\mathbf{A}_1 & || & \ldots & || & \mathbf{s}\mathbf{A}_k \\ \mathbf{s}(\mathbf{A}_1 + \mathbf{G}) & || & \ldots & || & \mathbf{s}(\mathbf{A}_k + \mathbf{G}) \end{bmatrix} + \mathbf{e}$$

- $\mathsf{BGG.Dec}(\mathsf{BGG.sk}_f, \mathsf{BGG.ct})$ uses Equation 3 to compute
  $$\left(\mathbf{s}[\mathbf{A}_1 + x_1\mathbf{G}|| \ldots ||\mathbf{A}_k + x_k\mathbf{G}] + \mathbf{e}\right) \cdot \mathbf{H}_{f,x,\vec{\mathbf{A}}} \approx \mathbf{s}(\mathbf{A}_f + f(\vec{x})\mathbf{G})$$
  Thus, when $f(x) = 0$ (which is interpreted as `true`), we can decrypt by computing $\mathbf{s}\mathbf{A}_f\mathbf{T}_f$ and checking if it is close to $\mathbf{u}$ (in which case, output $\mu = 0$) or not (in which case, output $\mu = 1$).

We remark that a version of the scheme where the encryption randomness is a matrix $\mathbf{S}$ (as opposed to a single vector $\mathbf{s}$) can be defined in a completely analogous way. Indeed, this will turn out to be useful in the next section.

We will not provide an analysis of the parameters of the scheme here (we will do that in Section 4 for our CP-ABE scheme instead). We remark that single-key KP-ABE schemes can be achieved in simpler ways than described above, for example directly using garbled circuits [SS10], however, the algebraic structure of the BGG$^+$-lite scheme is crucial to us down the line.

We prove in Appendix A the security of the scheme in the (selective, single-key) setting where the adversary obtains a single key for a function f and a ciphertext for an input x^* such that $f(x^*) = 1$ (which is interpreted as `false`). We remark that the proof is already in [BGG$^+$14] and is a slightly simpler version thereof.

4 Our CP-ABE Scheme and Correctness

We describe our CP-ABE scheme below.

$\mathsf{Setup}(1^\lambda)$. Let k be the length of the attributes in the CP-ABE scheme and d be the maximum depth of circuits that will be encrypted. Choose $2k$ uniformly random matrices $\mathbf{B}_{i,b} \in \mathbb{Z}_q^{m\times n}$ in the public parameters (for $i \in [k]$ and $b \in \{0,1\}$), and let the master secret key be a "joint trapdoor" for all the $\mathbf{B}_{i,b}$.

More precisely, run $\mathsf{TrapGen}(1^{2kn}, 1^m, q)$ to generate

$$(\mathbf{B}, \mathbf{T}) \leftarrow \mathsf{TrapGen}(1^{2kn}, 1^m, q)$$

where $\mathbf{B} \in \mathbb{Z}_q^{m\times 2kn}$. Let $\{\mathbf{B}_{i,b}\}_{i\in[k],b\in\{0,1\}}$ be blocks of columns of $\mathbf{B}$; thus,

$$\mathbf{B} = [\mathbf{B}_{1,0}||\mathbf{B}_{1,1}||\dots||\mathbf{B}_{k,0}||\mathbf{B}_{k,1}]$$

In particular, by the properties of the trapdoor, we know that $\mathbf{T}$ has small entries, and that for all i, b:

$$\mathbf{T}\mathbf{B}_{i,b} = \mathbf{0} \pmod q$$

Output

$$\mathsf{MPK} = \{\mathbf{B}_{i,b}\}_{i,b} \text{ and } \mathsf{MSK} = \mathbf{T}$$

$\mathsf{Enc}(\mathsf{MPK}, f, \mu)$. To generate a ciphertext for a function f encrypting a message $\mu \in \{0,1\}$, do the following.

- Pick fresh ABE parameters for the BGG$^+$-lite KP-ABE scheme.

  $$(\mathsf{BGG.MPK}, \mathsf{BGG.sk}_f) \leftarrow \mathsf{BGG.KeyGen}(1^\lambda, 1^k, f)$$

  We recall that $\mathsf{BGG.MPK}$ consists of $k+1$ matrices $\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_k \in \mathbb{Z}_q^{n\times\ell}$ and $\mathsf{BGG.sk}_f$ is a matrix $\mathbf{T}_f$ with small entries such that $\mathbf{A}_f\mathbf{T}_f = \mathbf{A}_0 \pmod q$.

- Run the offline phase of the matrix-BGG$^+$-Lite encryption algorithm. That is, let

  $$(\mathbf{C}_0, \{\mathbf{C}_{i,b}\}) \leftarrow \mathsf{BGG.OfflineEnc}(\mathsf{BGG.MPK})$$

  We recall that $\mathbf{C}_{i,b} = \mathbf{S}(\mathbf{A}_i + b\mathbf{G}) + \mathbf{E}_{i,b}$ and $\mathbf{C}_0 = \mathbf{S}\mathbf{A}_0 + \mathbf{E}_0$ (when $\mu = 0$) or uniformly random (when $\mu = 1$). Here, $\mathbf{S} \in \mathbb{Z}_q^{m\times n}$ is a uniformly random matrix and $\mathbf{E}_0, \{\mathbf{E}_{i,b}\}$ are short error matrices of the appropriate dimension.

Define the matrices

$$\hat{\mathbf{C}}_{i,b} := \mathbf{B}_{i,b}\hat{\mathbf{S}}_{i,b} + \mathbf{C}_{i,b} = \mathbf{B}_{i,b}\hat{\mathbf{S}}_{i,b} + \mathbf{E}_{i,b} + \mathbf{S}(\mathbf{A}_i + b\mathbf{G})$$

and

$$\hat{\mathbf{C}}_0 := \mathbf{B}_0\hat{\mathbf{S}}_0 + \mathbf{C}_0 = \mathbf{B}_0\hat{\mathbf{S}}_0 + \mathbf{E}_0 + \mathbf{S}\mathbf{A}_0$$

where the matrices $\hat{\mathbf{S}}_{i,b}$ and $\hat{\mathbf{S}}_0$ live in $\mathbb{Z}_q^{n\times\ell}$ and are randomly chosen.
The ciphertext is

$$CT_f := \left(\mathsf{BGG.sk}_f, \hat{\mathbf{C}}_0, \{\hat{\mathbf{C}}_{i,b}\}_{i,b}\right)$$

$\mathsf{KeyGen}(\mathsf{MSK}, x)$**.** The key sk_x for an attribute vector $x = (x_1, ..., x_k)$ is a short vector $\mathbf{t}_x$ such that $\mathbf{t}_x\mathbf{B}_0 = 0 \pmod q$ and for every i,

$$\mathbf{t}_x\mathbf{B}_{i,x_i} = 0 \pmod q$$

Such a vector can be generated by running $\mathsf{TrapSamp}$ with the matrix $\mathbf{B}_x := [\mathbf{B}_{1,x_1}||\ldots||\mathbf{B}_{k,x_k}]$ and the trapdoor $\mathbf{T}$. For example, we can use the trapdoor $\mathbf{T}$ to find a short vector $\mathbf{t}_x$ such that $\mathbf{t}_x\mathbf{B}_0 = 0 \pmod q$, $\mathbf{t}_x\mathbf{B}_{i,x_i} = 0 \pmod q$ and $\mathbf{t}_x\mathbf{B}_{i,1-x_i}$ are uniformly random and independent.

$\mathsf{Dec}(x, \mathsf{sk}_x, f, CT_f)$**.** To decrypt, compute

$$\mathbf{t}_x \cdot [\hat{\mathbf{C}}_0||\hat{\mathbf{C}}_{1,x_1}||\ldots||\hat{\mathbf{C}}_{k,x_k}] \cdot \begin{bmatrix} -\mathbf{I} \\ \mathbf{H}_{f,x}\mathbf{T}_f \end{bmatrix} \tag{4}$$

and output 1 if the resulting vector has small entries and 1 otherwise.

Correctness. Let us rewrite the decryption expression 4. Here we will refer to the BGG$^+$-Lite decryption algorithm from Section 3. We have

$$\begin{aligned}
&[\mathbf{C}_0||\mathbf{C}_{1,x_1}||\ldots||\mathbf{C}_{k,x_k}] \cdot \begin{bmatrix} -\mathbf{I} \\ \mathbf{H}_{f,x}\mathbf{T}_f \end{bmatrix} \\
&\approx -\mathbf{C}_0 + \mathbf{S}[\mathbf{A}_1 + x_1\mathbf{G}||\ldots||\mathbf{A}_k + x_k\mathbf{G}] \cdot \mathbf{H}_{f,x}\mathbf{T}_f \\
&\approx -\mathbf{C}_0 + \mathbf{S}[\mathbf{A}_f + f(x)\mathbf{G}]\mathbf{T}_f \\
&\approx -\mathbf{C}_0 + \mathbf{S}\mathbf{A}_0
\end{aligned}$$

where the first equation is by the definition of $\mathbf{C}_{i,b}$, the second by the homomorphic evaluation property (equation 3), and the third by the fact that $\mathbf{A}_f\mathbf{T}_f = \mathbf{A}_0$. Now, if $\mu = 0$, $\mathbf{C}_0 \approx \mathbf{S}\mathbf{A}_0$ and the computation above results in a matrix with small entries; and if $\mu = 1$, it is random and therefore unlikely to have small entries.

Let us now turn to the decryption equation.

$$\begin{aligned}
&\mathbf{t}_x \cdot [\hat{\mathbf{C}}_0||\hat{\mathbf{C}}_{1,x_1}||\ldots||\hat{\mathbf{C}}_{k,x_k}] \cdot \begin{bmatrix} -\mathbf{I} \\ \mathbf{H}_{f,x}\mathbf{T}_f \end{bmatrix} \\
&\quad\approx \mathbf{t}_x[\mathbf{B}_0\hat{\mathbf{S}}_0 + \mathbf{C}_0 \;||\; \mathbf{B}_{1,x_1}\hat{\mathbf{S}}_{1,x_1} + \mathbf{C}_{1,x_1} \;||\; \ldots \;||\; \mathbf{B}_{k,x_k}\hat{\mathbf{S}}_{k,x_k} + \mathbf{C}_{k,x_k}] \cdot \begin{bmatrix} -\mathbf{I} \\ \mathbf{H}_{f,x}\mathbf{T}_f \end{bmatrix} \\
&\quad\approx \mathbf{t}_x \cdot [\mathbf{C}_0||\mathbf{C}_{1,x_1}||\ldots||\mathbf{C}_{k,x_k}] \cdot \begin{bmatrix} -\mathbf{I} \\ \mathbf{H}_{f,x}\mathbf{T}_f \end{bmatrix} \\
&\quad\approx \mathbf{t}_x \cdot (-\mathbf{C}_0 + \mathbf{S}\mathbf{A}_0)
\end{aligned}$$

This resulting vector either has small entries (if $\mu = 0$) or not (if $\mu = 1$) w.h.p.

Parameter Settings. Looking at the decryption equation, we see that the largest asymptotic growth of error happens due to the error growth in BGG+ decryption, i.e., multiplying by the matrix $\mathbf{H}_{f,x}$ which increases the error by a multiplicative factor of $(n \log q)^{O(d)}$ where d is the depth of the circuit computin f. Thus, setting $q \gg (n \log q)^{O(d)}$ and $m = cn \log q$ for a sufficiently large constant $c > 1$ gives us decryption correctness.

Efficiency. The key parameter of interest to us is the size of the ciphertext. This consists of the BGG+ secret key as well as the k elements of the BGG+ ciphertext. In total, the length is $\mathsf{poly}(n, \log q, k)$ which is polynomial in k (the input length of f) and d (the circuit depth of f), but otherwise independent of the circuit size of f. The size of the secret key in the scheme is $\mathsf{poly}(n, \log q, k)$ which is polynomial in the security parameter and the attribute-length k.

4.1 Broadcast Encryption

Recall that in broadcast encryption, we have a setup algorithm that generates a master public/secret key pair for a universe of $N = 2^k$ users, an encryption algorithm that encrypts a message m to a subset $S \subseteq [N]$ of users, a key generation algorithm that gives each user $i \in [N]$ its private key sk_i, and finally a decryption algorithm that gets a user private key sk_i, a ciphertext ct as well as a description of the intended broadcast set S, and outputs μ if and only if $i \in S$.

A construction of broadcast encryption now follows immediately from the succinct CP-ABE scheme by instantiating the function f in the CP-ABE encryption by the indicator function f_S which, on input i, outputs 0 if and only if $i \in S$. The function f_S can be implemented by a circuit of size $O(|S| \cdot \mathsf{poly}(\log k))$ and depth $O(\log |S|) = O(\log N)$.

- The size of the ciphertexts in the scheme is $\mathsf{poly}(n, k, d) = \mathsf{poly}(n, \log N)$.
- The size of the keys in the scheme is $\mathsf{poly}(n, \log N)$ as well, by inspection.

In other words, this gives us a broadcast encryption scheme with polylogarithmic size ciphertexts and keys.

5 Assessment of Security

In this section we attempt to provide a more quantitative version of the high-level discussion in Section 1.3, and attempt to point out what we know about the security of our candidate (recall that we do not have a closed-form assumption that implies the security of the scheme). To avoid unnecessary repetition we will assume that the reader of this section already read Section 1.3 and we may refrain from repeating the high level discussion.

We formalize the notion of a *key recovery attack* where the adversary obtains a set of keys and is tasked with producing a key for a user that is not in the set, and show that our scheme is secure against this kind of attacks. We recall the notation of our scheme from Section 4, and in particular the matrix $\mathbf{B}_x = [\mathbf{B}_0 \| \mathbf{B}_{1,x_1} || \dots || \mathbf{B}_{k,x_k}]$.

Definition 5.1 (Key-Recovery Attack)**.** *In a key-recovery attack, and adversary and a challenger play the following game.*

1. *A challenger generates master public and secret key, and sends the master public key to the adversary.*

2. *The challenger can ask for keys for user x_i and the challenger generates sk_{x_i} and sends it to the adversary. This step can be repeated polynomially many times as the adversary pleases. Let X denote the set of x_i for which a key was requested.*

3. *The adversary sends $x^* \notin X$ and a vector $\mathbf{t}^*$.*

4. *The adversary wins if $\mathbf{t}^*$ is a valid secret key for the user x^*, i.e. $||\mathbf{t}^*|| < q/\mathrm{poly}(\lambda)$ for a sufficiently large polynomial (to be specified) and $\mathbf{t}^*\mathbf{B}_x = 0 \pmod q$.*

We say that our scheme is secure against key recovery if the probability of any efficient adversary to win in this game is negligible. We also consider the fully selective setting in which the adversary provides (X, x^*) before the beginning of the game.

We note that our definition is not completely generic and refers specifically to the possibility of the adversary to generate keys that are of a certain form (a short vector in the lattice specified by $\mathbf{B}_x$).

Recall that an adversary that succeeds in a key recovery attack can easily break the security of the scheme, simply by requesting a challenge ciphertext for a function f that is not satisfied by any element in X but is satisfied by x^*, and then using $\mathbf{t}^*$ in order to decrypt the challenge ciphertext.

Resilience to Key-Recovery and Intersection Attacks. We now show that with the right parameters, our scheme is resilient to key-recovery attacks. Our proof uses the following game that is based on the LWE problem.

Definition 5.2 (Related-Trapdoor Robust LWE)**.** *The related trapdoor LWE game is parameterized by the same parameters as our CP-ABE scheme and is defined as the following game between an adversary and a challenger.*

1. *The adversary chooses $x^* \in \{0,1\}^k$ from the message space and sends it to the challenger.*

2. *The challenger samples $(\mathbf{B}, \mathbf{T})$ as in the setup of the CP-ABE scheme and sends $\mathbf{B}$ to the adversary.*

3. *The adversary can request trapdoors $\mathbf{T}_x$ for matrices $\mathbf{B}_x$ for $x \neq x^*$ of the adversary's choice, from the message space. This is done polynomially many times adaptively.*

4. *The challenger flips a coin β. If $\beta = 0$, it generates LWE instances as follows. It samples $\mathbf{s}, \mathbf{e}$ and outputs $\mathbf{c} = \mathbf{B}_{x^*}\mathbf{s} + \mathbf{e}$. If $\beta = 1$ the challenger produces a uniform vector of the same dimension.*

5. *The adversary produces a guess β' as to the value of β, and wins in the game if $\beta' = \beta$.*

An adversary is successful against related trapdoor LWE if it can win in the game with probability noticeably greater than $1/2$.

Theorem 5.3. *The Related-Trapdoor LWE problem is hard under the hardness of LWE.*

Proof. We will show how to generate the matrix $\mathbf{B}$, given x^*, so that it is possible for every $x \neq x^*$ to generate a trapdoor for $\mathbf{B}_x$, but it is still not possible to solve LWE with respect to $\mathbf{B}_{x^*}$. To do this, we generate the matrices according to the same template:

$$\mathbf{B}_{i,b} = \begin{bmatrix} \mathbf{D}_{i,b} \\ \mathbf{R}\mathbf{D}_{i,b} + \mathbf{G}_{i,b} \end{bmatrix} ,$$

and likewise for $\mathbf{B}_0$ with $\mathbf{D}_0$ and $\mathbf{G}_0$. The $\mathbf{G}$ matrices are "gadgets" that we will define in a particular way in order for security to hold. The $\mathbf{D}$ matrices are uniform and the $\mathbf{R}$ (which is common to all i, b) is a random matrix with small entries, say in $\{0, 1\}$.

The dimensions of the matrices will be determined as follows. We will show below how to define the $\mathbf{G}$ matrices so as to have the security properties that we are looking for. This will determine the dimensions of the lower block in the $\mathbf{B}$ matrices. The remaining dimension of the upper block (i.e. the height of the $\mathbf{D}$ matrices) is determined so that a single $\mathbf{R}$ has sufficient entropy to make all $\mathbf{R}\mathbf{D}_{i,b}$ jointly statistically indistinguishable from uniform. We note that so long as the dimensions of the gadget matrices are polynomial, all other matrices will have polynomially related dimensions.

We note that following the template above, the matrices $\mathbf{B}_x$ also take the form

$$\mathbf{B}_x = \begin{bmatrix} \mathbf{D}_x \\ \mathbf{R}\mathbf{D}_x + \mathbf{G}_x \end{bmatrix} ,$$

where $\mathbf{D}_x, \mathbf{G}_x$ are defined analogously to $\mathbf{B}_x$.

We define the gadgets $\mathbf{G}_{i,b}$ so that for a party that knows $\mathbf{R}$ it is possible to generate trapdoors for all $\mathbf{B}_x$ for $x \neq x^*$, but not for $\mathbf{B}_{x^*}$ (and later we show that indeed LWE remains hard on the latter).

Let $\mathbf{u}_i$ denote the ith indicator column vector. We recall the standard $n \times n \log q$ gadget matrix $\mathbf{G}$, and define $\mathbf{G}_{i,b} = \mathbf{v}_{i,b} \otimes \mathbf{G}^T$ where $\mathbf{v}_{i,b} = \mathbf{u}_{2*(i-1)+b}$. Namely, the matrix $\mathbf{G}_{i,b}$ is a column-block matrix, where all but the $(2i - 2 + b)$th block are 0, and the $(2i - 2 + b)$th block contains $\mathbf{G}^T$. The matrix $\mathbf{G}_0$ is defined as $\mathbf{G}_0 = -\sum_i \mathbf{G}_{i,x_i^*} = \mathbf{v}_0 \otimes \mathbf{G}^T$, where $\mathbf{v}_0 = -\sum_i (\mathbf{u}_{2i+x_i^*})$. Let us further consider, for all x the matrix $\mathbf{V}_x$ whose columns are the vectors $\mathbf{v}_0, \mathbf{v}_{1,x_1}, \ldots, \mathbf{v}_{2k,x_k}$. It holds that $\mathbf{G}_x = \mathbf{V}_x \otimes \mathbf{G}^T$. We notice that for all $x \neq x^*$ the matrix $\mathbf{V}_x$ has full rank, but for $x = x^*$ it is rank deficient. Having $\mathbf{V}_x$ with small entries and full rank allows to generate a trapdoor for $\mathbf{B}_x$ using the techniques of [MP12]. Therefore, we show that it is possible to answer the trapdoor queries for all $x \neq x^*$. The total dimension of $\mathbf{G}_{i,b}$ is therefore $2kn \log q \times n$.

Now, let us consider an LWE instance respective to $\mathbf{B}_{x^*}$. The LWE vector $\mathbf{c}$ can be written as

$$\mathbf{c} = \mathbf{B}_0 \mathbf{s}_0 + \sum_{i \in [k]} \mathbf{B}_{i,x_i^*} \mathbf{s}_i + \mathbf{e} ,$$

where $\mathbf{s}_i$ are the blocks of the LWE secret $\mathbf{s}$. We notice that the problem only becomes easier if all $\mathbf{s}_i$ are equal (but the marginal distribution is still uniform). This is because such an instance can always be rerandomized to take the form above. It therefore suffices to show that

$$\mathbf{c} = (\mathbf{B}_0 + \sum_{i \in [k]} \mathbf{B}_{i,x_i^*}) \mathbf{s}_0 + \mathbf{e}$$

is indistinguishable from uniform. To see this, notice that

$$\mathbf{B}_0 + \sum_{i\in[k]} \mathbf{B}_{i,x_i^*} = \begin{bmatrix} \mathbf{D}_0 + \sum_{i\in[k]} \mathbf{D}_{i,x_i^*} \\ \mathbf{R}(\mathbf{D}_0 + \sum_{i\in[k]} \mathbf{D}_{i,x_i^*}) + (\mathbf{G}_0 + \sum_{i\in[k]} \mathbf{G}_{i,x_i^*}) \end{bmatrix}$$
$$= \begin{bmatrix} \mathbf{I} \\ \mathbf{R} \end{bmatrix} (\mathbf{D}_0 + \sum_{i\in[k]} \mathbf{D}_{i,x_i^*}) \ ,$$

due to the way $\mathbf{G}_0$ was generated. We are therefore left with an LWE witness relative to the matrix $(\mathbf{D}_0 + \sum_{i\in[k]} \mathbf{D}_{i,x_i^*})$, multiplied by a short matrix, which can be simulated as in [BGG+14]. □

Given the above, security against fully selective key recovery follows essentially by definition. For the BE application, the non-selective setting also follows since the space of possible x^* is polynomially bounded, and the proof can generate trapdoors for all $x \neq x^*$ so there is no need to know X in advance.

Expanding The Family of Attacks. We notice that the above can be extended to a template of resilience to different attacks, if we organize the gadget matrices in the proof of related-trapdoor robust LWE in different ways, we can rule out different methods of the adversary's "mix-and-match" capabilities. We claim that this has a potential to cover a large variety of attacks, in particular ones where the adversary creates linear combinations with short coefficients of the ciphertext pieces $\hat{\mathbf{C}}_{i,b}$ and then tries to use the keys that it obtained in order to cancel out the $\mathbf{B}$ part of the resulting linear combination, in order to recover meaningful information about the "internal" LWE secret $\mathbf{S}$ and by extension about the encrypted message. However, this does not seem to be an "ultimate" proof technique, even for the restricted case of linear combinations with short coefficients, since in the full-fledged scenario, the number of "forbidden" combinations (analogous to x^* in the current outline) might become very large and in particular grow with the number of parties in the BE scenario. In such a case, it will not be possible to create gadgets of low enough dimension that adhere to all of these constraints simply due to rank considerations.

References

[ABB10a] Shweta Agrawal, Dan Boneh, and Xavier Boyen. Efficient lattice (H)IBE in the standard model. In Henri Gilbert, editor, *Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Monaco / French Riviera, May 30 - June 3, 2010. Proceedings*, volume 6110 of *Lecture Notes in Computer Science*, pages 553–572. Springer, 2010.

[ABB10b] Shweta Agrawal, Dan Boneh, and Xavier Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In Tal Rabin, editor, *Advances in Cryptology - CRYPTO 2010, 30th Annual Cryptology Conference, Santa Barbara, CA, USA, August 15-19, 2010. Proceedings*, volume 6223 of *Lecture Notes in Computer Science*, pages 98–115. Springer, 2010.

[Ajt96] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In Gary L. Miller, editor, *Proceedings of the Twenty-Eighth Annual ACM Symposium on*

the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996, pages 99–108. ACM, 1996.

[Ajt99] Miklós Ajtai. Generating hard instances of the short basis problem. In Jirí Wiedermann, Peter van Emde Boas, and Mogens Nielsen, editors, *Automata, Languages and Programming, 26th International Colloquium, ICALP'99, Prague, Czech Republic, July 11-15, 1999, Proceedings*, volume 1644 of *Lecture Notes in Computer Science*, pages 1–9. Springer, 1999.

[Ale03] Michael Alekhnovich. More on average case vs approximation complexity. In *44th Symposium on Foundations of Computer Science (FOCS 2003), 11-14 October 2003, Cambridge, MA, USA, Proceedings*, pages 298–307. IEEE Computer Society, 2003.

[AWY20] Shweta Agrawal, Daniel Wichs, and Shota Yamada. Optimal broadcast encryption from LWE and pairings in the standard model. In Rafael Pass and Krzysztof Pietrzak, editors, *Theory of Cryptography - 18th International Conference, TCC 2020, Durham, NC, USA, November 16-19, 2020, Proceedings, Part I*, volume 12550 of *Lecture Notes in Computer Science*, pages 149–178. Springer, 2020.

[AY20] Shweta Agrawal and Shota Yamada. Optimal broadcast encryption from pairings and LWE. Eurocrypt 2020, 2020.

[BFKL93] Avrim Blum, Merrick L. Furst, Michael J. Kearns, and Richard J. Lipton. Cryptographic primitives based on hard learning problems. In Douglas R. Stinson, editor, *Advances in Cryptology - CRYPTO '93, 13th Annual International Cryptology Conference, Santa Barbara, California, USA, August 22-26, 1993, Proceedings*, volume 773 of *Lecture Notes in Computer Science*, pages 278–291. Springer, 1993.

[BGG+14] Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings*, pages 533–556, 2014.

[BGW05] Dan Boneh, Craig Gentry, and Brent Waters. Collusion resistant broadcast encryption with short ciphertexts and private keys. In Victor Shoup, editor, *Advances in Cryptology - CRYPTO 2005: 25th Annual International Cryptology Conference, Santa Barbara, California, USA, August 14-18, 2005, Proceedings*, volume 3621 of *Lecture Notes in Computer Science*, pages 258–275. Springer, 2005.

[BV16] Zvika Brakerski and Vinod Vaikuntanathan. Circuit-abe from LWE: unbounded attributes and semi-adaptive security. In *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part III*, pages 363–384, 2016.

[BWZ14] Dan Boneh, Brent Waters, and Mark Zhandry. Low overhead broadcast encryption from multilinear maps. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara,*

CA, USA, August 17-21, 2014, Proceedings, Part I, volume 8616 of *Lecture Notes in Computer Science*, pages 206–223. Springer, 2014.

[BZ17] Dan Boneh and Mark Zhandry. Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. *Algorithmica*, 79(4):1233–1285, 2017.

[CHKP12] David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. *J. Cryptology*, 25(4):601–639, 2012.

[FN93] Amos Fiat and Moni Naor. Broadcast encryption. In Douglas R. Stinson, editor, *Advances in Cryptology - CRYPTO '93, 13th Annual International Cryptology Conference, Santa Barbara, California, USA, August 22-26, 1993, Proceedings*, volume 773 of *Lecture Notes in Computer Science*, pages 480–491. Springer, 1993.

[FN16] Amos Fiat and Moni Naor. 2016 Paris Kanellakis Award., 2016.

[GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Cynthia Dwork, editor, *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 197–206. ACM, 2008.

[GVW13] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 545–554. ACM, 2013.

[MM11] Daniele Micciancio and Petros Mol. Pseudorandom knapsacks and the sample complexity of LWE search-to-decision reductions. In Phillip Rogaway, editor, *Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2011. Proceedings*, volume 6841 of *Lecture Notes in Computer Science*, pages 465–484. Springer, 2011.

[MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings*, volume 7237 of *Lecture Notes in Computer Science*, pages 700–718. Springer, 2012.

[NNL01] Dalit Naor, Moni Naor, and Jeffery Lotspiech. Revocation and tracing schemes for stateless receivers. In Joe Kilian, editor, *Advances in Cryptology - CRYPTO 2001, 21st Annual International Cryptology Conference, Santa Barbara, California, USA, August 19-23, 2001, Proceedings*, volume 2139 of *Lecture Notes in Computer Science*, pages 41–62. Springer, 2001.

[Pei09] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In Michael Mitzenmacher, editor, *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 333–342. ACM, 2009.

[Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 84–93. ACM, 2005.

[Sch87] Claus-Peter Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theor. Comput. Sci.*, 53:201–224, 1987.

[SS10] Amit Sahai and Hakan Seyalioglu. Worry-free encryption: functional encryption with public keys. In Ehab Al-Shaer, Angelos D. Keromytis, and Vitaly Shmatikov, editors, *Proceedings of the 17th ACM Conference on Computer and Communications Security, CCS 2010, Chicago, Illinois, USA, October 4-8, 2010*, pages 463–472. ACM, 2010.

A Selective, Single-Key Security of BGG$^+$-Lite

Lemma 1. *For every function f and every input x^* such that $f(x^*) = 1$, the following distributions are indistinguishable under the LWE assumption:*

$$\left(\mathsf{BGG.MPK}, \mathsf{BGG.sk}_f, \mathsf{BGG.ct}_0 \leftarrow \mathsf{Enc}(\mathsf{BGG.MPK}, x^*, 0)\right) \approx_c \left(\mathsf{BGG.MPK}, \mathsf{BGG.sk}_f, \mathsf{BGG.ct}_1 \leftarrow \mathsf{Enc}(\mathsf{BGG.MPK}, x^*, 1)\right)$$

Proof. We proceed by a hybrid argument.

Hybrid 0 is the distribution on the right, encrypting $\mu = 0$.

In **Hybrid** 1, we change $\mathsf{BGG.MPK}$ to be of the form $\mathbf{A}_i = \mathbf{A}\mathbf{R}_i - x_i^*\mathbf{G}$ where $\mathbf{A}$ is a random LWE matrix, and $\mathbf{R}_i$ are random matrices with small entries. We note that the second component of the ciphertext looks like

$$\mathbf{s}[\mathbf{A}_1 + x_1^*\mathbf{G}|| \ldots ||\mathbf{A}_k + x_k^*\mathbf{G}] + \mathbf{e}_i = \mathbf{s}\mathbf{A} \cdot [\mathbf{R}_1|| \ldots ||\mathbf{R}_k] + \mathbf{e}$$

Furthermore

$$\mathbf{A}_0 = \mathbf{A}_f\mathbf{T}_f = (\mathbf{A}\mathbf{R}_f + f(x^*)\mathbf{G}) \cdot \mathbf{T}_f = \mathbf{A}\mathbf{R}_f\mathbf{T}_f + \mathbf{G}\mathbf{T}_f$$

Hybrids 0 and 1 are statistically indistinguishable by an application of the leftover hash lemma.

In **Hybrid** 2, we pick a random matrix $\mathbf{A}'_0$ and compute $\mathbf{T}_f := \mathbf{G}^{-1}(\mathbf{A}'_0)$. Use this $\mathbf{T}_f$ as the functional secret key. Hybrids 1 and 2 are perfectly indistinguishable.

Note that now,

$$\mathbf{A}_0 = \mathbf{A}_f\mathbf{T}_f = \mathbf{A}\mathbf{R}_f\mathbf{T}_f + \mathbf{G}\mathbf{T}_f = \mathbf{A}\mathbf{R}_f\mathbf{T}_f + \mathbf{A}'_0$$

In **Hybrid** 3, we notice that the first part of the ciphertext is

$$\mathbf{u} = \mathbf{s}\mathbf{A}_0 = \mathbf{s}\mathbf{A}\mathbf{R}_f\mathbf{T}_f + \mathbf{s}\mathbf{A}'_0$$

and the second part is

$$\mathbf{s}\mathbf{A} \cdot [\mathbf{R}_1|| \ldots ||\mathbf{R}_k] + \mathbf{e}$$

Simulate both of these using $\mathbf{s}\mathbf{A}_0 + \mathbf{e}_0$ and $\mathbf{s}\mathbf{A}'_0 + \mathbf{e}'_0$, LWE samples relative to matrices $\mathbf{A}_0$ and $\mathbf{A}'_0$ respectively (using in addition $\mathbf{R}_i$ and $\mathbf{T}_f$ and noise-flooding appropriately).

Hybrids 2 and 3 are statistically indistinguishable.

In **Hybrid** 4, we replace the first component of the ciphertext by a truly random vector. Hybrids 3 and 4 are computationally indistinguishable by an application of the LWE assumption.

Now, retracing back through hybrids 2, 1 and back to 0, we arrive at the BGG encryption of $\mu = 0$, completing the proof. □