

Homomorphic Encryption from Learning with Errors: Conceptually-Simpler, Asymptotically-Faster, Attribute-Based

Craig Gentry*, Amit Sahai**, and Brent Waters***

Abstract. We describe a comparatively simple fully homomorphic encryption (FHE) scheme based on the learning with errors (LWE) problem. In previous LWE-based FHE schemes, multiplication is a complicated and expensive step involving “relinearization”. In this work, we propose a new technique for building FHE schemes that we call the *approximate eigenvector* method. In our scheme, for the most part, homomorphic addition and multiplication are just matrix addition and multiplication. This makes our scheme both asymptotically faster and (we believe) easier to understand.

In previous schemes, the homomorphic evaluator needs to obtain the user’s “evaluation key”, which consists of a chain of encrypted secret keys. Our scheme has no evaluation key. The evaluator can do homomorphic operations without knowing the user’s public key at all, except for some basic parameters. This fact helps us construct the first identity-based FHE scheme. Using similar techniques, we show how to compile a recent attribute-based encryption scheme for circuits by Gorbunov et al. into an attribute-based FHE scheme that permits data encrypted under the same index to be processed homomorphically.

1 Introduction

Fully homomorphic encryption (FHE) schemes [RAD78, Gen09, Gen10, vDGHV10] [SV10, GH11b, CMNT11, BV11a, BV11b, GH11a, BGV12, CNT12, GHS12a, GHS12b]

* IBM Research. cbgentry@us.ibm.com. This work was supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior National Business Center (DoI/NBC) contract number D11PC20202. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DoI/NBC, or the U.S. Government.

** UCLA. sahai@cs.ucla.edu. Research supported in part from a DARPA/ONR PROCEED award, NSF grants 1228984, 1136174, 1118096, 1065276, 0916574 and 0830803, a Xerox Faculty Research Award, a Google Faculty Research Award, an equipment grant from Intel, and an Okawa Foundation Research Grant. This material is based upon work supported by the Defense Advanced Research Projects Agency through the U.S. Office of Naval Research under Contract N00014-11-1-0389. The views expressed are those of the author and do not reflect the official policy or position of the Department of Defense, the National Science Foundation, or the U.S. Government.

*** UT Austin. bwaters@cs.utexas.edu. Supported by NSF CNS-0915361 and CNS-0952692, CNS-1228599 DARPA via Office of Naval Research under Contract N00014-11-1-0382, DARPA N11AP20006, the Alfred P. Sloan Fellowship, and Microsoft Faculty Fellowship, and Packard Foundation Fellowship. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Department of Defense or the U.S. Government.

[LATV12,Bra12] “have been simplified enough so that their description can fit, well, in a blog post” [BB12b,BB12a]. In this paper, we try to make FHE even simpler.

1.1 Previous FHE Schemes Based on Learning with Errors

Currently, perhaps the simplest leveled¹ FHE scheme based on the learning with errors (LWE) assumption [Reg05] is by Brakerski [Bra12]. In fact, Barak and Brakerski do give a remarkably clear exposition of this scheme in a blog post [BB12a]. However, while the scheme’s key generation, encryption, decryption, and homomorphic addition procedures are easy to describe, they note that “multiplication is more tricky”.

In Brakerski’s scheme, similar to previous FHE schemes based on LWE [BV11b,BGV12], the ciphertext \mathbf{c} and secret key \mathbf{s} are n -dimensional vectors whose dot product $\langle \mathbf{c}, \mathbf{s} \rangle \approx \mu$ equals the message μ , up to some small “error” that is removed by rounding. Homomorphic multiplication uses an identity regarding dot products of tensor products of vectors: namely, $\langle \mathbf{u}_1 \otimes \mathbf{u}_2, \mathbf{v}_1 \otimes \mathbf{v}_2 \rangle = \langle \mathbf{u}_1, \mathbf{v}_1 \rangle \cdot \langle \mathbf{u}_2, \mathbf{v}_2 \rangle$. Thus, if ciphertexts \mathbf{c}_1 and \mathbf{c}_2 satisfy $\langle \mathbf{c}_1, \mathbf{s} \rangle \approx \mu_1$ and $\langle \mathbf{c}_2, \mathbf{s} \rangle \approx \mu_2$, then $\langle \mathbf{c}_1 \otimes \mathbf{c}_2, \mathbf{s} \otimes \mathbf{s} \rangle \approx \mu_1 \cdot \mu_2$, where $\mathbf{c}_1 \otimes \mathbf{c}_2$ is interpreted as the new ciphertext and $\mathbf{s} \otimes \mathbf{s}$ as the new secret key, each having dimension $\Theta(n^2)$. Since multiplying-by-tensoring blows up the ciphertext size, it can only be used for a constant number of steps. For efficiency, the evaluator must *relinearize* [BV11b] the ciphertext after tensoring. Relinearization is a procedure that takes the *long* ciphertext that encrypts $\mu_1 \cdot \mu_2$ under the long key $\mathbf{s} \otimes \mathbf{s}$, and compresses it into a *normal-sized* n -dimensional ciphertext that encrypts $\mu_1 \cdot \mu_2$ under a normal-sized n -dimensional key \mathbf{s}' . To relinearize, the evaluator multiplies the long ciphertext vector by a special $n \times \Theta(n^2)$ relinearization matrix. This relinearization matrix is part of the “evaluation key” that the evaluator must obtain from the public key to perform homomorphic evaluation.

The relinearization step [BV11b] is ingenious and is perhaps the main insight that led to FHE based on LWE. However, relinearization is not particularly natural, nor is it easy to give an intuitive description of how and why it works. Moreover, relinearization is expensive. Each relinearization matrix has size $\Omega(n^3)$, and the public key must contain L of them to evaluate circuits of maximum multiplicative depth L . Computationally, relinearization requires $\Omega(n^3)$ operations, where each operation has cost polynomial in L .

This situation raises the question: Can we construct a LWE-based FHE scheme with a *natural* multiplication procedure? For ciphertexts c_1 and c_2 , can we construct a scheme where homomorphic addition and multiplication are just

¹ “Leveled” FHE is a relaxation of “pure” FHE [Gen09]. For fixed parameters, a pure FHE scheme can evaluate arbitrary circuits. In a leveled FHE scheme, the parameters of the scheme may depend on the *depth*, but not the *size*, of the circuits that the scheme can evaluate. We focus on leveled FHE schemes, and typically omit the term “leveled”. One can transform our leveled FHE schemes to pure ones by using Gentry’s bootstrapping theorem and assuming “circular security” [Gen09].

$c_1 + c_2$ and $c_1 \cdot c_2$, where ‘+’ and ‘·’ are natural algebraic operations over some ring, and where the new ciphertexts have the “same form” as the old ones; for example, $c_1 \cdot c_2$ is not a “long” ciphertext? Can we eliminate the need for an “evaluation key” in general, and the relinearization matrices in particular? If so, LWE-based FHE might become easier to explain. If we can simplify LWE-based FHE while also improving its efficiency and supporting new applications, then even better.

1.2 Our Results

Our main results are:

- **Conceptually simpler FHE based on LWE:** We fully describe our scheme here in the Introduction, and think our new approach will prove valuable pedagogically and theoretically.
- **Asymptotically faster FHE based on LWE:** We eliminate relinearization and the large relinearization matrices, with their $\Omega(n^3)$ complexity. Instead, ciphertexts are matrices that are added and multiplied naturally. In principle, matrix multiplication uses sub-cubic computation: e.g., Strassen and Williams achieved $n^{2.807}$ and $n^{2.3727}$ respectively [Str69,Wil12].
- **Identity-based FHE:** We solve an open problem mentioned in previous works [Nac10,GHV10,Bra12,CHT13] – namely, to construct an identity-based FHE scheme, in which there are no user-specific keys that must be obtained by the encrypter or evaluator. Informally speaking, in an identity-based FHE scheme, a user that has only the public parameters should be able to perform *both* encryption and homomorphism operations. The homomorphism operations should allow a user to take two ciphertexts encrypted to the same target identity, and homomorphically combine them to produce another ciphertext under the same target identity. Previously, only “weak” identity-based FHE schemes were known, where the evaluator needs a user-specific evaluation key, and thus the homomorphism is *not* exploitable by a user that only has the public parameters. Our scheme solves the problem by eliminating evaluation keys entirely.

We obtain our identity-based FHE scheme by presenting a “compiler” that transforms any LWE-based IBE scheme in the literature that satisfies certain properties, into a fully homomorphic identity-based encryption scheme. Several LWE-based IBE schemes in the literature satisfy the properties needed for our compiler [GPV08,ABB10a,ABB10b,CHKP10].

- **Attribute-based FHE:** Recently Gorbunov et al. [GVW13] constructed an attribute-based encryption (ABE) for circuits based on LWE. Our compiler for LWE-based IBE also works for their ABE scheme, with relatively minor modifications. We obtain an ABE scheme in which messages encrypted under the same index can be processed homomorphically without any evaluation

key in a polynomial depth circuit, and still be decrypted by any party that was entitled to decrypt the original ciphertexts.²

Our FHE scheme retains advantages of other LWE-based FHE schemes, such as making bootstrapping optional [BGV12], (with bootstrapping) basing security on LWE for quasi-polynomial factors versus sub-exponential factors [BGV12], eliminating “modulus switching” [Bra12], and basing security directly on the hardness of classical GapSVP [Bra12].

We do not want to oversell our asymptotic result; we now provide some additional context: In general, FHE schemes based on LWE have much worse performance (certainly asymptotically) than schemes based on ring LWE (RLWE) [LPR10,BV11a,GHS12a], and even RLWE-based schemes cannot yet be considered practical [GHS12b]. Moreover, sub-cubic matrix multiplication algorithms may not beat cubic ones by much in practice. Rather, we view our asymptotic result mainly as evidence of how fundamentally new our techniques are. We note that it is straightforward to construct an RLWE-based version of our scheme, but its performance is worse than the best known RLWE-based schemes [BGV12,Bra12,GHS12a,GHS12b] by log factors. On the other hand, while our techniques may not reduce evaluation complexity as much as we would like, they reduce the space complexity significantly (from quasi-cubic to quasi-quadratic), which is a significant issue for LWE-based FHE schemes in practice.

As with all current FHE schemes without bootstrapping, the parameters and per-gate complexity of evaluation depend on the multiplicative depth L of the circuit. “Bootstrapping” [Gen09], together with an assumption of circular security, remains the only known way of making these performance metrics independent of L , and while the overhead of bootstrapping is high, it becomes an attractive option once L passes some threshold. However, our scheme loses some of its advantages once bootstrapping is used. First, to apply bootstrapping, the evaluator needs to obtain the user’s secret key encrypted under its public key – in effect, an evaluation key – and therefore we no longer achieve identity-based/attribute-based FHE in this context. Second, this encrypted secret key has quasi-cubic size in our scheme, and while this can be mitigated by public key compression techniques [CNT12], it eliminates the space complexity advantages of our scheme. Essentially, bootstrapping returns us to the realm of “unnatural” operations, with all of its disadvantages. It remains a fascinating open problem to find some “natural” alternative to bootstrapping, and (relatedly) to achieve “pure” FHE without an assumption of circular security.

1.3 An Overview of Our FHE Scheme

Our main insight is that we can achieve LWE-based homomorphic encryption where homomorphic addition and multiplication correspond directly to matrix addition and multiplication.

² Independently, Garg et al. [GGH⁺13b] also recently constructed an ABE scheme for circuits using multilinear maps [GGH13a,CLT13], but our techniques do not work as effectively with their scheme.

Homomorphic Operations Let us skip key generation and encryption for the moment, and jump directly to the homomorphic operations (and decryption).

In our scheme, for some modulus q and dimension parameter N to be specified later, a ciphertext C is a $N \times N$ matrix over \mathbb{Z}_q , with “small” entries (much smaller than q) and the secret key \mathbf{v} is a N -dimensional vector over \mathbb{Z}_q with at least one “big” coefficient v_i . We restrict the message μ to be a “small” integer. We say C encrypts μ when $C \cdot \mathbf{v} = \mu \cdot \mathbf{v} + \mathbf{e}$, where \mathbf{e} is a “small” error vector. To decrypt, we extract the i -th row C_i from C , compute $x \leftarrow \langle C_i, \mathbf{v} \rangle = \mu \cdot v_i + e_i$, and output $\mu = \lfloor x/v_i \rfloor$. In a nutshell, the essence of our scheme is that the secret key \mathbf{v} is an *approximate eigenvector* of the ciphertext matrix C , and the message μ is the *eigenvalue*.

Now, let us see why matrix addition and multiplication are correct homomorphic operations. Suppose C_1 and C_2 encrypt μ_1 and μ_2 in that $C_i \cdot \mathbf{v} = \mu_i \cdot \mathbf{v} + \mathbf{e}_i$ for small \mathbf{e}_i . Let $C^+ = C_1 + C_2$ and $C^\times = C_1 \cdot C_2$. For addition, we have $C^+ \cdot \mathbf{v} = (\mu_1 + \mu_2) \cdot \mathbf{v} + (\mathbf{e}_1 + \mathbf{e}_2)$, where the error likely has grown a little, as usual in FHE schemes. But assuming the error is still “small”, the sum of the ciphertext matrices encrypts the sum of the messages. For multiplication, we have

$$\begin{aligned} C^\times \cdot \mathbf{v} &= C_1 \cdot (\mu_2 \cdot \mathbf{v} + \mathbf{e}_2) = \mu_2 \cdot (\mu_1 \cdot \mathbf{v} + \mathbf{e}_1) + C_1 \cdot \mathbf{e}_2 = \mu_1 \cdot \mu_2 \cdot \mathbf{v} + \mu_2 \cdot \mathbf{e}_1 + C_1 \cdot \mathbf{e}_2 \\ &= \mu_1 \cdot \mu_2 \cdot \mathbf{v} + \textit{small} \end{aligned}$$

where the final error vector is hopefully “small”, since μ_2 , C_1 , \mathbf{e}_1 , and \mathbf{e}_2 are all small. If so, the product of the ciphertext matrices encrypts the product of the messages. Interestingly, $C_2 \cdot C_1$ is also an encryption of $\mu_1 \cdot \mu_2$, even though matrix multiplication is not commutative.

To simplify further, it might be helpful to imagine an *error-free* version of the scheme, where $C_i \cdot \mathbf{v} = \mu_i \cdot \mathbf{v}$ *exactly*. In this case, the key \mathbf{v} is an (exact) eigenvector of ciphertext matrices, and the message μ_i is the eigenvalue. In general, if matrices C_1 and C_2 have a common eigenvector \mathbf{v} with eigenvalues μ_1 and μ_2 , then $C_1 \cdot C_2$ and $C_2 \cdot C_1$ have eigenvector \mathbf{v} with eigenvalue $\mu_1 \cdot \mu_2$.

Of course, in our scheme, the secret key \mathbf{v} is only an *approximate* eigenvector, not an *exact* one. Introducing error is necessary to base the security of our scheme on LWE. The cost of making \mathbf{v} only an *approximate* eigenvector is that certain terms in our scheme must be “small” to ensure that homomorphic operations do not disrupt the essential form of the ciphertexts. We call our new approach to LWE-based (homomorphic) encryption the *approximate eigenvector* method.

Bounding the Error and Somewhat Homomorphic Encryption Although we have not fully specified the scheme, let us go ahead and estimate how homomorphic it is. The scheme above works correctly until the coefficients of the error vector begin to approach q in magnitude. How many homomorphic operations can we perform before that happens?

Suppose C_1 and C_2 are B -bounded ciphertexts, in the sense that μ_i and the coefficients of C_i and \mathbf{e}_i all have magnitude at most some bound B . Then, C^+ is $2B$ -bounded, and C^\times is $(N + 1)B^2$ -bounded. In short, the error level grows

worse than B^{2^L} , *doubly exponentially with the multiplicative depth L* of the circuit being evaluated. Alternatively, if one wants to consider the *degree* (rather than *depth*) of functions that can be evaluated, if we evaluate a multivariate polynomial $P(x_1, \dots, x_t)$ of total degree d , on B -bounded ciphertexts as input, the final ciphertext is $|P|(N+1)^{d-1}B^d$ -bounded, where $|P|$ is the ℓ_1 -norm of P 's coefficient vector. Taking q to comfortably exceed this bound, we (roughly) can evaluate polynomials of degree $\log_{NB} q$. Since q/B must be subexponential (at most) in N for security reasons, our scheme-so-far can only evaluate polynomials of (sublinear) polynomial degree in N (only logarithmic depth). In short, our scheme-so-far is a *somewhat homomorphic encryption* (SWHE) scheme [Gen09] that can evaluate log-depth or polynomial degree. Though not yet fully homomorphic, it is by far the most homomorphic LWE-based encryption scheme that uses only “natural” homomorphic operations.

Flattening Ciphertexts and Fully Homomorphic Encryption To obtain a leveled FHE scheme that can evaluate circuits of polynomial *depth* without bootstrapping or techniques like relinearization, we need to ensure better bounds on the growth of the error. Let us say that a ciphertext C is B -strongly-bounded if its associated μ and the coefficients of C all have magnitude *at most 1*, while the coefficients of its \mathbf{e} all have magnitude at most B . If we evaluate a NAND gate on B -strongly-bounded ciphertexts C_1, C_2 to obtain a new ciphertext $C_3 \leftarrow I_N - C_1 \cdot C_2$ (where I_N is the N -dimensional identity matrix), then the message remains in $\{0, 1\}$, and the coefficients of C_3 's error vector have magnitude at most $(N+1)B$. If we could somehow additionally ensure that C_3 's coefficients have magnitude at most 1 so that strong-boundedness is preserved, then we could evaluate a circuit of *depth L* while keeping the error magnitude at most $(N+1)^L B$. Setting q/B to be subexponential in N , we could evaluate a circuit of polynomial *depth* rather than merely polynomial *degree*. In short, we would have a leveled FHE scheme.

Here we describe a operation called ciphertext *flattening* that keeps ciphertexts strongly bounded, so that we obtain leveled FHE.

Flattening uses some simple transformations from [BV11b, BGV12, Bra12] that modify vectors without affecting dot products. Let \mathbf{a}, \mathbf{b} be vectors of some dimension k over \mathbb{Z}_q . Let $\ell = \lfloor \log_2 q \rfloor + 1$ and $N = k \cdot \ell$. Let $\text{BitDecomp}(\mathbf{a})$ be the N -dimensional vector $(a_{1,0}, \dots, a_{1,\ell-1}, \dots, a_{k,0}, \dots, a_{k,\ell-1})$, where $a_{i,j}$ is the j -th bit in a_i 's binary representation, bits ordered least significant to most significant. For $\mathbf{a}' = (a_{1,0}, \dots, a_{1,\ell-1}, \dots, a_{k,0}, \dots, a_{k,\ell-1})$, let $\text{BitDecomp}^{-1}(\mathbf{a}') = (\sum 2^j \cdot a_{1,j}, \dots, \sum 2^j \cdot a_{k,j})$ be the inverse of BitDecomp , but well-defined even when the input is not a 0/1 vector. For N -dimensional \mathbf{a}' , let $\text{Flatten}(\mathbf{a}') = \text{BitDecomp}(\text{BitDecomp}^{-1}(\mathbf{a}'))$, a N -dimensional vector with 0/1 coefficients. When A is a matrix, let $\text{BitDecomp}(A)$, BitDecomp^{-1} , or $\text{Flatten}(A)$ be the matrix formed by applying the operation to each row of A separately. Finally, let $\text{Powersof2}(\mathbf{b}) = (b_1, 2b_1, \dots, 2^{\ell-1}b_1, \dots, b_k, 2b_k, \dots, 2^{\ell-1}b_k)$, a N -dimensional vector. Here are some obvious facts:

- $\langle \text{BitDecomp}(\mathbf{a}), \text{Powersof2}(\mathbf{b}) \rangle = \langle \mathbf{a}, \mathbf{b} \rangle$.

- For any N -dimensional \mathbf{a}' , $\langle \mathbf{a}', \text{Powersof2}(\mathbf{b}) \rangle = \langle \text{BitDecomp}^{-1}(\mathbf{a}'), \mathbf{b} \rangle = \langle \text{Flatten}(\mathbf{a}'), \text{Powersof2}(\mathbf{b}) \rangle$.

An interesting feature of `Flatten` is that it makes the coefficients of a vector or matrix *small*, without affecting its product with `Powersof2(b)`, and without knowing \mathbf{b} .

To facilitate ciphertext flattening, we give a special form to our secret key \mathbf{v} . Specifically, we set $\mathbf{v} = \text{Powersof2}(\mathbf{s})$ for some secret vector \mathbf{s} (to be specified later). This form is consistent with our earlier requirement that \mathbf{v} have some big coefficient v_i for decryption; indeed, since \mathbf{v} 's coefficients go up by $\lfloor \log_2 q \rfloor$ powers of 2, it *must* have a big coefficient suitable to recover $\mu \in \{0, 1\}$.

Now, for any $N \times N$ matrix C , we have $\text{Flatten}(C) \cdot \mathbf{v} = C \cdot \mathbf{v}$. So, after we compute an initial ciphertext $C_3 \leftarrow I_N - C_1 \cdot C_2$ for the NAND gate, we set $C^{\text{NAND}} = \text{Flatten}(C_3)$ to obtain a ciphertext that has 0/1 coefficients and is strongly bounded. Thus, we obtain leveled FHE without relinearization, under a fixed approximate eigenvector secret key.

Key Generation, Encryption, and Reduction to LWE Let us finally circle back to key generation and encryption. We want to base security on LWE. So, for key generation, we generate an LWE instance. For suitable parameters $q, n, m = O(n \log q)$, an LWE instance over \mathbb{Z}_q consists of a $m \times (n + 1)$ matrix A such that there exists a $(n + 1)$ -dimensional vector \mathbf{s} whose first coefficient is 1 where $\mathbf{e} = A \cdot \mathbf{s}$ is a “small” error vector. (See Section 2 for a formal definition of LWE.) In our scheme, A is public and \mathbf{s} is secret. We set our approximate eigenvector to be $\mathbf{v} = \text{Powersof2}(\mathbf{s})$, a vector of dimension $N = (n + 1) \cdot \ell$ for $\ell = \lfloor \log_2 q \rfloor + 1$.

To encrypt $\mu \in \mathbb{Z}_q$, the encrypter generates a random $N \times m$ matrix R with 0/1 entries, and sets $C = \text{Flatten}(\mu \cdot I_N + \text{BitDecomp}(R \cdot A))$, where I_N is the N -dimensional identity matrix. Since `Flatten` does not affect the product with \mathbf{v} , we have:

$$C \cdot \mathbf{v} = \mu \cdot \mathbf{v} + \text{BitDecomp}(R \cdot A) \cdot \mathbf{v} = \mu \cdot \mathbf{v} + R \cdot A \cdot \mathbf{s} = \mu \cdot \mathbf{v} + \textit{small}$$

`Flatten` ensures that the coefficients of C are small, and therefore that C has the proper form of a ciphertext that permits our homomorphic operations. Decryption works as mentioned previously.

To show that security is based on LWE, it is now enough to show that C is statistically independent of μ when A is a uniformly random $m \times (n + 1)$ matrix over \mathbb{Z}_q . Let $C' = \text{BitDecomp}^{-1}(C)$. Recall that C is `Flatten`'d, and so $C = \text{Flatten}(C) = \text{BitDecomp}(C')$. Therefore, C reveals nothing more than C' . But $C' = \text{BitDecomp}^{-1}(\mu \cdot I_N) + R \cdot A$, and $R \cdot A$ is statistically uniform by the leftover hash lemma when $m = O(n \log q)$ is chosen appropriately.

1.4 Roadmap

After finishing some preliminaries in Section 2, we describe our new FHE construction more formally in Section 3. In Section 4, we provide an overview of our identity-based and attribute-based FHE schemes.

2 Preliminaries

2.1 The Learning with Errors (LWE) Problem and GapSVP

The learning with errors (LWE) problem was introduced by Regev [Reg05].

Definition 1 (LWE). For security parameter λ , let $n = n(\lambda)$ be an integer dimension, let $q = q(\lambda) \geq 2$ be an integer, and let $\chi = \chi(\lambda)$ be a distribution over \mathbb{Z} . The $\text{LWE}_{n,q,\chi}$ problem is to distinguish the following two distributions: In the first distribution, one samples (\mathbf{a}_i, b_i) uniformly from \mathbb{Z}_q^{n+1} . In the second distribution, one first draws $\mathbf{s} \leftarrow \mathbb{Z}_q^n$ uniformly and then samples $(\mathbf{a}_i, b_i) \in \mathbb{Z}_q^{n+1}$ by sampling $\mathbf{a}_i \leftarrow \mathbb{Z}_q^n$ uniformly, $e_i \leftarrow \chi$, and setting $b_i = \langle \mathbf{a}_i, \mathbf{s} \rangle + e_i$. The $\text{LWE}_{n,q,\chi}$ assumption is that the $\text{LWE}_{n,q,\chi}$ problem is infeasible.

Sometimes it is convenient to view the vectors $b_i \parallel \mathbf{a}_i$ as the rows of a matrix A , and to redefine \mathbf{s} as $(1, -\mathbf{s})$. Then, either A is uniform, or there is a vector \mathbf{s} whose first coefficient is 1 such that $A \cdot \mathbf{s} = \mathbf{e}$, where the coefficients of \mathbf{e} come from the distribution χ .

For lattice dimension parameter n and number d , GapSVP_γ is the problem of distinguishing whether a n -dimensional lattice has a vector shorter than d or no vector shorter than $\gamma(n) \cdot d$. The two theorems below capture reductions, quantum and classical, from GapSVP to LWE for certain parameters. We state the result in terms of B -bounded distributions.

Definition 2 (B -bounded distributions). A distribution ensemble $\{\chi_n\}_{n \in \mathbb{N}}$, supported over the integers, is called B -bounded if $\Pr_{e \leftarrow \chi_n}[|e| > B] = \text{negl}(n)$.

Theorem 1 ([Reg05, Pei09, MM11, MP12], stated as Corollary 2.1 from [Bra12]). Let $q = q(n) \in \mathbb{N}$ be either a prime power or a product of small (size $\text{poly}(n)$) distinct primes, and let $B \geq \omega(\log n) \cdot \sqrt{n}$. Then there exists an efficient sampleable B -bounded distribution χ such that if there is an efficient algorithm that solves the average-case LWE problem for parameters n, q, χ , then:

- There is an efficient quantum algorithm that solves $\text{GapSVP}_{\tilde{O}(nq/B)}$ on any n -dimensional lattice.
- If $q \geq \tilde{O}(2^{n/2})$, then there is an efficient classical algorithm for $\text{GapSVP}_{\tilde{O}(nq/B)}$ on any n -dimensional lattice.

In both cases, if one also considers distinguishers with sub-polynomial advantage, then we require $B \geq \tilde{O}(n)$ and the resulting approximation factor is slightly larger than $\tilde{O}(n^{1.5}q/B)$.

Theorem 2 (Informal Theorem 1.1 of [BLP⁺13]). Solving n -dimensional LWE with $\text{poly}(n)$ modulus implies an equally efficient solution to a worst-case lattice problem (e.g., GapSVP) in dimension \sqrt{n} .

2.2 Identity-Based and Attribute-Based Homomorphic Encryption

In a homomorphic encryption scheme $\text{HE} = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Eval})$, the message space is some ring, and Eval homomorphically evaluates arithmetic circuits

over this ring (with addition and multiplication gates). We omit formal definitions and theorems regarding homomorphic encryption, which can be found in referenced papers.

An identity-based HE scheme $\text{IBHE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec}, \text{Eval})$ has all of the properties of a normal IBE scheme $\text{IBE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ [Sha84,BF03]. Setup generates master keys (MSK, MPK) , $\text{KeyGen}(\text{MSK}, \text{ID})$ outputs a secret key sk_{ID} for identity ID , $\text{Enc}(\text{MPK}, \text{ID}, m)$ outputs an encryption c of m under ID , and $\text{Dec}(\text{sk}_{\text{ID}}, c)$ decrypts c (if it is under ID). Standard security properties apply. For example, an IBE scheme is expected to be *collusion-resistant* – in particular, the adversary can ask for many secret keys, as long as the challenge ciphertext is under an unqueried identity.

For some function family \mathcal{F} , IBHE 's procedure $c \leftarrow \text{Eval}(\text{MPK}, \text{ID}, f, c_1, \dots, c_t)$ homomorphically evaluates any $f \in \mathcal{F}$ on ciphertexts $\{c_i \leftarrow \text{Enc}(\text{MPK}, \text{ID}, m_i)\}$ under the same ID . Ultimately, $\text{Dec}(\text{sk}_{\text{ID}}, c) = f(m_1, \dots, m_t)$. We define identity-based (leveled) *fully* homomorphic encryption (IBFHE) in the expected way.

The definition of IBHE can be extended to a multi-identity setting – specifically, Eval could work over ciphertexts under multiple identities. For security to make sense, Dec would require cooperation of all parties whose identities were used in Eval . In this paper, we restrict our attention to the single-identity setting.

An attribute-based HE scheme $\text{ABHE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec}, \text{Eval})$ has all of the properties of a normal ABE scheme $\text{ABE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ [SW05,GPSW06]. For some relation R , some function family \mathcal{F} and any $f \in \mathcal{F}$, and any ciphertexts $\{c_i \leftarrow \text{Enc}(\text{MPK}, x, m_i)\}$ encrypted under common index x , the ciphertext $c \leftarrow \text{Eval}(\text{MPK}, x, f, c_1, \dots, c_t)$ can be decrypted (to $f(m_1, \dots, m_t)$) using a key sk_y for any y for which $R(x, y) = 1$. In an ABE scheme *for circuits*, R can be a circuit of polynomial depth. We define attribute-based (leveled) *fully* homomorphic encryption (ABFHE) in the expected way.

Similar to IBHE , ABHE can be extended so that Eval operates on ciphertexts under multiple indices x_1, \dots, x_k . Regarding decryption, there are different possibilities. For example, the result can only be decrypted using some sk_y for which $R(x_1, y) = \dots = R(x_k, y) = 1$. Alternatively, the result can be cooperatively decrypted using $\text{sk}_{y_1}, \dots, \text{sk}_{y_\ell}$ such that for every x_i there is some j such that $R(x_i, y_j) = 1$. We restrict our attention to the single-index setting.

2.3 Other Preliminaries

For n , q , and $\ell = \lfloor \log q \rfloor + 1$, we define the procedures BitDecomp , BitDecomp^{-1} , Flatten and Powersof2 as described in the Introduction. I_N denotes the N -dimensional identity matrix.

3 Our LWE-Based FHE Scheme

3.1 Basic Encryption Scheme

Here, we formally describe our basic encryption scheme (without homomorphic operations). This description matches the description outlined in the Introduction. In our description, we split up KeyGen into three parts Setup , SecretKeyGen

and `PublicKeyGen`. We provide two decryption algorithms `Dec` and `MPDec`. `Dec` is sufficient to recover the message μ when it is in a small space (e.g., $\{0, 1\}$). `MPDec` is an algorithm by Micciancio and Peikert [MP12] that can recover any $\mu \in \mathbb{Z}_q$.

- `Setup`($1^\lambda, 1^L$): Choose a modulus q of $\kappa = \kappa(\lambda, L)$ bits, lattice dimension parameter $n = n(\lambda, L)$, and error distribution $\chi = \chi(\lambda, L)$ appropriately for LWE that achieves at least 2^λ security against known attacks. Also, choose parameter $m = m(\lambda, L) = O(n \log q)$. Let $params = (n, q, \chi, m)$. Let $\ell = \lceil \log q \rceil + 1$ and $N = (n + 1) \cdot \ell$.
- `SecretKeyGen`($params$): Sample $\mathbf{t} \leftarrow \mathbb{Z}_q^n$. Output $sk = \mathbf{s} \leftarrow (1, -t_1, \dots, -t_n) \in \mathbb{Z}_q^{n+1}$. Let $\mathbf{v} = \text{Powersof2}(\mathbf{s})$.
- `PublicKeyGen`($params, sk$): Generate a matrix $B \leftarrow \mathbb{Z}_q^{m \times n}$ uniformly and a vector $\mathbf{e} \leftarrow \chi^m$. Set $\mathbf{b} = B \cdot \mathbf{t} + \mathbf{e}$. Set A to be the $(n + 1)$ -column matrix consisting of \mathbf{b} followed by the n columns of B . Set the public key $pk = A$. (*Remark:* Observe that $A \cdot \mathbf{s} = \mathbf{e}$.)
- `Enc`($params, pk, \mu$): To encrypt a message $\mu \in \mathbb{Z}_q$, sample a uniform matrix $R \in \{0, 1\}^{N \times m}$ and output the ciphertext C given below.

$$C = \text{Flatten}(\mu \cdot I_N + \text{BitDecomp}(R \cdot A)) \in \mathbb{Z}_q^{N \times N}.$$

- `Dec`($params, sk, C$): Observe that the first ℓ coefficients of \mathbf{v} are $1, 2, \dots, 2^{\ell-1}$. Among these coefficients, let $v_i = 2^i$ be in $(q/4, q/2]$. Let C_i be the i -th row of C . Compute $x_i \leftarrow \langle C_i, \mathbf{v} \rangle$. Output $\mu' = \lfloor x_i / v_i \rfloor$.
- `MPDec`($params, sk, C$) (for q a power of 2): Observe that $q = 2^{\ell-1}$ and the first $\ell - 1$ coefficients of \mathbf{v} are $1, 2, \dots, 2^{\ell-2}$, and therefore if $C \cdot \mathbf{v} = \mu \cdot \mathbf{v} + \text{small}$, then the first $\ell - 1$ coefficients of $C \cdot \mathbf{v}$ are $\mu \cdot \mathbf{g} + \text{small}$, where $\mathbf{g} = (1, 2, \dots, 2^{\ell-2})$. Recover $\text{LSB}(\mu)$ from $\mu \cdot 2^{\ell-2} + \text{small}$, then recover the next-least-significant-bit from $(\mu - \text{LSB}(\mu)) \cdot 2^{\ell-3} + \text{small}$, etc. (See [MP12] for the general q case.)

`Dec` is a `BitDecomp`'d version of Regev's decryption procedure, applied to one row of the ciphertext, which is a `BitDecomp`'d Regev ciphertext. (The extra rows will come into play in the homomorphic operations). If C is properly generated, then by the elementary properties of `BitDecomp` and `Powersof2`, we have

$$C \cdot \mathbf{v} = \mu \cdot \mathbf{v} + R \cdot A \cdot \mathbf{s} = \mu \cdot \mathbf{v} + R \cdot \mathbf{e}.$$

`Dec` only uses the i -th coefficient of the above expression, which is $x_i = \mu \cdot v_i + \langle R_i, \mathbf{e} \rangle$. The error $\langle R_i, \mathbf{e} \rangle$ has magnitude at most $\|\mathbf{e}\|_1$. In general, if $x_i = \mu \cdot v_i + e'$ for some error e' of magnitude at most $q/8$, and if $v_i \in (q/4, q/2]$, then x_i/v_i differs from μ by at most $(q/8)/v_i < 1/2$, and `Dec` uses rounding to output the correct value of μ . (In the basic scheme, we set χ to ensure that the error is so bounded with overwhelming probability.)

For the basic scheme (without homomorphic operations), one can take n to be quasi-linear in the security parameter λ and $\kappa = O(\log n)$. When allowing homomorphic operations, L represents the circuit complexity of the functions that the scheme correctly evaluates (roughly, L is the multiplicative depth); we provide a detailed analysis later of how L affects the other parameters.

3.2 Security

Observe that $\text{BitDecomp}^{-1}(C) = \mu \cdot G + R \cdot A$, where $G = \text{BitDecomp}^{-1}(I_N)$ is (the transpose of) the “primitive matrix” used by Micciancio and Peikert [MP12] in their construction of lattice trapdoors, and the rows of $R \cdot A$ are simply Regev [Reg05] encryptions of 0 for dimension n . Assuming $\text{BitDecomp}^{-1}(C)$ hides μ , C does as well, since C can be derived by applying BitDecomp . Thus, the security of our basic encryption scheme follows directly from the following lemma, used to prove the security of Regev’s encryption scheme [Reg05].

Lemma 1 (Implicit in [Reg05]). *Let $\text{params} = (n, q, \chi, m)$ be such that the $\text{LWE}_{n,q,\chi}$ assumption holds. Then, for $m = O(n \log q)$ and A, R as generated above, the joint distribution $(A, R \cdot A)$ is computationally indistinguishable from uniform over $\mathbb{Z}_q^{m \times (n+1)} \times \mathbb{Z}_q^{N \times (n+1)}$.*

Concretely, it suffices to take $m > 2n \log q$ [Reg05].

Like Brakerski [Bra12], we can also base security on GapSVP via a classical reduction from LWE [Pei09, BLP⁺13]. Specifically, Peikert [Pei09] gives a classical reduction of GapSVP to LWE, with the caveat that q must be exponential in n . Brakerski notes that exponential q was unusable in previous FHE schemes, since the ratio of q to the error level B of “fresh” ciphertexts cannot be exponential in n for security reasons (since LLL [LLL82] could be used to break such a scheme), and since B must be very small to permit many homomorphic operations. As in Brakerski’s scheme, we do not have that problem. The error bound B of fresh ciphertexts in our scheme does not need to be small to permit many homomorphic operations; we only require q/B to be sub-exponential, and we can therefore permit q to be exponential. Alternatively, we can use a sub-exponential q and base security on GapSVP via Brakerski et al.’s [BLP⁺13] recent classical reduction to LWE that works even for polynomial-size moduli, with the caveat that, in their reduction, the dimension of the GapSVP instances may be much smaller than the dimension of the LWE instances.

3.3 Homomorphic Operations

Recall that we already described some basic “homomorphic” operations BitDecomp , BitDecomp^{-1} , Flatten , and Powersof2 . These will play an important role in analyzing the homomorphic operations supported by our scheme. We remark that BitDecomp could alternatively decompose with respect to bases other than 2, or according to the Chinese Remainder Theorem.

We provide additional homomorphic operations MultConst , Add , Mult , NAND as follows.

- $\text{MultConst}(C, \alpha)$: To multiply a ciphertext $C \in \mathbb{Z}_q^{N \times N}$ by known constant $\alpha \in \mathbb{Z}_q$, set $M_\alpha \leftarrow \text{Flatten}(\alpha \cdot I_N)$ and output $\text{Flatten}(M_\alpha \cdot C)$. Observe that:

$$\begin{aligned} \text{MultConst}(C, \alpha) \cdot \mathbf{v} &= M_\alpha \cdot C \cdot \mathbf{v} = M_\alpha \cdot (\mu \cdot \mathbf{v} + \mathbf{e}) = \mu \cdot (M_\alpha \cdot \mathbf{v}) + M_\alpha \cdot \mathbf{e} \\ &= \alpha \cdot \mu \cdot \mathbf{v} + M_\alpha \cdot \mathbf{e} \end{aligned}$$

Thus, the error increases by a factor of at most N , regardless of what element $\alpha \in \mathbb{Z}_q$ is used for multiplication. As in “classical” additively homomorphic encryption schemes, we could alternatively perform multiplication-by-constant α by recursively applying **Add**. But this multiplies the error size by at least α , whereas **MultConst** increases the error by at most a factor of N , regardless of α . An example application of **MultConst** is that we can perform homomorphic fast Fourier transformations (FFTs) natively over \mathbb{Z}_q without error growth dependent on q . Previously, the error growth depended on the size of the field underlying the FFT [GHS12a,GHS12b], restricting the choice of field.

- **Add**(C_1, C_2): To add ciphertexts $C_1, C_2 \in \mathbb{Z}_q^{N \times N}$, output **Flatten**($C_1 + C_2$). The correctness of this operation is immediate. Note that the addition of messages is over the full base ring \mathbb{Z}_q .
- **Mult**(C_1, C_2): To multiply ciphertexts $C_1, C_2 \in \mathbb{Z}_q^{N \times N}$, output **Flatten**($C_1 \cdot C_2$). Observe that:

$$\begin{aligned} \text{Mult}(C_1, C_2) \cdot \mathbf{v} &= C_1 \cdot C_2 \cdot \mathbf{v} = C_1 \cdot (\mu_2 \cdot \mathbf{v} + \mathbf{e}_2) + \mu_2 \cdot (\mu_1 \mathbf{v} + \mathbf{e}_1) + C_1 \cdot \mathbf{e}_2 \\ &= \mu_1 \cdot \mu_2 \cdot \mathbf{v} + \mu_2 \cdot \mathbf{e}_1 + C_1 \cdot \mathbf{e}_2 \end{aligned}$$

As in **Add**, the multiplication operator is over the full base field \mathbb{Z}_q . In **Mult**, the new error depends on the old errors, the ciphertext C_1 , and the message μ_2 . The dependence on the old errors seems unavoidable (and normal for LWE-based HE schemes), and observe that C_1 contributes at most a factor N blowup of error, since all components of C_1 are restricted to $\{0, 1\}$. The error growth based on the message μ_2 , however, presents a concern. In general, we must address this concern by using homomorphic operations in a way that restricts the message space to small messages. One way to do this is to consider Boolean circuits using only **NAND** operations: this would restrict the message space to $\{0, 1\}$. We elaborate below.

- **NAND**(C_1, C_2): To **NAND** ciphertexts $C_1, C_2 \in \mathbb{Z}_q^{N \times N}$ that are known to encrypt messages $\mu_1, \mu_2 \in \{0, 1\}$, output **Flatten**($I_N - C_1 \cdot C_2$). Observe that:

$$\text{NAND}(C_1, C_2) \cdot \mathbf{v} = (I_N - C_1 \cdot C_2) \cdot \mathbf{v} = (1 - \mu_1 \cdot \mu_2) \cdot \mathbf{v} - \mu_2 \cdot \mathbf{e}_1 - C_1 \cdot \mathbf{e}_2$$

Note here that the **NAND** homomorphic operation maintains the invariant that if the input messages are in $\{0, 1\}$, then the output ciphertext will also encrypt $\{0, 1\}$, thus guaranteeing small messages. Note that since $\mu_2 \in \{0, 1\}$, the error is increased by a factor of at most $N + 1$.

Circuits. By iteratively applying the homomorphic operations above, different types of (bounded-depth) circuits may be homomorphically computed while maintaining correctness of decryption.

The simplest case to analyze is the case of Boolean circuits computed over encryptions of $\{0, 1\}$ values. In this case, the circuit can be converted to use only **NAND** gates, and through appropriate leveled application of the **NAND** homomorphic operation, the final ciphertext’s error will be bounded by $(N +$

$1)^L \cdot B$, where L is the NAND-depth of the circuit, and B is the original bound on the error of a fresh encryption of $\{0, 1\}$.

More generally, with more care, we may consider arithmetic circuits over \mathbb{Z}_q that make use of gates that perform addition, multiplication, or multiplication by a known constant. However, as we have seen in the case of multiplication gates, the error growth may depend on the values being encrypted in intermediate computations. One way to deal with this is to focus on situations where (1) all input values are known to encrypt values bounded by some value T , and (2) the arithmetic circuit is chosen to guarantee that all intermediate values are also bounded by T' whenever the circuit inputs are constrained to values bounded by T . In such a situation, the final ciphertext's error will be bounded by $(N+T')^L \cdot B$, where L is the depth of the arithmetic circuit, and B is the original bound on the error of fresh encryptions of values smaller than T . For example, in this way, we can homomorphically evaluate polynomials of degree d in this large-message-space variant when the initial messages are bounded by roughly $q^{1/d}$, achieving a scheme that is “somewhat homomorphic” [Gen09]. Another example application would be to convert encryptions of a polynomially bounded set of small values to encryptions of binary values, by using an appropriate arithmetic circuit for the conversion. Once converted to encryptions of binary values, a NAND-based Boolean circuit could be used for further computations.

3.4 Parameters, Performance and Optimizations

Suppose that Flatten'd ciphertexts C_1, C_2 encrypt $\mu_1, \mu_2 \in \{0, 1\}$ under approximate eigenvector \mathbf{v} with B -bounded error – that is, $C_i \cdot \mathbf{v} = \mu_i \cdot \mathbf{v} + \mathbf{e}_i$ where $|\mathbf{e}_i|_\infty \leq B$. Then $C^{\text{NAND}} \leftarrow \text{NAND}(C_1, C_2)$ encrypts $\text{NAND}(\mu_1, \mu_2) \in \{0, 1\}$ under \mathbf{v} with $(N+1)B$ -bounded error. As long as $q/B > 8(N+1)^L$, we can evaluate a depth- L circuit of NANDs over B -bounded ciphertexts to obtain a $q/8$ -bounded ciphertext, which Dec will decrypt correctly.

As in previous LWE-based FHE schemes, n (hence N) must increase linearly with $\log(q/B)$ to maintain fixed 2^λ security against known attacks, so q/B grows more like $\exp(L \log L)$. We will brush such issues under the rug and view n as a fixed parameter. Choosing χ so that B is not too large, and since in practice there is no reason to have $\kappa = \log q$ grow super-linearly with n , we have $\kappa = O(L \log N) = O(L(\log n + \log \kappa)) = O(L \log n)$, similar to [BGV12, Bra12]. Given that the NAND procedure is dominated by multiplication of two $N \times N$ matrices for $N = O(n\kappa) = \tilde{O}(nL)$, we have the following theorem to characterize the performance of our FHE scheme.

Theorem 3. *For dimension parameter n and depth parameter L , our FHE scheme evaluates depth- L circuits of NAND gates with $\tilde{O}((nL)^\omega)$ field operations per gate, where $\omega < 2.3727$ is the matrix multiplication exponent.*

This compares favorably with previous LWE-based FHE schemes, which all have at least $\tilde{O}(n^3L)$ field operations per gate [BV11b, BGV12, Bra12].

Theorem 3 hides some factors, both good and bad. On the good size, it hides the fact that ciphertext matrices in our scheme have 0/1 entries, and

therefore can be multiplied faster than if they were general matrices over \mathbb{Z}_q . In previous LWE-based FHE schemes, the field operations involve multiplying a small number with a general number of \mathbb{Z}_q , which has complexity $\tilde{O}(\kappa) = \tilde{O}(L)$. So, previous LWE-based FHE schemes have real complexity $\tilde{O}(n^3L^2)$ whereas ours remains $\tilde{O}((nL)^\omega)$. On the bad side, Theorem 3 hides logarithmic factors in the dimension of the ciphertext matrices, since $N = O(n\kappa) = O(nL \log n)$. We note that typically n will dominate L , since for very deep circuits, one would want to use Gentry’s bootstrapping technique [Gen09] to make the per-gate computation *independent* of L .

Since bootstrapping involves homomorphically evaluating the decryption function, and since Dec is essentially Regev decryption [Reg05], bootstrapping works as in previous LWE-based FHE schemes. In particular, we can use techniques from [BV11b] to reduce the dimension and modulus-size of the ciphertext before bootstrapping, so that the complexity of decryption (and hence bootstrapping) is completely independent of the depth L of the circuit that was evaluated to arrive at that ciphertext. Regev decryption can be evaluated in $O(\log n)$ depth. Due to the logarithmic depth, one can take q/B to be quasi-polynomial in n , and base security on LWE for quasi-polynomial factors.

4 Our Identity-Based and Attribute-Based FHE Schemes

Identity-based encryption (IBE) [Sha84,BF03] and attribute-based encryption (ABE) [SW05,GPSW06] are designed to provide more flexible access control of encrypted data than a traditional public key infrastructure. Traditionally, IBE and ABE do not offer any computation over the encrypted data. However, access control of encrypted data remains important even (or especially) when the data is encrypted homomorphically. (See [CHT13] for a nice discussion of applications.)

Unfortunately, while there are some IBE schemes that allow simple homomorphic operations [GHV10,CHT13], it has remained a stubborn open problem [Nac10,GHV10,Bra12,CHT13] to construct an IBE scheme that allows fully or even “somewhat” homomorphic encryption. Previously it was mentioned [Bra12,CHT13] that instead of building an FHE scheme on Regev’s encryption scheme as we do in Section 3, one can alternatively use the “dual-Regev” system [GPV08], for which it is known how to generate identity-based keys (see also [ABB10a,ABB10b,CHKP10]). However, making the encryption/decryption keys identity-based only solves half of the problem, and yields only a “weak” form of identity-based FHE. In all previous FHE schemes, there is also an “evaluation key” required for homomorphic evaluation. This evaluation key is user-specific and is not “identity-based”, in the sense that it cannot be computed non-interactively from the user’s identity. But having to *obtain* this evaluation key undermines the main appeal of IBE: its non-interactivity. Thus, identity-based FHE (IBFHE) has remained wide open, and attribute-based FHE (ABFHE) seems even more difficult to construct.

Interestingly, however, our new FHE scheme does not have evaluation keys. To perform evaluation, the evaluator only needs to know some basic parameters of the scheme (like n , m and ℓ).

The absence of evaluation keys allows us to construct the first IBFHE scheme. We describe a simple “compiler” that transforms any LWE-based IBE scheme (that satisfies certain natural properties) into a IBFHE. All LWE-based IBE schemes that we know of (e.g., [GPV08,ABB10a,ABB10b,CHKP10]) can be described so as to have the required properties.

1. **Property 1 (Ciphertext and decryption key vectors):** The decryption key for identity ID , and a ciphertext for ID , are vectors $\mathbf{s}_{ID}, \mathbf{c}_{ID} \in \mathbb{Z}_q^{n'}$ for some n' . The first coefficient of \mathbf{s}_{ID} is 1.
2. **Property 2 (Small Dot Product):** If \mathbf{c}_{ID} encrypts 0, then $\langle \mathbf{c}_{ID}, \mathbf{s}_{ID} \rangle$ is “small”.
3. **Property 3 (Security):** Encryptions of 0 are indistinguishable from uniform vectors over \mathbb{Z}_q (under LWE).

Theorem 4. *We can compile an IBE scheme E with the above properties into a related IBFHE scheme.*

Proof. The IBFHE uses E ’s Setup and KeyGen algorithms, supplementing E ’s MPK with the basic parameters for our FHE scheme (such as m, ℓ). Let $N = (n + 1) \cdot \ell$ for $\ell = \lfloor \log q \rfloor + 1$, as usual. To encrypt $\mu \in \{0, 1\}$, the encrypter generates N encryptions of 0 using $E.\text{Enc}$, sets C'_{ID} to be the $N \times (n + 1)$ matrix whose rows are these ciphertexts, and outputs $C_{ID} = \text{Flatten}(\mu \cdot I_N + \text{BitDecomp}(C'_{ID}))$. Suppose \mathbf{s}_{ID} is the decryption key for ID , as above, and let $\mathbf{v}_{ID} = \text{Powersof2}(\mathbf{s}_{ID})$. The decrypter runs our FHE decryption algorithm $\text{Dec}(\mathbf{v}_{ID}, C_{ID})$ to recover μ . Homomorphic operations are as in Section 3.3.

Decryption is correct, since $C_{ID} \cdot \mathbf{v}_{ID} = \mu \cdot \mathbf{v}_{ID} + C'_{ID} \cdot \mathbf{s}_{ID} = \mu \cdot \mathbf{v}_{ID} + \text{small}$, where $C'_{ID} \cdot \mathbf{s}_{ID}$ is a small vector by Property 2. In this setting Dec recovers $\mu \in \{0, 1\}$. Any adversary that breaks the semantic security of our IBFHE scheme can distinguish C'_{ID} from a uniform matrix over \mathbb{Z}_q , and therefore distinguish LWE by Property 3.

For ABFHE, our approach begins by re-interpreting the decryption process in the Gorbunov et al. (GVW) ABE scheme [GVW13]. To decrypt a ABE ciphertext under x with sk_y for which $R(x, y) = 1$, we view the decrypter as deriving a “sub-key” $\mathbf{s}_{x,y}$ associated to x . This sub-key will satisfy something similar to Property 2 above – i.e., if \mathbf{c}_x encrypts 0 under x , then $\langle \mathbf{c}_x, \mathbf{s}_{x,y} \rangle$ is “small”. Viewing GVW in this way allows us to apply our compiler above.

We provide more details of our identity-based and attribute-based FHE constructions in the full version of the paper.

Acknowledgments

We gratefully thank Shai Halevi for his collaboration on this work. We also thank Boaz Barak and the anonymous CRYPTO reviewers for their many helpful comments.

References

- [ABB10a] Shweta Agrawal, Dan Boneh, and Xavier Boyen. Efficient lattice (h)ibe in the standard model. In *EUROCRYPT*, pages 553–572, 2010.
- [ABB10b] Shweta Agrawal, Dan Boneh, and Xavier Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical ibe. In *CRYPTO*, pages 98–115, 2010.
- [BB12a] Boaz Barak and Zvika Brakerski. Building the swiss army knife. Windows on Theory Blog, <http://windowsontheory.org/2012/05/02/building-the-swiss-army-knife>, 2012.
- [BB12b] Boaz Barak and Zvika Brakerski. The swiss army knife of cryptography. Windows on Theory Blog, <http://windowsontheory.org/2012/05/01/the-swiss-army-knife-of-cryptography>, 2012.
- [BF03] Dan Boneh and Matt Franklin. Identity-based encryption from the Weil pairing. *SIAM J. of Computing*, 32(3):586–615, 2003. Extended abstract in Crypto’01.
- [BGV12] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. Fully homomorphic encryption without bootstrapping. In *Innovations in Theoretical Computer Science (ITCS’12)*, 2012. Available at <http://eprint.iacr.org/2011/277>.
- [BLP⁺13] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In *STOC*, pages 575–584, 2013.
- [Bra12] Zvika Brakerski. Fully homomorphic encryption without modulus switching from classical gapsvp. In *CRYPTO*, pages 868–886, 2012.
- [BV11a] Zvika Brakerski and Vinod Vaikuntanathan. Fully homomorphic encryption from ring-LWE and security for key dependent messages. In *CRYPTO*, volume 6841, page 501, 2011.
- [BV11b] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In *FOCS*, pages 97–106, 2011. References are to full version: <http://eprint.iacr.org/2011/344>.
- [CHKP10] David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. In *EUROCRYPT*, pages 523–552, 2010.
- [CHT13] Michael Clear, Arthur Hughes, and Hitesh Tewari. Homomorphic encryption with access policies: Characterization and new constructions. In *AFRICACRYPT*, pages 61–87, 2013.
- [CLT13] Jean-Sébastien Coron, Tancrede Lepoint, and Mehdi Tibouchi. Practical multilinear maps over the integers. *IACR Cryptology ePrint Archive*, 2013:183, 2013. To appear in CRYPTO 2013.
- [CMNT11] Jean-Sébastien Coron, Avradip Mandal, David Naccache, and Mehdi Tibouchi. Fully homomorphic encryption over the integers with shorter public keys. In *CRYPTO*, pages 487–504, 2011.
- [CNT12] Jean-Sébastien Coron, David Naccache, and Mehdi Tibouchi. Public key compression and modulus switching for fully homomorphic encryption over the integers. In *Advances in Cryptology - EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 446–464. Springer, 2012.
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, pages 169–178, 2009.
- [Gen10] Craig Gentry. Toward basing fully homomorphic encryption on worst-case hardness. In *CRYPTO*, pages 116–137, 2010.

- [GGH13a] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In *EUROCRYPT*, pages 1–17, 2013.
- [GGH⁺13b] Sanjam Garg, Craig Gentry, Shai Halevi, Amit Sahai, and Brent Waters. Attribute-based encryption for circuits from multilinear maps. *IACR Cryptology ePrint Archive*, 2013:128, 2013. To appear in CRYPTO 2013.
- [GH11a] Craig Gentry and Shai Halevi. Fully homomorphic encryption without squashing using depth-3 arithmetic circuits. In *FOCS*, pages 107–109, 2011.
- [GH11b] Craig Gentry and Shai Halevi. Implementing gentry’s fully-homomorphic encryption scheme. In *EUROCRYPT*, volume 6632 of *Lecture Notes in Computer Science*, pages 129–148. Springer, 2011.
- [GHS12a] Craig Gentry, Shai Halevi, and Nigel P. Smart. Fully homomorphic encryption with polylog overhead. In *EUROCRYPT*, pages 465–482, 2012.
- [GHS12b] Craig Gentry, Shai Halevi, and Nigel P. Smart. Homomorphic evaluation of the aes circuit. In *CRYPTO*, pages 850–867, 2012.
- [GHV10] Craig Gentry, Shai Halevi, and Vinod Vaikuntanathan. A simple bgn-type cryptosystem from lwe. In *EUROCRYPT*, pages 506–522, 2010.
- [GPSW06] Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. Attribute-based encryption for fine-grained access control of encrypted data. In *ACM CCS*, pages 89–98, 2006.
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206. ACM, 2008.
- [GVW13] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. In *STOC*, pages 545–554, 2013.
- [LATV12] Adriana López-Alt, Eran Tromer, and Vinod Vaikuntanathan. On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In *STOC*, pages 1219–1234, 2012.
- [LLL82] A. K. Lenstra, H. W. Lenstra, and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261:515–534, 1982. 10.1007/BF01457454.
- [LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In *EUROCRYPT*, pages 1–23, 2010.
- [MM11] Daniele Micciancio and Petros Mol. Pseudorandom knapsacks and the sample complexity of lwe search-to-decision reductions. In *CRYPTO*, pages 465–484, 2011.
- [MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *EUROCRYPT*, pages 700–718, 2012.
- [Nac10] David Naccache. Is theoretical cryptography any good in practice? Invited talk at Crypto/CHES 2010, available at <http://www.iacr.org/workshops/ches/ches2010>, 2010.
- [Pei09] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In *STOC*, pages 333–342, 2009.
- [RAD78] Ron Rivest, Leonard Adleman, and Michael L. Dertouzos. On data banks and privacy homomorphisms. In *Foundations of Secure Computation*, pages 169–180, 1978.
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, pages 84–93, 2005.
- [Sha84] Adi Shamir. Identity-based cryptosystems and signature schemes. In *CRYPTO*, pages 47–53, 1984.

- [Str69] Volker Strassen. Gaussian elimination is not optimal. *Numer. Math.*, 13:354–356, 1969.
- [SV10] Nigel P. Smart and Frederik Vercauteren. Fully homomorphic encryption with relatively small key and ciphertext sizes. In *Public Key Cryptography - PKC'10*, volume 6056 of *Lecture Notes in Computer Science*, pages 420–443. Springer, 2010.
- [SW05] Amit Sahai and Brent Waters. Fuzzy identity-based encryption. In *EUROCRYPT*, volume 3494 of *Lecture Notes in Computer Science*, pages 457–473, 2005.
- [vDGHV10] Marten van Dijk, Craig Gentry, Shai Halevi, and Vinod Vaikuntanathan. Fully homomorphic encryption over the integers. In *EUROCRYPT*, pages 24–43, 2010.
- [Wil12] Virginia Vassilevska Williams. Multiplying matrices faster than coppersmith-winograd. In *STOC*, pages 887–898, 2012.