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

Craig Gentry ${ }^{1, \star}$, Amit Sahai ${ }^{2, \star \star}$, and Brent Waters ${ }^{3, \star \star \star}$<br>${ }^{1}$ IBM Research<br>cbgentry@us.ibm.com<br>${ }^{2}$ UCLA<br>sahai@cs.ucla.edu<br>${ }^{3}$ UT Austin<br>bwaters@cs.utexas.edu


#### 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


[^0]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 identitybased 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 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 leveled ${ }^{1}$ 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 $\boldsymbol{c}$ and secret key $s$ are $n$-dimensional vectors whose dot product $\langle\boldsymbol{c}, \boldsymbol{s}\rangle \approx \mu$ equals the message $\mu$, up to some small "error" that is removed by rounding. Homomorphic multiplication uses an identity regarding dot products of tensor products of vectors: namely, $\left\langle\boldsymbol{u}_{1} \otimes \boldsymbol{u}_{2}, \boldsymbol{v}_{1} \otimes \boldsymbol{v}_{2}\right\rangle=$ $\left\langle\boldsymbol{u}_{1}, \boldsymbol{v}_{1}\right\rangle \cdot\left\langle\boldsymbol{u}_{2}, \boldsymbol{v}_{2}\right\rangle$. Thus, if ciphertexts $\boldsymbol{c}_{1}$ and $\boldsymbol{c}_{2}$ satisfy $\left\langle\boldsymbol{c}_{1}, \boldsymbol{s}\right\rangle \approx \mu_{1}$ and $\left\langle\boldsymbol{c}_{2}, \boldsymbol{s}\right\rangle \approx \mu_{2}$, then $\left\langle\boldsymbol{c}_{1} \otimes \boldsymbol{c}_{2}, \boldsymbol{s} \otimes \boldsymbol{s}\right\rangle \approx \mu_{1} \cdot \mu_{2}$, where $\boldsymbol{c}_{1} \otimes \boldsymbol{c}_{2}$ is interpreted as the new ciphertext and $s \otimes s$ as the new secret key, each having dimension $\Theta\left(n^{2}\right)$. 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 $\boldsymbol{s} \otimes \boldsymbol{s}$, and compresses it into a normal-sized n-dimensional ciphertext that encrypts $\mu_{1} \cdot \mu_{2}$ under a normal-sized $n$-dimensional key $s^{\prime}$. To relinearize, the evaluator multiplies the long ciphertext vector by a special $n \times \Theta\left(n^{2}\right)$ relinearization matrix.

[^1]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 $\Omega\left(n^{3}\right)$, and the public key must contain $L$ of them to evaluate circuits of maximum multiplicative depth $L$. Computationally, relinearization requires $\Omega\left(n^{3}\right)$ 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 $c_{1}+c_{2}$ and $c_{1} \cdot c_{2}$, where ' + ' and ' $\because$ ' 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 $\Omega\left(n^{3}\right)$ 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 identitybased FHE scheme, in which there are no user-specific keys that must be obtained by the encrypter or evaluator. Informally speaking, in an identitybased 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 userspecific 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 $\square^{2}$

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 quasiquadratic), 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/attributebased 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

[^2]"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.

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 $\boldsymbol{v}$ is a $N$-dimensional vector over $\mathbb{Z}_{q}$ with at least one "big" coefficient $v_{i}$. We restrict the message $\mu$ to be a "small" integer. We say $C$ encrypts $\mu$ when $C \cdot \boldsymbol{v}=\mu \cdot \boldsymbol{v}+\boldsymbol{e}$, where $\boldsymbol{e}$ is a "small" error vector. To decrypt, we extract the $i$-th row $C_{i}$ from $C$, compute $x \leftarrow\left\langle C_{i}, \boldsymbol{v}\right\rangle=\mu \cdot v_{i}+e_{i}$, and output $\mu=\left\lfloor x / v_{i}\right\rceil$. In a nutshell, the essence of our scheme is that the secret key $\boldsymbol{v}$ is an approximate eigenvector of the ciphertext matrix $C$, and the message $\mu$ is the eigenvalue.

Now, let us see why matrix addition and multiplication are correct homomorphic operations. Suppose $C_{1}$ and $C_{2}$ encrypt $\mu_{1}$ and $\mu_{2}$ in that $C_{i} \cdot \boldsymbol{v}=\mu_{i} \cdot \boldsymbol{v}+\boldsymbol{e}_{i}$ for small $e_{i}$. Let $C^{+}=C_{1}+C_{2}$ and $C^{\times}=C_{1} \cdot C_{2}$. For addition, we have $C^{+} \cdot \boldsymbol{v}=\left(\mu_{1}+\mu_{2}\right) \cdot \boldsymbol{v}+\left(\boldsymbol{e}_{1}+\boldsymbol{e}_{2}\right)$, 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

$$
\begin{aligned}
C^{\times} \cdot \boldsymbol{v} & =C_{1} \cdot\left(\mu_{2} \cdot \boldsymbol{v}+\boldsymbol{e}_{2}\right)=\mu_{2} \cdot\left(\mu_{1} \cdot \boldsymbol{v}+\boldsymbol{e}_{1}\right)+C_{1} \cdot \boldsymbol{e}_{2}=\mu_{1} \cdot \mu_{2} \cdot \boldsymbol{v}+\mu_{2} \cdot \boldsymbol{e}_{1}+C_{1} \cdot \boldsymbol{e}_{2} \\
& =\mu_{1} \cdot \mu_{2} \cdot \boldsymbol{v}+\text { small }
\end{aligned}
$$

where the final error vector is hopefully "small", since $\mu_{2}, C_{1}, \boldsymbol{e}_{1}$, and $\boldsymbol{e}_{2}$ 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 $C_{i} \cdot \boldsymbol{v}=\mu_{i} \cdot \boldsymbol{v}$ exactly. In this case, the key $\boldsymbol{v}$ is an (exact) eigenvector of ciphertext matrices, and the message $\mu_{i}$ is the eigenvalue. In general, if matrices $C_{1}$ and $C_{2}$ have a common eigenvector $\boldsymbol{v}$ with eigenvalues $\mu_{1}$ and $\mu_{2}$, then $C_{1} \cdot C_{2}$ and $C_{2} \cdot C_{1}$ have eigenvector $\boldsymbol{v}$ with eigenvalue $\mu_{1} \cdot \mu_{2}$.

Of course, in our scheme, the secret key $\boldsymbol{v}$ 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 $\boldsymbol{v}$ 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 $\mu_{i}$ and the coefficients of $C_{i}$ and $e_{i}$ all have magnitude at most some bound $B$. Then, $C^{+}$ is $2 B$-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 $P\left(x_{1}, \ldots, x_{t}\right)$ 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 $\ell_{1}$-norm of $P$ 's coefficient vector. Taking $q$ to comfortably exceed this bound, we (roughly) can evaluate polynomials of degree $\log _{N B} q$. 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 $\mu$ and the coefficients of $C$ all have magnitude at most 1 , while the coefficients of its $\boldsymbol{e}$ 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 $\boldsymbol{a}, \boldsymbol{b}$ be vectors of some dimension $k$ over $\mathbb{Z}_{q}$. Let $\ell=\left\lfloor\log _{2} q\right\rfloor+1$ and $N=k \cdot \ell$. Let BitDecomp $(\boldsymbol{a})$ be the $N$-dimensional vector $\left(a_{1,0}, \ldots, a_{1, \ell-1}, \ldots, a_{k, 0}, \ldots, a_{k, \ell-1}\right)$, where $a_{i, j}$ is the $j$-th bit in $a_{i}$ 's binary representation, bits ordered least significant to most significant. For $\boldsymbol{a}^{\prime}=\left(a_{1,0}, \ldots, a_{1, \ell-1}, \ldots, a_{k, 0}, \ldots, a_{k, \ell-1}\right)$, let $\operatorname{BitDecomp}^{-1}\left(\boldsymbol{a}^{\prime}\right)=$ $\left(\sum 2^{j} \cdot a_{1, j}, \ldots, \sum 2^{j} \cdot a_{k, j}\right)$ be the inverse of BitDecomp, but well-defined even
when the input is not a $0 / 1$ vector. For $N$-dimensional $\boldsymbol{a}^{\prime}$, let Flatten $\left(\boldsymbol{a}^{\prime}\right)=$ BitDecomp(BitDecomp ${ }^{-1}\left(\boldsymbol{a}^{\prime}\right)$ ), a $N$-dimensional vector with $0 / 1$ coefficients. When $A$ is a matrix, let $\operatorname{BitDecomp}(A), \operatorname{BitDecomp}{ }^{-1}$, or $\operatorname{Flatten}(A)$ be the matrix formed by applying the operation to each row of $A$ separately. Finally, let Powersof2 $(\boldsymbol{b})=\left(b_{1}, 2 b_{1}, \ldots, 2^{\ell-1} b_{1}, \ldots, b_{k}, 2 b_{k}, \ldots, 2^{\ell-1} b_{k}\right)$, a $N$-dimensional vector. Here are some obvious facts:
$-\langle\operatorname{BitDecomp}(\boldsymbol{a})$, Powersof2 $(\boldsymbol{b})\rangle=\langle\boldsymbol{a}, \boldsymbol{b}\rangle$.

- For any $N$-dimensional $\boldsymbol{a}^{\prime},\left\langle\boldsymbol{a}^{\prime}\right.$, Powersof2 $\left.(\boldsymbol{b})\right\rangle=\left\langle\operatorname{BitDecomp}^{-1}\left(\boldsymbol{a}^{\prime}\right), \boldsymbol{b}\right\rangle=$ $\left\langle\right.$ Flatten $\left(\boldsymbol{a}^{\prime}\right)$, Powersof2 $\left.(\boldsymbol{b})\right\rangle$.

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 $\boldsymbol{b}$.

To facilitate ciphertext flattening, we give a special form to our secret key $\boldsymbol{v}$. Specifically, we set $\boldsymbol{v}=\operatorname{Powersof2(s)~for~some~secret~vector~} \boldsymbol{s}$ (to be specified later). This form is consistent with our earlier requirement that $\boldsymbol{v}$ have some big coefficient $v_{i}$ for decryption; indeed, since $\boldsymbol{v}$ 's coefficients go up by $\left\lfloor\log _{2} q\right\rfloor$ 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 Flatten $(C) \cdot \boldsymbol{v}=C \cdot \boldsymbol{v}$. So, after we compute an initial ciphertext $C_{3} \leftarrow I_{N}-C_{1} \cdot C_{2}$ for the NAND gate, we set $C^{\text {NAND }}=$ Flatten $\left(C_{3}\right)$ 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 $s$ whose first coefficient is 1 where $\boldsymbol{e}=A \cdot \boldsymbol{s}$ is a "small" error vector. (See Section 2 for a formal definition of LWE.) In our scheme, $A$ is public and $s$ is secret. We set our approximate eigenvector to be $\boldsymbol{v}=\operatorname{Powersof} 2(\boldsymbol{s})$, a vector of dimension $N=(n+1) \cdot \ell$ for $\ell=\left\lfloor\log _{2} q\right\rfloor+1$.

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

$$
C \cdot \boldsymbol{v}=\mu \cdot \boldsymbol{v}+\operatorname{BitDecomp}(R \cdot A) \cdot \boldsymbol{v}=\mu \cdot \boldsymbol{v}+R \cdot A \cdot s=\mu \cdot \boldsymbol{v}+\operatorname{small}
$$

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 $\mu$ when $A$ is a uniformly random $m \times(n+1)$ matrix over $\mathbb{Z}_{q}$. Let $C^{\prime}=\operatorname{BitDecomp}^{-1}(C)$. Recall that $C$ is Flatten'd, and so
$C=\operatorname{Flatten}(C)=\operatorname{BitDecomp}\left(C^{\prime}\right)$. Therefore, $C$ reveals nothing more than $C^{\prime}$. But $C^{\prime}=\operatorname{BitDecomp}^{-1}\left(\mu \cdot I_{N}\right)+R \cdot A$, 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 Section4, 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 $\lambda$, 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 $\mathrm{LWE}_{n, q, \chi}$ problem is to distinguish the following two distributions: In the first distribution, one samples $\left(\boldsymbol{a}_{i}, b_{i}\right)$ uniformly from $\mathbb{Z}_{q}^{n+1}$. In the second distribution, one first draws $\boldsymbol{s} \leftarrow \mathbb{Z}_{q}^{n}$ uniformly and then samples $\left(\boldsymbol{a}_{i}, b_{i}\right) \in \mathbb{Z}_{q}^{n+1}$ by sampling $\boldsymbol{a}_{i} \leftarrow \mathbb{Z}_{q}^{n}$ uniformly, $e_{i} \leftarrow \chi$, and setting $b_{i}=\left\langle\boldsymbol{a}_{i}, \boldsymbol{s}\right\rangle+e_{i}$. The $\mathrm{LWE}_{n, q, \chi}$ assumption is that the $\mathrm{LWE}_{n, q, \chi}$ problem is infeasible.

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

For lattice dimension parameter $n$ and number $d$, GapSVP $\gamma_{\gamma}$ 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 $\left\{\chi_{n}\right\}_{n \in \mathbb{N}}$, supported over the integers, is called $B$-bounded if $\operatorname{Pr}_{e \leftarrow \chi_{n}}[|e|>B]=\operatorname{neg} \mid(n)$.

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 poly $(n))$ distinct primes, and let $B \geq \omega(\log n) \cdot \sqrt{n}$. Then there exists an efficient sampleable $B$-bounded distribution $\chi$ such that if there is an efficient algorithm that solves the average-case LWE problem for parameters $n, q, \chi$, then:

- There is an efficient quantum algorithm that solves $G a p S V P_{\tilde{O}(n q / B)}$ on any n-dimensional lattice.
- If $q \geq \tilde{O}\left(2^{n / 2}\right)$, then there is an efficient classical algorithm for $\operatorname{GapSVP}_{\tilde{O}(n q / B)}$ 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 $\tilde{O}\left(n^{1.5} q / B\right)$.

Theorem 2 (Informal Theorem 1.1 of $\left[\overline{\mathbf{B L P}^{+} \mathbf{1 3}}\right]$ ). Solving $n$-dimensional LWE with 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 $\mathrm{HE}=($ KeyGen, Enc, Dec, 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 IBHE = (Setup, KeyGen, Enc, Dec, Eval) has all of the properties of a normal IBE scheme IBE = (Setup, KeyGen, Enc, Dec) Sha84, BF03. Setup generates master keys (MSK, MPK), KeyGen(MSK, ID) outputs a secret key sk ${ }_{I D}$ for identity ID, Enc(MPK, ID, $m$ ) outputs an encryption $c$ of $m$ under ID, and $\operatorname{Dec}\left(\mathrm{sk}_{\mathrm{ID}}, c\right)$ decrypts $c$ (if it is under ID). Standard security properties apply. For example, an IBE scheme is expected to be collusionresistant - 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 $c \leftarrow \operatorname{Eval}\left(\mathrm{MPK}, \mathrm{ID}, f, c_{1}, \ldots, c_{t}\right)$ homomorphically evaluates any $f \in \mathcal{F}$ on ciphertexts $\left\{c_{i} \leftarrow \operatorname{Enc}\left(\mathrm{MPK}, \mathrm{ID}, m_{i}\right)\right\}$ under the same ID. Ultimately, $\operatorname{Dec}\left(\mathrm{sk}_{\mathrm{ID}}, c\right)=f\left(m_{1}, \ldots, m_{t}\right)$. We define identitybased (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 ABHE = (Setup, KeyGen, Enc, Dec, Eval) has all of the properties of a normal ABE scheme $\mathrm{ABE}=$ (Setup, KeyGen, Enc, Dec) SW05, GPSW06. For some relation $R$, some function family $\mathcal{F}$ and any $f \in \mathcal{F}$, and any ciphertexts $\left\{c_{i} \leftarrow \operatorname{Enc}\left(\mathrm{MPK}, x, m_{i}\right)\right\}$ encrypted under common index $x$, the ciphertext $c \leftarrow \operatorname{Eval}\left(\mathrm{MPK}, x, f, c_{1}, \ldots, c_{t}\right)$ can be decrypted (to $f\left(m_{1}, \ldots, m_{t}\right)$ ) using a key $\mathrm{sk}_{y}$ 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 attributebased (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 $x_{1}, \ldots, x_{k}$. Regarding decryption, there are different possibilities. For example, the result can only be decrypted using some $\mathrm{sk}_{y}$ for which $R\left(x_{1}, y\right)=\cdots=R\left(x_{k}, y\right)=1$. Alternatively, the result can be cooperatively decrypted using $\mathrm{sk}_{y_{1}}, \ldots, \mathrm{sk}_{y_{\ell}}$ such that for every $x_{i}$ there is some $j$ such that $R\left(x_{i}, y_{j}\right)=1$. 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 $\mu$ 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 $\left(1^{\lambda}, 1^{L}\right)$ : 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 $2^{\lambda}$ security against known attacks. Also, choose parameter $m=m(\lambda, L)=O(n \log q)$. Let params $=(n, q, \chi, m)$. Let $\ell=\lfloor\log q\rfloor+1$ and $N=(n+1) \cdot \ell$.
- SecretKeyGen(params): Sample $\boldsymbol{t} \leftarrow \mathbb{Z}_{q}^{n}$. Output $s k=s \leftarrow$ $\left(1,-t_{1}, \ldots,-t_{n}\right) \in \mathbb{Z}_{q}^{n+1}$. Let $\boldsymbol{v}=$ Powersof2 $(s)$.
- PublicKeyGen (params, sk): Generate a matrix $B \leftarrow \mathbb{Z}_{q}^{m \times n}$ uniformly and a vector $\boldsymbol{e} \leftarrow \chi^{m}$. Set $\boldsymbol{b}=B \cdot \boldsymbol{t}+\boldsymbol{e}$. Set $A$ to be the $(n+1)$-column matrix consisting of $\boldsymbol{b}$ followed by the $n$ columns of $B$. Set the public key $p k=A$. (Remark: Observe that $A \cdot \boldsymbol{s}=\boldsymbol{e}$.)
- 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.

$$
C=\operatorname{Flatten}\left(\mu \cdot I_{N}+\operatorname{BitDecomp}(R \cdot A)\right) \in \mathbb{Z}_{q}^{N \times N}
$$

- Dec $($ params $, s k, C)$ : Observe that the first $\ell$ coefficients of $\boldsymbol{v}$ are $1,2, \ldots, 2^{\ell-1}$. 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 $x_{i} \leftarrow\left\langle C_{i}, \boldsymbol{v}\right\rangle$. Output $\mu^{\prime}=\left\lfloor x_{i} / v_{i}\right\rceil$.
- MPDec (params, sk,C) (for $q$ a power of 2 ): Observe that $q=2^{\ell-1}$ and the first $\ell-1$ coefficients of $\boldsymbol{v}$ are $1,2, \ldots, 2^{\ell-2}$, and therefore if $C \cdot \boldsymbol{v}=$ $\mu \cdot \boldsymbol{v}+$ small, then the first $\ell-1$ coefficients of $C \cdot \boldsymbol{v}$ are $\mu \cdot \boldsymbol{g}+$ small, where $\boldsymbol{g}=\left(1,2, \ldots, 2^{\ell-2}\right)$. Recover $\operatorname{LSB}(\mu)$ from $\mu \cdot 2^{\ell-2}+$ small, then recover the next-least-significant-bit from $(\mu-\operatorname{LSB}(\mu)) \cdot 2^{\ell-3}+$ 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

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

Dec only uses the $i$-th coefficient of the above expression, which is $x_{i}=\mu \cdot v_{i}+$ $\left\langle R_{i}, \boldsymbol{e}\right\rangle$. The error $\left\langle R_{i}, \boldsymbol{e}\right\rangle$ has magnitude at most $\|\boldsymbol{e}\|_{1}$. In general, if $x_{i}=\mu \cdot v_{i}+e^{\prime}$ for some error $e^{\prime}$ of magnitude at most $q / 8$, and if $v_{i} \in(q / 4, q / 2]$, then $x_{i} / v_{i}$ differs from $\mu$ by at most $(q / 8) / v_{i}<1 / 2$, and Dec uses rounding to output the correct value of $\mu$. (In the basic scheme, we set $\chi$ 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 $\lambda$ 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 BitDecomp ${ }^{-1}(C)=\mu \cdot G+R \cdot A$, where $G=\operatorname{BitDecomp}^{-1}\left(I_{N}\right)$ 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 BitDecomp ${ }^{-1}(C)$ hides $\mu, C$ 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 params $=(n, q, \chi, m)$ be such that the $\mathrm{LWE}_{n, q, \chi}$ 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 $m>2 n \log q$ Reg05.
Like Brakerski Bra12, we can also base security on GapSVP via a classical reduction from LWE [Pei09, $\mathrm{BLP}^{+} 13$. 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 $\mathrm{BLP}^{+} 13$ 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.

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

$$
\begin{aligned}
\operatorname{MultConst}(C, \alpha) \cdot \boldsymbol{v} & =M_{\alpha} \cdot C \cdot \boldsymbol{v}=M_{\alpha} \cdot(\mu \cdot \boldsymbol{v}+\boldsymbol{e})=\mu \cdot\left(M_{\alpha} \cdot \boldsymbol{v}\right)+M_{\alpha} \cdot \boldsymbol{e} \\
& =\alpha \cdot \mu \cdot \boldsymbol{v}+M_{\alpha} \cdot \boldsymbol{e}
\end{aligned}
$$

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 $\alpha$ by recursively applying Add. But this multiplies the error size by at least $\alpha$, whereas MultConst increases the error by at most a factor of $N$, regardless of $\alpha$. 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.
$-\operatorname{Add}\left(C_{1}, C_{2}\right)$ : To add ciphertexts $C_{1}, C_{2} \in \mathbb{Z}_{q}^{N \times N}$, output Flatten $\left(C_{1}+C_{2}\right)$. The correctness of this operation is immediate. Note that the addition of messages is over the full base ring $\mathbb{Z}_{q}$.

- Mult $\left(C_{1}, C_{2}\right)$ : To multiply ciphertexts $C_{1}, C_{2} \in \mathbb{Z}_{q}^{N \times N}$, output Flatten $\left(C_{1}\right.$. $C_{2}$ ). Observe that:

$$
\begin{aligned}
\operatorname{Mult}\left(C_{1}, C_{2}\right) \cdot \boldsymbol{v} & =C_{1} \cdot C_{2} \cdot \boldsymbol{v}=C_{1} \cdot\left(\mu_{2} \cdot \boldsymbol{v}+\boldsymbol{e}_{2}\right)+\mu_{2} \cdot\left(\mu_{1} \boldsymbol{v}+\boldsymbol{e}_{1}\right)+C_{1} \cdot \boldsymbol{e}_{2} \\
& =\mu_{1} \cdot \mu_{2} \cdot \boldsymbol{v}+\mu_{2} \cdot \boldsymbol{e}_{1}+C_{1} \cdot \boldsymbol{e}_{2}
\end{aligned}
$$

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 $\mu_{2}$. 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 $\mu_{2}$, 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.
$-\operatorname{NAND}\left(C_{1}, C_{2}\right)$ : 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 $\left(I_{N}-C_{1} \cdot C_{2}\right)$. Observe that: $\operatorname{NAND}\left(C_{1}, C_{2}\right) \cdot \boldsymbol{v}=\left(I_{N}-C_{1} \cdot C_{2}\right) \cdot \boldsymbol{v}=\left(1-\mu_{1} \cdot \mu_{2}\right) \cdot \boldsymbol{v}-\mu_{2} \cdot \boldsymbol{e}_{1}-C_{1} \cdot \boldsymbol{e}_{2}$

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 encryption of $\{0,1\}$, 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 $T^{\prime}$ whenever the circuit inputs are constrained to values bounded by $T$. In such a situation, the final ciphertext's error will be bounded by $\left(N+T^{\prime}\right)^{L} \cdot B$, 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-messagespace 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 $\boldsymbol{v}$ with $B$-bounded error - that is, $C_{i} \cdot \boldsymbol{v}=\mu_{i} \cdot \boldsymbol{v}+\boldsymbol{e}_{i}$ where $\left|\boldsymbol{e}_{i}\right|_{\infty} \leq B$. Then $C^{\text {NAND }} \leftarrow \operatorname{NAND}\left(C_{1}, C_{2}\right)$ encrypts $\operatorname{NAND}\left(\mu_{1}, \mu_{2}\right) \in\{0,1\}$ under $\boldsymbol{v}$ 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 $2^{\lambda}$ 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 $\chi$ 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 $N=O(n \kappa)=\tilde{O}(n L)$, 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 $\tilde{O}\left((n L)^{\omega}\right)$ 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 $\tilde{O}\left(n^{3} L\right)$ 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 $\tilde{O}\left(n^{3} L^{2}\right)$ whereas ours remains $\tilde{O}\left((n L)^{\omega}\right)$. On the bad side, Theorem 3 hides logarithmic factors in the dimension of the ciphertext matrices, since $N=O(n \kappa)=O(n L \log n)$. 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 userspecific and is not "identity-based", in the sense that it cannot be computed noninteractively 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 $\ell$ ).

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 $s_{\mathrm{ID}}, \boldsymbol{c}_{\mathrm{ID}} \in \mathbb{Z}_{q}^{n^{\prime}}$ for some $n^{\prime}$. The first coefficient of $s_{\text {ID }}$ is 1 .
2. Property 2 (Small Dot Product): If $\boldsymbol{c}_{\mathrm{ID}}$ encrypts 0 , then $\left\langle\boldsymbol{c}_{\mathrm{ID}}, \boldsymbol{s}_{\mathrm{ID}}\right\rangle$ 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 $m, \ell)$. 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.Enc, sets $C_{\mathrm{ID}}^{\prime}$ to be the $N \times(n+1)$ matrix whose rows are these ciphertexts, and outputs $C_{\mathrm{ID}}=\operatorname{Flatten}\left(\mu \cdot I_{N}+\operatorname{BitDecomp}\left(C_{\mathrm{ID}}^{\prime}\right)\right)$. Suppose $\boldsymbol{s}_{\mathrm{ID}}$ is the decryption key for ID, as above, and let $\boldsymbol{v}_{\mathrm{ID}}=\operatorname{Powersof} 2\left(\boldsymbol{s}_{\mathrm{ID}}\right)$. The decrypter runs our FHE decryption algorithm $\operatorname{Dec}\left(\boldsymbol{v}_{\mathrm{ID}}, C_{\mathrm{ID}}\right)$ to recover $\mu$. Homomorphic operations are as in Section 3.3

Decryption is correct, since $C_{\mathrm{ID}} \cdot \boldsymbol{v}_{\mathrm{ID}}=\mu \cdot \boldsymbol{v}_{\mathrm{ID}}+C_{\mathrm{ID}}^{\prime} \cdot \boldsymbol{s}_{\mathrm{ID}}=\mu \cdot \boldsymbol{v}_{\mathrm{ID}}+$ small, where $C_{\mathrm{ID}}^{\prime} \cdot s_{\mathrm{ID}}$ 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 $C_{\mathrm{ID}}^{\prime}$ 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 $y_{y}$ for which $R(x, y)=1$, we view the decrypter as deriving a
"sub-key" $\boldsymbol{s}_{x, y}$ associated to $x$. This sub-key will satisfy something similar to Property 2 above - i.e., if $\boldsymbol{c}_{x}$ encrypts 0 under $x$, then $\left\langle\boldsymbol{c}_{x}, \boldsymbol{s}_{x, y}\right\rangle$ 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] Agrawal, S., Boneh, D., Boyen, X.: Efficient lattice (h)ibe in the standard model. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 553-572. Springer, Heidelberg (2010)
[ABB10b] Agrawal, S., Boneh, D., Boyen, X.: Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 98-115. Springer, Heidelberg (2010)
[BB12a] Barak, B., Brakerski, Z.: Building the swiss army knife. Windows on Theory Blog (2012), http://windowsontheory.org/2012/05/02/ building-the-swiss-army-knife
[BB12b] Barak, B., Brakerski, Z.: The swiss army knife of cryptography. Windows on Theory Blog (2012), http://windowsontheory.org/2012/05/01/ the-swiss-army-knife-of-cryptography
[BF03] Boneh, D., Franklin, M.: Identity-based encryption from the Weil pairing. SIAM J. of Computing 32(3), 586-615 (2003); Extended abstract in Kilian, J. (ed.): CRYPTO 2001. LNCS, vol. 2139, pp. 586-615. Springer, Heidelberg (2001)
[BGV12] Brakerski, Z., Gentry, C., Vaikuntanathan, V.: Fully homomorphic encryption without bootstrapping. In: Innovations in Theoretical Computer Science, ITCS 2012 (2012), http://eprint.iacr.org/2011/277
$\left[\mathrm{BLP}^{+} 13\right]$ Brakerski, Z., Langlois, A., Peikert, C., Regev, O., Stehlé, D.: Classical hardness of learning with errors. In: STOC, pp. 575-584 (2013)
[Bra12] Brakerski, Z.: Fully homomorphic encryption without modulus switching from classical gapSVP. In: Safavi-Naini, R. (ed.) CRYPTO 2012. LNCS, vol. 7417, pp. 868-886. Springer, Heidelberg (2012)
[BV11a] Brakerski, Z., Vaikuntanathan, V.: Fully homomorphic encryption from ring-LWE and security for key dependent messages. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 505-524. Springer, Heidelberg (2011)
[BV11b] Brakerski, Z., Vaikuntanathan, V.: Efficient fully homomorphic encryption from (standard) LWE. In: FOCS, pp. 97-106 (2011), http://eprint.iacr.org/2011/344
[CHKP10] Cash, D., Hofheinz, D., Kiltz, E., Peikert, C.: Bonsai trees, or how to delegate a lattice basis. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 523-552. Springer, Heidelberg (2010)
[CHT13] Clear, M., Hughes, A., Tewari, H.: Homomorphic encryption with access policies: Characterization and new constructions. In: Youssef, A., Nitaj, A., Hassanien, A.E. (eds.) AFRICACRYPT 2013. LNCS, vol. 7918, pp. 61-87. Springer, Heidelberg (2013)
[CLT13] Coron, J.-S., Lepoint, T., Tibouchi, M.: Practical multilinear maps over the integers. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 476-493. Springer, Heidelberg (2013)
[CMNT11] Coron, J.-S., Mandal, A., Naccache, D., Tibouchi, M.: Fully homomorphic encryption over the integers with shorter public keys. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 487-504. Springer, Heidelberg (2011)
[CNT12] Coron, J.-S., Naccache, D., Tibouchi, M.: Public key compression and modulus switching for fully homomorphic encryption over the integers. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 446-464. Springer, Heidelberg (2012)
[Gen09] Gentry, C.: Fully homomorphic encryption using ideal lattices. In: STOC, pp. 169-178 (2009)
[Gen10] Gentry, C.: Toward basing fully homomorphic encryption on worst-case hardness. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 116137. Springer, Heidelberg (2010)
[GGH13a] Garg, S., Gentry, C., Halevi, S.: Candidate multilinear maps from ideal lattices. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 1-17. Springer, Heidelberg (2013)
$\left[\mathrm{GGH}^{+}\right.$13b] Garg, S., Gentry, C., Halevi, S., Sahai, A., Waters, B.: Attribute-based encryption for circuits from multilinear maps. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 479-499. Springer, Heidelberg (2013)
[GH11a] Gentry, C., Halevi, S.: Fully homomorphic encryption without squashing using depth-3 arithmetic circuits. In: FOCS, pp. 107-109 (2011)
[GH11b] Gentry, C., Halevi, S.: Implementing gentry's fully-homomorphic encryption scheme. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 129-148. Springer, Heidelberg (2011)
[GHS12a] Gentry, C., Halevi, S., Smart, N.P.: Fully homomorphic encryption with polylog overhead. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 465-482. Springer, Heidelberg (2012)
[GHS12b] Gentry, C., Halevi, S., Smart, N.P.: Homomorphic evaluation of the AES circuit. In: Safavi-Naini, R. (ed.) CRYPTO 2012. LNCS, vol. 7417, pp. 850-867. Springer, Heidelberg (2012)
[GHV10] Gentry, C., Halevi, S., Vaikuntanathan, V.: A simple BGN-type cryptosystem from LWE. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 506-522. Springer, Heidelberg (2010)
[GPSW06] Goyal, V., Pandey, O., Sahai, A., Waters, B.: Attribute-based encryption for fine-grained access control of encrypted data. In: ACM CCS, pp. 89-98 (2006)
[GPV08] Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: STOC, pp. 197-206. ACM (2008)
[GVW13] Gorbunov, S., Vaikuntanathan, V., Wee, H.: Attribute-based encryption for circuits. In: STOC, pp. 545-554 (2013)
[LATV12] López-Alt, A., Tromer, E., Vaikuntanathan, V.: On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In: STOC, pp. 1219-1234 (2012)
[LLL82] Lenstra, A.K., Lenstra, H.W., Lovász, L.: Factoring polynomials with rational coefficients. Mathematische Annalen 261, 515-534 (1982), doi:10.1007/BF01457454
[LPR10] Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 1-23. Springer, Heidelberg (2010)
[MM11] Micciancio, D., Mol, P.: Pseudorandom knapsacks and the sample complexity of LWE search-to-decision reductions. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 465-484. Springer, Heidelberg (2011)
[MP12] Micciancio, D., Peikert, C.: Trapdoors for lattices: Simpler, tighter, faster, smaller. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 700-718. Springer, Heidelberg (2012)
[Nac10] Naccache, D.: Is theoretical cryptography any good in practice? Invited talk at Crypto/CHES 2010 (2010),
http://www.iacr.org/workshops/ches/ches2010
[Pei09] Peikert, C.: Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In: STOC, pp. 333-342 (2009)
[RAD78] Rivest, R., Adleman, L., Dertouzos, M.L.: On data banks and privacy homomorphisms. In: Foundations of Secure Computation, pp. 169-180 (1978)
[Reg05] Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: STOC, pp. 84-93 (2005)
[Sha84] Shamir, A.: Identity-based cryptosystems and signature schemes. In: Blakely, G.R., Chaum, D. (eds.) CRYPTO 1984. LNCS, vol. 196, pp. 47-53. Springer, Heidelberg (1985)
[Str69] Strassen, V.: Gaussian elimination is not optimal. Numer. Math. 13, 354-356 (1969)
[SV10] Smart, N.P., Vercauteren, F.: Fully homomorphic encryption with relatively small key and ciphertext sizes. In: Nguyen, P.Q., Pointcheval, D. (eds.) PKC 2010. LNCS, vol. 6056, pp. 420-443. Springer, Heidelberg (2010)
[SW05] Sahai, A., Waters, B.: Fuzzy identity-based encryption. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 457-473. Springer, Heidelberg (2005)
[vDGHV10] van Dijk, M., Gentry, C., Halevi, S., Vaikuntanathan, V.: Fully homomorphic encryption over the integers. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. EUROCRYPT, pp. 24-43. Springer, Heidelberg (2010)
[Wil12] Williams, V.V.: Multiplying matrices faster than coppersmith-winograd. In: STOC, pp. 887-898 (2012)


[^0]:    * 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.
    ** 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.
    *** 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.

[^1]:    ${ }^{1}$ "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.

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

