

Homomorphic Encryption Using RNS Theorem

Shishir Kumar Choudhary and Ankit Gupta,
8th sem, CSE, YDIT, Bengaluru.

Email: shishirkundan@gmail.com and greamsx@gmail.com

ABSTRACT

The prospect of outsourcing an increasing amount of data storage and management to cloud services raises many new privacy concerns for individuals and businesses alike. The privacy concerns can be satisfactorily addressed if users encrypt the data they send to the cloud. If the encryption scheme is homomorphic, the cloud can still perform meaningful computations on the data, even though it is encrypted. In fact, we now know a number of constructions of fully homomorphic encryption schemes that allow arbitrary computation on encrypted data. In the last two years, solutions for fully homomorphic encryption schemes have been proposed and improved upon, but it is hard to ignore the elephant in the room, namely efficiency { can homomorphic encryption ever be efficient enough to be practical? Certainly, it seems that all known fully homomorphic encryption schemes have a long way to go before they can be used in practice. Given this state of affairs, our contribution is two-fold. First, we exhibit a number of real-world applications, in the medical, financial, and the advertising domains, which require only that the encryption scheme is "somewhat" homomorphic. Somewhat homomorphic encryption schemes, which support a limited number of homomorphic operations, can be much faster, and more compact than fully homomorphic encryption schemes. Secondly, we show a proof-of-concept implementation of the recent somewhat homomorphic encryption scheme of Brakerski and Vaikuntanathan whose security relies on the ring learning with errors" (Ring LWE) problem. The scheme is very efficient, and has reasonably short ciphertexts. Our unoptimized implementation in magma enjoys comparable efficiency to even optimized pairing-based schemes with the same level of security and homomorphic capacity. We also show a number of application-specific optimizations to the encryption scheme, most notably the ability to efficiently convert between different message encodings in a cipher-text.

Keywords: FHE, cloud storage homomorphic encryption scheme, SHE

1. INTRODUCTION

The development of cloud storage and computing platforms allows users to outsource storage and computation on their data, and allows businesses to offload the task of maintaining data-centers. However, concerns over loss of privacy and business value of private data is an overwhelming barrier to the adoption of cloud services by consumers and businesses alike. An excellent way to assuage these privacy concerns is to store all data in the cloud encrypted, and perform computations on encrypted

data. To this end, we need an encryption scheme that allows meaningful computation on encrypted data, namely a homomorphic encryption scheme. Homomorphic encryption schemes that allow simple computations on encrypted data have been known for a long time. For example, the encryption systems of Goldwasser and Micali [GM82], El Gamal [El-84] and Paillier [Pai99] support either adding or multiplying encrypted ciphertexts, but not both operations at the same time. Boneh, Goh and Nissim [3] were the first to construct a scheme capable of performing both operations at the same time { their scheme handles an arbitrary number of additions and just one multiplication. More recently, in a breakthrough work, Gentry [Gen09, Gen10] constructed a fully homomorphic encryption scheme (FHE) capable of evaluating an arbitrary number of additions and multiplications (and thus, compute any function) on encrypted data. The main point of this paper is to show to what extent current schemes can actually be used to compute functions of practical interest on encrypted data. Since the appearance of Gentry's scheme, there has been much informal discussion in the industry as to whether fully homomorphic encryption is implementable and practical. While the initial solution may not have been practical, subsequent developments produced other schemes [DGHV10, SV10, SS10] leading up to the most recent solutions of Brakerski and Vaikuntanathan [BV11b, BV11a], an implementation of which we consider in this paper. The scheme is efficient and simple, produces short ciphertexts, and its security is based on the "ring learning with errors" (Ring LWE) problem [LPR10]. While the

performance of the state-of-the art FHE implementations is itself a question of interest (and has indeed been considered recently in, e.g., [GH11, SV11]), our focus here is on describing concrete practical applications and concrete useful functions to be computed, most of which require only a limited number of multiplications of ciphertexts (as well as a possibly very large number of additions of ciphertexts). For these applications, it is enough to consider an implementation of a somewhat homomorphic encryption" (SHE) scheme, namely, one which allows a fixed number of multiplications of ciphertexts. These SHE schemes are building blocks for the FHE schemes of, e.g., [Gen09, DGHV10, BV11b, BV11a], and provide much better efficiency guarantees than their fully homomorphic counter-parts.

1.1 Practical Applications of Homomorphic Encryption

We describe a number of concrete applications and functions to be implemented to provide cloud services in the medical, financial, and advertising sectors. (We provide a sketch of the applications here, and refer the reader to Section 2 for detailed descriptions). For a cloud service managing electronic medical records (EMR), consider a futuristic scenario where devices continuously collect vital health information, and stream them to a server who then computes some statistics (over these measurements, and over the course of time) and presumably decides on the course of treatment (e.g., whether the dosage of medicine should be changed). The volume of the data involved is large, and thus, the patient presumably does not want to store and manage all this data locally; she may prefer to use cloud storage and computation. To protect patient privacy, all the data is uploaded in encrypted form, and thus the cloud must perform operations on the encrypted data in order to return (encrypted) alerts, predictions, or summaries of the results to the patient. We describe scenarios such as the above, which require computing simple statistical functions such as the mean, standard deviation, as well as logistical regressions that are typically used for prediction of likelihoods of certain desirable or undesirable outcomes. For these functions, it suffices to have a somewhat homomorphic encryption system which computes many additions and a small number of multiplications on ciphertexts: for example, averages require no multiplications, standard deviation requires one multiplication, and predictive analysis such as logistical regression requires a few multiplications (depending on the precision required). Other applications we describe in the financial and advertising sector use similar functions, except that in

those sectors, the function itself may also be private or proprietary.

2. CLOUD SERVICES

Adoption of cloud services by consumers and businesses is limited by concerns over the loss of privacy or business value of their private data. In this section we will describe concrete and valuable applications of Fully Homomorphic Encryption which can help preserve customer privacy while outsourcing various kinds of computation to the cloud. In all of these scenarios, we imagine a future of streaming data from multiple sources, uploaded in encrypted form to the cloud, and processed by the cloud to provide valuable services to the content owner. There are two aspects of the computation to consider: the data itself, and the function to be computed on this data. We consider cases where one or both of these are private or proprietary and should not be shared with the cloud. In all of these applications, we consider a single content owner, who is the consumer for the cloud service. All data that is encrypted and sent to the cloud is public-key encrypted to the content-owner's public key, using the semantically secure somewhat homomorphic encryption scheme from [5] described later in this paper.

2.1 Medical Applications: Private data and Public functions

In [6], a private cloud medical records storage system (Patient Controlled Encryption) was proposed, in which all data for a patient's medical record is encrypted by the healthcare providers before being uploaded to the patient's record in the cloud storage system. The patient controls sharing and access to the record by sharing secret keys with specific providers (features include a hierarchical structure of the record, ability to search the encrypted data, and various choices for how to handle key distribution). However this system does not provide for the cloud to do any computation other than search (exact keyword match, or possibly conjunctive searches). With our FHE implementation, we add the ability for the cloud to do computation on the encrypted data on behalf of the patient. Imagine a future where monitors or other devices may be constantly streaming data on behalf of the patient to the cloud. With FHE, the cloud can compute functions on the encrypted data and send the patient updates, alerts, or recommendations based on the received data. The functions to be computed in this scenario may include averages, standard deviations or other statistical functions such as logistical regression which can help predict the likelihood of certain

dangerous health episodes. Encrypted input to the functions could include blood pressure or heart monitor or blood sugar readings, for example, along with information about the patient such as age, weight, gender, and other risk factors. The functions computed may not need to be private in this case since they may be a matter of public health and thus public.

2.2 Financial Applications: Private data and Private functions

In the financial industry there is a potential application scenario in which both the data and the function to be computed on the data is private and proprietary. As an example, data about corporations, their stock price or their performance or inventory is often relevant to making investment decisions.

Figure 1 below shows how homomorphic encryption is done in a banking system. Data may even be streamed on a continuous basis reflecting the most up-to-date information necessary for making

decisions for trading purposes. Functions which do computations on this data may be proprietary, based on new predictive models for stock price performance and these models may be the product of costly research done by financial analysts, so a company may want to keep these models private to preserve their advantage and their investment. With FHE, some functions can be evaluated privately as follows. The customer uploads an encrypted version of the function to the cloud, for example a program where some of the evaluations involve encrypted inputs which are specified. The streaming data is encrypted to the customer's public key and uploaded to the cloud. The cloud service evaluates the private function by applying the encrypted description of the program to the encrypted inputs it receives. After processing, the cloud returns the encrypted output to the customer.

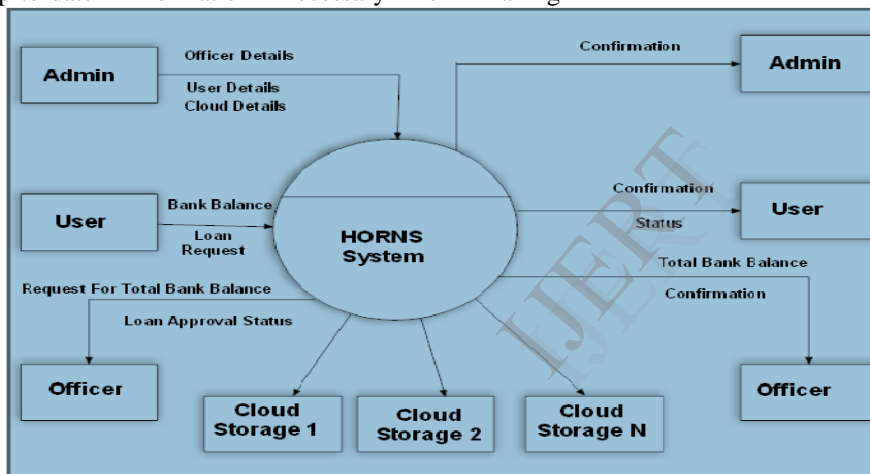


Figure 1 - High level diagram of homomorphic encryption for banking system

2.3 Advertising and Pricing

Imagine an advertiser, for example a cosmetics company, who wants to use contextual information to target advertising to potential customers. The consumer uses a mobile phone as a computing device, and the device constantly uploads contextual information about the consumer, including location, the time of day, information from email or browsing activity such as keywords from email or browser searches. In the future, imagine that information is uploaded potentially constantly from video devices: either pictures of objects of interest such as brands or faces which are automatically identified, or from a video stream from a camera on the body which is identifying context in the room (objects, people, workplace vs. home vs. store). When contextual

information is uploaded to the cloud server and made accessible to the cosmetics company, the company computes some function of the contextual data and determines. This targeted advertisement is sent back to the consumer's phone. Some examples of where context is important for advertising or providing targeted coupons: beer commercials during sports events, or, you are near a Starbucks in the morning and a coffee discount coupon for the Starbucks nearby is sent to your phone, or, cosmetics companies market different products for different times of day (e.g. Friday night going out vs. Sunday morning hanging out with the family), advisor coupons for shows if you are in New York near Broadway in the evening. Other (private) contextual data might be: your income, your profession, your

purchasing history, your travel history, your address, etc. Encrypted version: The problem with these scenarios is the invasion of privacy resulting from giving that much detailed information about the consumer to the server or to the advertising company. Now, imagine an encrypted version of this entire picture. All the contextual data is encrypted and then uploaded to the server; the advertiser uploads encrypted ads to the server; the server computes a function on the encrypted inputs which determines which encrypted ad to send to the consumer; this function could be either private/proprietary or not. All contextual data and all ads are encrypted to the consumer's public key. Then the cloud can operate and compute on this data, and the consumer can decrypt the received ad. As long as the cloud service provider does not collude with the advertisers, and semantically secure FHE encryption is employed, the cloud and the advertisers don't learn anything about the consumer's data.

2.4 Functions to be computed with FHE

We can compute the following functions with a somewhat homomorphic encryption scheme:

- Average of n terms: $\text{fcg} = \frac{1}{n} \sum_{i=1}^n c_i$, where c_i is the average. Standard deviation: $\sqrt{\frac{1}{n} \sum_{i=1}^n (c_i - \text{fcg})^2}$, returned as a pair which is the numerator and denominator of the expression, before taking the square root.
- Logistical regression: $x = \frac{1}{1 + e^{-\sum_{i=1}^n w_i x_i}}$, where w_i is the weighting constant or regression coefficient for the variable x_i , and the prediction is $f(x) = \frac{1}{1 + e^{-x}}$.

A couple of remarks are in order. First, we set the parameter choices for the encryption system based on the expected number of multiplication operations to be done to compute the given functions. These parameter choices determine the efficiency and security of the system. Thus parameters for the system need to be changed as the functions to be computed change. Secondly, so far we do not have a way to efficiently do divisions of real numbers or square roots. Thus in the above computations, numerators and denominators need to be returned as separate encryptions. If the cloud and the advertiser collude, then the cloud may be able to learn some information about whether the user likes the ad or not, which reveals information about his preferences. This constitutes a form of CCA attack, which might endanger the security of the FHE.

3. THE ENCRYPTION SCHEME

We describe the ring learning with errors (Ring LWE) assumption of [LPR10] in Section 3.1, and present the "Somewhat" homomorphic encryption scheme of Brakerski and Vaikuntanathan [5] based on Ring

LWE in Section 3.2. We then report on an instantiation of the parameters, as well as the running times and sizes of the keys and ciphertext in Section 5.

3.1 The Ring LWE Assumption

In this section, we describe a variant of the ring learning with errors assumption of Lyubashevsky, Peikert and Regev. In the RLWE assumption, we consider rings for some degree n integer polynomial and a prime integer the ring of degree n polynomials modulo $f(x)$ with coefficients in \mathbb{Z}_q . Addition in these rings is done component-wise in their coefficients (thus, their additive group is isomorphic to \mathbb{Z}^n and \mathbb{Z}_q^n respectively). Multiplication is simply polynomial multiplication modulo (and also q , in the case of the ring R_q). Thus an element can be viewed as a degree n polynomial over \mathbb{Z}_q . One can represent such an element using the vector of its coefficients. For an element we let note its ℓ_1 norm. The RLWE assumption is parameterized by an integer polynomial of degree n (which defines the ring R), a prime integer q and an error distribution over R , and is defined as follows. Let $s \in R_q$ be a uniformly random ring element. The assumption is that given any polynomial number of samples of the form (a_i, b_i) , where a_i is uniformly random and $b_i = s a_i + e_i$ where e_i is drawn from the error distribution, the b_i 's are computationally indistinguishable from uniform in R_q . As shown in [10], this is equivalent to a variant where the secret s is sampled from the noise distribution rather than being uniform in R_q . It is also easy to see that the assumption is equivalent to a variant where the noise e_i are multiples of some integer t that is relatively prime to q . We consider the RLWE problem for specific choices of the polynomial $f(x)$ and the error distribution χ . Namely, we set $f(x)$ to be the cyclotomic polynomial $x^n + 1$ for n a power of two. In addition to many other useful properties, the fact that means that multiplication of ring elements does not increase their norm by too much (see Lemma 3.2 below). The error distribution is the discrete Gaussian distribution for sample from this distribution defines a polynomial $p(x)$. We present some elementary facts about the Gaussian error distribution, and multiplication over the ring. The result fact bounds the (Euclidean and therefore, the ℓ_1) length of a vector drawn from a discrete Gaussian of standard deviation σ by $n \sigma$. The second fact says that multiplication in the ring $\mathbb{Z}[x] = \sum_{i=0}^{n-1} h_i x^i$ increases the norm of the constituent elements only by a modest amount.

3.2 Somewhat Homomorphic Encryption

The somewhat homomorphic encryption scheme SHE = (SH:Keygen; SH:Enc; SH:Add; SH:Mult; SH:Dec) is associated with a number of parameters:

the dimension n , which is a power of 2, the cyclotomic polynomial $f(x) = x^n + 1$, the modulus q , which is a prime such that $q \equiv 1 \pmod{2n}$. Together, n ; q and $f(x)$ define the rings $R = \mathbb{Z}[x]/(f(x))$ and $R_q = R/qR = \mathbb{Z}_q[x]/(f(x))$. The error parameter, which defines a discrete Gaussian error distribution with standard deviation σ a prime $t < q$, which defines the message space of the scheme as $\mathcal{M} = \mathbb{Z}_t[x]/(f(x))$, the ring of integer polynomials modulo $f(x)$ and t , and a number $D > 0$, which defines a bound on the maximum number of multiplications that can be performed correctly using the scheme. These parameters will be chosen (depending on the security parameter) in such a way as to guarantee correctness and security of the scheme.

3.2.1 The Scheme

Keygen sample a ring element s and define the secret key, Sample a uniformly random ring element $a_1 \in R_q$ and an error e and compute the public key. Publish pk and keep secret. Recall that our message space is \mathcal{M} . Namely, we encode our message as a degree n polynomial with coefficients in \mathbb{Z}_t . Given the public key $pk = (a_0; a_1)$ and a message $m \in \mathcal{M}$, the encryption algorithm samples and computes the ciphertext. We now show how to compute the addition and multiplication operations homomorphically. To compute an arbitrary function f homomorphically, we construct an arithmetic circuit for f (made of addition and multiplication operations over \mathbb{Z}_t), and then use Mult to iteratively compute f on encrypted inputs. Although the ciphertexts reduced by SH:Enc contain two ring elements, the homomorphic operations (in particular, multiplication) increases the number of ring elements in the ciphertext. In general, the SH:Add and mult operations get as input two ciphertexts and the output of SH:Add contains \max ring elements, whereas the output of SH:Multi contains ring elements. Assume that otherwise pad the shorter ciphertext with zeroes. Homomorphic addition is done by simple component-wise addition of the ciphertexts. Namely, compute and output $c + e$ either of the ciphertexts with zero. Let v be a symbolic variable and consider the expression.

3.2.2 Correctness and Security

We show the correctness of decryption and homomorphic evaluation in the following lemmas. The statement of the lemma also serves as the setting of the modulus q (in terms of n ; t ; n and D) that ensures that the scheme can perform D multiplications and A additions. Lemma 3.3. The encryption scheme SHE is correct, and can compute D multiplications followed by A additions, Proof First, note that the ciphertext can be written in the polynomial where

each coefficient has at most with overwhelming probability. This is because e ; u ; f ; s and g are all polynomials whose coefficients are drawn from a discrete Gaussian with standard deviation σ and multiplying two such polynomials (mod $x^n + 1$) produces a polynomial whose coefficients are of size at most with overwhelming probability (by the Central Limit theorem). Our experiments show that this number is in fact smaller, and is of the order of 2. Before we prove correctness of the homomorphic operations, we state an invariant that holds for all ciphertexts produced either by the encryption algorithm, or as a result of a homomorphic evaluation. The invariant is that for a ciphertext where s is the secret key, small error m is the message. Clearly, the invariant holds for a fresh ciphertext produced by SH:Enc (by the calculation above), assuming that. Furthermore, if the invariant holds, then the decryption algorithm succeeds. This is because the decryption algorithm outputs $fct(s) \pmod{t}$ which is indeed the message, assuming the bound on the error. The bound on the error essentially ensures that the quantity that the decryption algorithm recovers does not wrap around mod q . Correctness of homomorphic addition is easy to see. We have two ciphertexts and that satisfy the invariants, the sum of the two ciphertexts is $c + e$ where the addition is done componentwise. This satisfies the invariant as well, assuming that the larger error $t(e + e_0)$ is smaller than q . In general, adding A ciphertexts with error at most each results in error at most A . In practice, as our experiments show, this is likely to be smaller, namely of the order of $2pA$.

3.2.3 An Optimization to Reduce Ciphertext Size

The homomorphic multiplication operation described above increases the number of ring elements in a ciphertext. Brakerski and Vaikuntanathan [BV11a] describe a transformation called re-linearization that reduces the ciphertext back to two ring elements. We describe this optimization below, implement it and report on the performance numbers. Essentially, the idea is the following: assume that we run SH:Multi on two ciphertexts (each containing two ring elements) produced by the encryption algorithm. The resulting ciphertext ct contains three ring elements that satisfy the "invariant" $fct(mt) = c_2s^2 + c_1s + c_0 = t emult + mm_0$. This is a quadratic equation in s , and thus, SH:Multi turned two "linear ciphertexts" into a "quadratic ciphertexts". The goal of re-linearization is to bring this back down to a linear ciphertext. To this end, we publish some homomorphism keys to aid re-linearization. This could be thought of as part of the public key, but the homomorphism key is only used for re-linearization (following an SH:Multi operation).

The homomorphism key hk is computed. In a sense, quasi-encryptions these are They are not real encryptions since t_1, t_2 may not lie in the message space of the encryption scheme, namely R_t . The homomorphic multiplication generates a ciphertext $ct = (c_0; c_1; c_2)$, starting from two 2-element ciphertexts. Re-linearization is performed after every homomorphic multiplication, and proceeds as follows. On the one hand, re-linearization reduces the length of the ciphertext considerably. On the other side, the public parameters become much larger. They now consist of an additional $\log t$ ring elements, totaling to $\log t \cdot n \cdot \log t = n(\log t)^2 = \lg t$ bits. Re-linearization also affects the running time of the homomorphic multiplication. In particular, one needs to additionally perform roughly $\log t \cdot q$ polynomial multiplications and additions. These side effects are most pronounced for small t , where our experiments indicate considerable overhead. Quite encouragingly, though, the benefits of re-linearization seem to dominate the side-effects for large t (see Section 5 for more details). The security of the encryption scheme given the homomorphism keys relies on the circular security of the encryption scheme when encrypting quadratic functions of the secret key.

4. MESSAGE ENCODING TECHNIQUES

The ease of performing homomorphic operations depends crucially on the specific message-encoding used in the ciphertexts. Consider the following two examples. If we wish to compare two encrypted integers $x, y \in \mathbb{Z}_t$ homomorphically then it seems best to encrypt them as (x, y) . The summation runs from $i = 0$ to $i = d \log t \cdot q - 1$. We omit the indices for brevity. Bitwise rather than as an element of \mathbb{Z}_t . The former approach translates to computing a polynomial of degree $\lg t$ over the encrypted bits, whereas the latter seems to require a polynomial of degree $O(t)$. If we wish to compute the mean of k integers, then it seems most natural to encode them as elements of \mathbb{Z}_t (for a large enough t). Computing the mean homomorphically then involves only cheap homomorphic additions over \mathbb{Z}_t . On the other hand, if the numbers are encrypted bit-wise, then addition requires computation of expensive "carry" operations that involve homomorphic multiplication over \mathbb{Z}_2 . We describe two tricks for encoding messages. The first trick shows how to efficiently encode integers in a ciphertext so as to enable efficient computation of their sums and products over the integers. This is useful in computing the mean, the standard deviation and other private statistics efficiently. The second trick shows how to "pack" n encryptions of bits into a single encryption of the n -bit string. Some

homomorphic operations, e.g., comparison of integers or private information retrieval, seem to require bit-wise encryptions of the input. Once the answers are computed, though, they can be packed into a single encryption using this trick.

4.1 Efficient Encoding of Integers for Arithmetic Operations

Given a list of integers if our goal is to compute their sum or product over the integers homomorphically, the obvious (and sub-optimal) choice is to encrypt them directly. Namely, for every m in the list, To ensure that we obtain $\sum m_i$ over the integers (and not mod t), we are forced to choose t to be rather large, namely $t > \sum m_i$, which could be rather prohibitive. We show a method of encrypting integers more efficiently by encoding them in the polynomial ring, in essence enabling a smaller choice of t and better efficiency. In particular, for small enough $m_i < 2n$, we show that it suffices to choose n in order to add n integers. Being able to work with a small t in turn enables us to choose other parameters, e.g., q and n to be correspondingly smaller. The idea is very simple: break each m into (at most n) bits $(m(0); \dots; m(n-1))$, create a degree- $(n-1)$ polynomial $p_m(x) = \sum_{j=0}^{n-1} m(j)x^j$ and encrypt m as $(p_m(x))$. Adding these encryptions now adds up the polynomials $p_{m_i}(x)$ coefficient-wise. Note that each coefficient was a single bit to start with, and a sum of them grows to at most n . As long as $q > t$ this does not wrap around modulo t and upon decryption, we in fact get the polynomial. Now, the result is simply $\sum p_{m_i}(x)$. Extending this idea to support multiplication is a bit trickier. The problem stems from the fact that multiplying the polynomials $p_m(x)$ and $p_{m_0}(x)$ increases their degree. If their original degree was close to n to start with, we will only be able to obtain $p_m(x)p_{m_0}(x) \pmod{x^n + 1}$ upon decryption, which loses information about the product. The solution is to encode the messages m as polynomials of degree at most n/d , if we anticipate performing d multiplications. For our applications (e.g., computing standard deviations), this is an acceptable trade since we only anticipate doing a single multiplication (or, at most a small number of them in the case of computing higher-order regression functions).

5. EXTENSIONS AND FUTURE WORK

Implementing, Fully Homomorphic Encryption. The somewhat homomorphic encryption scheme of [5] can be turned into a fully homomorphic encryption scheme using the re-linearization and the dimension reduction techniques of [BV11a]. We leave the

problem of implementing the resulting fully homomorphic encryption scheme as an important future work. Implementing bootstrapping could also lead to a number of nice applications of homomorphic encryption, for example, to the problem of optimizing communication with the cloud described below. Optimizing communication with the cloud, we present a solution to help mitigate the problem of the large ciphertext size for the Ring-LWE based FHE solution. In any of the above applications, a client communicates with the cloud service and uploads its data encrypted under a FHE scheme, and the cloud operates on this data and returns encrypted outputs to the client. Each ciphertext has size $n \log(q)$, and for functions requiring a large number of multiplications, q and n could be very large (see the implementation section for sample choices of q and n). In short, all the communication over the network consists of short, non-homomorphic ciphertexts. At the server's end, the ciphertexts are first "upgraded" to homomorphic ciphertexts which are then computed on, and finally "downgraded" to short non-homomorphic ciphertexts which are then sent to the client.

[6] Melissa Chase, Kristin Lauter, Josh Benaloh, and Eric Horvitz. "Patient-controlled encryption":

REFERENCES

[1] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. "Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In Shai Halevi, editor, CRYPTO, volume 5677 of Lecture Notes in Computer Science, pages 595-618. Springer 2009".

[2] Web Bosma, John Cannon and Catherine Playoust. "The Magma algebra system I: The user language. J. Symbolic Compute, 24(3-4):235-265, 1997. Computational algebra and number theory (London, 1993)".

[3] Dan Boneh, Eu-Jin Goh, and Kobbi Nissim. "Evaluating 2-DNF formulas, on ciphertexts. In Theory of Cryptography - TCC'05, volume 3378 of Lecture Notes in Computer Science, pages 325-341. Springer, 2005".

[4] Zvika Brakerski, and Vinod Vaikuntanathan. "Efficient fully homomorphic encryption from (standard) LWE, In Submission 2011".

[5] Zvika Brakerski and Vinod Vaikuntanathan. "Fully homomorphic encryption from ring-LWE and security for key dependent messages. To Appear in CRYPTO 2011".