Processing Encrypted Data Using Homomorphic Encryption

Anthony Barnett[6], Charlotte Bonte[1], Carl Bootland[1], Joppe W. Bos[2], Wouter Castryck[1,3], Anamaria Costache[5], Louis Goubin[7,8], Ilia Iliashenko[1], Tancrède Lepoint*[10], Michele Minelli**[11,9], Pascal Paillier[7], Nigel P. Smart[5], Frederik Vercauteren[1,4], Srinivas Vivek[5], and Adrian Waller[6]

[1] imec-Cosic, Dept. Electrical Engineering, KU Leuven
[2] NXP Semiconductors
[3] Laboratoire Paul Painlevé, Université de Lille-1
[4] Open Security Research
[5] University of Bristol
[6] Thales UK, Research and Technology
[7] CryptoExperts
[8] Laboratoire de Mathématiques de Versailles
[9] INRIA
[10] SRI International
[11] Département d'informatique de l'ENS, CNRS, PSL Research University

1 Introduction

Fully Homomorphic Encryption (FHE) was initially introduced as a concept shortly after the development of the RSA cryptosystem, by Rivest et al. [54]. Although long sought after, the first functional scheme was only proposed over thirty years later by Gentry [34, 35] in 2009. The same blueprint to construct FHE has been followed in all subsequent work. First a scheme is constructed which can evaluate arithmetic circuits of a limited depth, a so-called Somewhat Homomorphic Encryption (SHE) scheme. If the complexity of the circuits which the SHE scheme can evaluate is slightly more than the complexity of the decryption circuit for the SHE scheme, then (by placing a SHE encryption of the scheme's private key inside the public key) one can bootstrap the SHE scheme into a FHE scheme. This bootstrapping operation is obtained by homomorphically evaluating the decryption circuit on input of the ciphertext to be bootstrapped and the encryption of the secret key.

* This work was done while the author was employed by CryptoExperts
** This work was done while the author was visiting CryptoExperts. The author was funded from the European Union's Horizon 2020 research and innovation programme under the Marie Sklodowska-Curie grant agreement No. 643161

However, in many application one can perform interesting privacy preserving computations using only a Somewhat Homomorphic Encryption scheme. A European funded research project, HEAT (Homomorphic Encryption Applications and Technology) was founded on the basis of exploring such privacy preserving applications. In this note we explore three such use cases, and the results obtained so far. In the first use case we examine the prediction of house hold electricity consumption within a Smart Grid infrastructure. Here the SHE scheme is used to evaluate a special form of Neural Network which is particularly suited to homomorphic evaluation. In our second application we consider processing of sensor data, and in particular satellite image data. For this application, we consider in this paper the initial processing of the image homomorphically, with future work being focused on learning from the encrypted image information. Finally, our third application looks at privacy preserving querying of crime data across national boundaries.

Before proceeding to our three use cases we outline, for the reader, the three generations of SHE schemes which have so far been developed. So far, there have been roughly three generations of SHE schemes. The first generation consisted of Gentry's original scheme, which was based on having two representations of a basis of an ideal of a number field, one easy basis and one hard basis. Gentry's original scheme was simplified and implemented in [36, 55], where the ideal was chosen to be principal, with the easy basis being the principal generator and the hard basis being the standard two element representation of this ideal. A second family in the first generation of schemes was based on the approximate-GCD problem, and consisted of so-called "integer based" schemes [26]. The first family in the initial generation schemes is now considered insecure due to work of Cramer et al [20], who extended the work of Campbell et al [13] to solve the problem of finding small generators of principal ideals in cyclotomic number fields. The second family, despite having numerous optimizations applied to it - such as [16, 17] - is still not considered competitive compared to the second generation schemes.

The second generation schemes were all based on the Learning With Errors (LWE) problem, and its generalisation to rings (the Ring-LWE problem) [10–12]. These schemes, generally referred to as BGV, were extensively optimized and implemented in a series of works by Gentry et al [37–40], with an implementation (HELib) being given in [42]. A variant of BGV, called FV, was presented in [32] which embeds the message into the upper bits of the underlying ring. The second generation systems

also include those based on the NTRU assumption [9, 48], although the security of these has since been called into question [3].

A third generation of schemes, based on standard LWE and encoding messages via matrix eigenvalues, was presented in [41]. These schemes have an interesting property of asymmetric noise growth; and as such have given rise to some interesting theoretical applications and a fast method to perform boostrapping [28]. However, they are particularly focused on bit-encryption and hence evaluation of binary circuits on encrypted data; thus in practice their efficiency does not match that of the second generation schemes.

2 Privacy-preserving forecasting techniques for the smart grid

Many countries around the world are investing significantly in *smart grid* solutions with the prospect of having a positive impact on the sustainability, reliability, flexibility, and efficiency of the power supply. The deployment of smart meters is already well underway. For example, in the United Kingdom the large energy suppliers were operating over $400,000$ smart gas and electricity meters, representing 0.9 percent of all the domestic meters operated by the large suppliers in 2014 [24]. This development is expected to intensify: the EU third energy package has as an objective to replace at least 80 percent of electricity meters with smart meters by 2020 [31]. This change will fundamentally re-engineer the (electricity) service industry.

The replacement of the classical meters with their smart variants has advantages for both the consumer and industry. Some of the key benefits include giving consumers the information to gain control over their energy consumption, lowering the cost for managing the supply of energy across industry, and producing detailed consumption information data from these smart meters which in turn enable a wide range of services [24]. It is expected that the meters have an update rate of every 15 minutes at least [30]. When generating such a large amount of consumer data a lot of privacy sensitive information is being disclosed. There are various initiatives (e.g. [52, 56]) which stress and outline the importance of having solutions for the smart grid where privacy protecting mechanisms are already built-in by design.

One of the areas where industry would like to use this smart data is to perform a forecast in order to buy energy generation contracts that cover their clients. Moreover, to ensure network capacity the network operators

require longer term forecasting [44, 56, 25]. This forecasting is typically done by taking as input the (aggregated) data from a number of households. Based on this consumption data, together with other variables such as the date and the current temperature and weather, a forecast is computed to predict the short, medium, or long term consumption. The energy providers or network operators only need to know the desired forecast information based on their (potentially proprietary) forecasting algorithm and model. There is no need to observe the individual consumer data. The computation on this aggregated data could be performed in a privacy-friendly manner: something which is currently not the case.

Additively homomorphic encryption schemes [51] and other tools have been proposed to enhance the privacy in the setting of computing detailed billing in the context of the smart grid [53, 50, 33, 46, 29, 45]. However, these approaches cannot be directly used in the setting of prediction algorithms since these more complex algorithms need to compute both additions and multiplications.

One popular class of prediction algorithms are based on artificial neural networks. However, one of the main ingredients in these forecasting algorithms is the computation of the activation function: in practice a sigmoid function is often used where the logistic function $t \mapsto 1/(1+e^{-t})$ is a popular choice. However, computing such a sigmoid function homomorphically is far from practical. One possible way to proceed is to simply ignore the sigmoidality requirement and to proceed with a truncated Taylor series approximating this function or, more generally, to use any non-linear polynomial function which is *simple*. This was investigated by Livni et al. [47] regardless of cryptographic applications. Recent work by Xie et al. [58] and Dowlin et al. [27] suggests to apply the same approach to homomorphically encrypted data. However, by computing artificial neural networks in this fashion it becomes just an organized manner of fitting a polynomial through the given data set. We have studied [8] the application of of Ivakhnenko's group method of data handling (GMDH) which was proposed back in 1970 [43].

The goal of GMDH is to approximate our target function $\tilde{f} : \mathbf{R}^{n_0} \to \mathbf{R}$ with a truncated Wiener series

$$a_0 + \sum_{i=1}^{n_0} a_i x_i + \sum_{i=1}^{n_0} \sum_{j=i}^{n_0} a_{ij} x_i x_j + \sum_{i=1}^{n_0} \sum_{j=i}^{n_0} \sum_{k=j}^{n_0} a_{ijk} x_i x_j x_k + \ldots,$$

which is also called a Kolmogorov-Gabor polynomial. The idea is to approach this by a finite superposition of quadratic polynomials

$$\nu_{ij} : \mathbf{R}^2 \to \mathbf{R} : (x, y) \mapsto b_{ij0} + b_{ij1} x + b_{ij2} y + b_{ij3} xy + b_{ij4} x^2 + b_{ij5} y^2$$

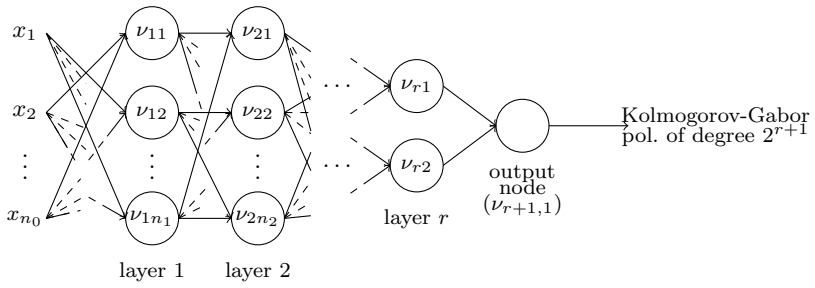

Fig. 1. Illustration of the Group Method of Data Handling.

as is illustrated in Figure 1. One can think of this as some sort of ANN, and indeed the diagram is sometimes called a 'polynomial neural network'. As a first main difference, however, note that the wiring is incomplete: each neuron has two inputs only.

The choice of the degree of the polynomial modulus used in all popular Somewhat Homomorphic Encryption schemes is dominated by security considerations, while with the current encoding techniques the correctness requirement allows for much smaller values. We have introduced [7] a generic encoding method using expansions with respect to a non-integral base, which exploits this large degree at the benefit of reducing the growth of the coefficients when performing homomorphic operations. This allows one to choose a smaller plaintext coefficient modulus which results in a significant reduction of the running time.

Let us recall and summarize the exact forecasting setting and the parameters we selected for the implementation. It is our goal to predict the energy consumption for the next half hour of an apartment complex of 10 households while not revealing any energy consumption information to the party computing on this data. In order to assess the practical performance we implemented this privacy-preserving homomorphic forecasting approach. Our implementation uses the FV-NFLlib software library [22] which implements the FV homomorphic encryption scheme which in turn uses the NFLlib software library (as described in [49] and released at [23]) for computing polynomial arithmetic. Our benchmark results are obtained when running the implementation on an average laptop equipped with an Intel Core i5-3427U CPU (running at 1.80GHz).

We used the data that was collected through the Irish smart metering electricity customer behaviour trials [15] which ran in 2009 and 2010 with over 5,000 Irish homes and businesses participating. Our GMDH network of three hidden layers with 8, 4 and 2 nodes, respectively. As input layer

a set of 51 nodes is used, where 48 nodes represent the half hour measurements that were made during the previous 24 hours. The remaining 3 inputs correspond to the temperature, the month, and the day of the week. The single output node then returns the predicted electricity consumption for the next half hour. Each node performs 8 multiplications, since there are at most 15 nodes being evaluated this means computing 120 multiplications and 75 additions homomorphically. The entire running time of our GMDH implementation to forecast on encrypted data in combination with our new encoding scheme the homomorphic forecasting can be done in only 2.5 seconds: making this approach suitable for industrial applications in the smart grid.

3 Privacy-preserving processing on sensor data

Signal and image processing algorithms are used widely, to improve the quality of sensor data, perform object detection and classification, and many other applications. The processing of sensor data quite clearly gives rise to concerns about privacy, and access to the data needs to be controlled. As an additional concern, those creating the processing algorithms often invest significant resources to do so, and wish to protect sensitive parameter details of the algorithms as their Intellectual Property. These motivate the need for being able to perform signal and image processing on encrypted data.

As a specific example, consider imaging satellites. Due to the nature of satellite operations, current systems tend to be owned and operated by one organization and built for a particular purpose. The development and operation of satellite systems is not cheap, and cost pressures are starting to demand the need for more flexibility in satellite missions. We are likely to see the operational trend moving towards the repurposing of satellites and sharing of infrastructure in the coming years. This change may enable new entities, e.g. from academia or industry, to take advantage of available satellite infrastructure. The satellite data may be sensitive in a number of respects. At present, the sure way for an organisation to prevent data leakage is to own and manage the entire system including the satellite vehicle, its payload, the ground station, and the data processing facilities. Homomorphic Encryption (HE) presents a technical opportunity by which the aforementioned facilities might be shared between organizations that do not necessarily trust each other, so reducing the cost of operating satellite systems. In this context, encrypted data can be sent between organizations, trusted or not, and remain secure.

A key tool in signal and image processing is the Fourier Transform. Consider the Discrete Fourier Transform (DFT) for a sequence of complex numbers $x_0, x_1, \ldots, x_{N-1}$,

$$X_k = \sum_{n=0}^{N-1} x_n \cdot e^{-2 \cdot \pi \cdot i \cdot k \cdot n / N}, \quad k = 0, \ldots, N-1.$$

This is a linear transform; it requires only the summation of sequence samples and multiplication by scalar values $e^{-2 \cdot p \cdot i \cdot k \cdot n / N}$ (sometimes called twiddle factors). We usually implement this via a Fast Fourier Transform (FFT) which reduces the computational complexity by the use of a few mathematical tricks.

Previous work has considered how FFTs can be implemented on encrypted data [4, 5]. However, this work is limited when considering practical applications as it only considers one FFT and makes use of the Paillier cryptosystem [51] that can only perform ciphertext additions and not ciphertext multiplications. For example, consider the Single-Look Synthetic Aperture Radar (SLSAR) processing chain shown in Figure 2. SLSAR is used to turn raw radar sensor data into images. It involves a series of operations, including FFTs, Inverse FFTs (IFFT) and Hadamard multiplication of matrices. Some of these multiplications may require sensitive inputs, to protect details of the algorithm. More complex algorithms, such as Multi-Look SAR which combines multiple sensor images, require additional ciphertext multiplications. These require a Somewhat Homomorphic Encryption (SHE) scheme that allows a limited number of ciphertext multiplications. This led authors to examine how to perform FFT operations using Somewhat Homomorphic Encryption, see [14] and [19].

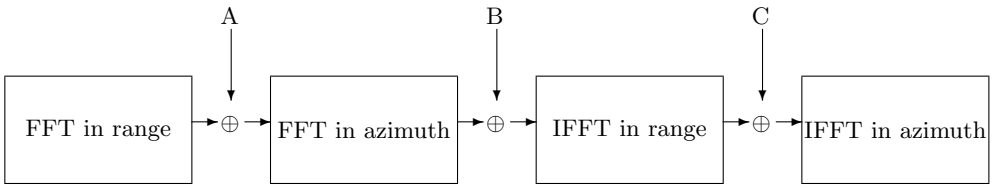

Fig. 2. SLSAR algorithm block diagram. The filter matrices A, B and C are multiplied point-wise with the output of the previous FFT

In [18] we introduce a new method to homomorphically compute on complex numbers using a Somewhat Homomorphic Encryption scheme.

For the evaluation of the FFT pipeline above our methods can also dispense with the associated approximations of complex numbers, and we find we can evaluate the DFT pipeline using exact operations on encodings of exact complex numbers. This methodology enables us to achieve a considerable improvement in the ability to homomorphically evaluate a DFT. Notice that despite the DFT being linear, the large number of additions and scalar multiplications means that the often heard mantra of "only multiplications matter" does not apply. We need to be careful not only of the growth of the coefficients of the ring elements which encode our values, but also of the homomorphic noise.

We are able to evaluate a single iteration of a FFT-Hadamard-iFFT pipeline of input size 8192 elements, as opposed to 1024 elements for [14] and [19]. In terms of latency we were able to evaluate a pipeline for 256 elements in 9.43 *seconds*, compared to a latency of 581 *minutes* for [19] and 87 *minutes* for [14]. Our amortized times are however much worse; since our method does not allow packing our amortized time for the same calculation is still 9.43 seconds, compared to 89.4 seconds for [19] and 0.31 seconds for [14]. So whilst we obtain faster latency (and exact computations), for high throughput calculations the method of [14] is still to be preferred.

More complex applications require even more complex processing chains. Our current work is considering how machine learning algorithms can be implemented on encrypted data as part of such processing chains, for image classification for example.

4 Privacy preserving processing of crime related data

Organized Crime is becoming increasingly diverse in its method, group structures and impact on society. A new criminal landscape is emerging, marked increasingly by highly mobile and flexible groups operating in multiple jurisdictions and criminal sectors. Internet and mobile technologies have emerged as key facilitators for organized crime. Although electronic communications have made organized crime activities less visible to authorities targeting criminal assets, the increasing usage of the Internet and of mobile communications offers new opportunities to investigators to detect signals and to pre-empt organized crime activities.

However, police forces have to comply with the law and democracy protects the rights of citizens and in particular the privacy of their personal data. Wide-range scanning for weak organized crime (OC) signals is typically incompatible with the legal constraints because it would un-

duly give power to the executive arm and thereby limit personal freedom. Such data privacy rules are even more stringent when it comes to the collaboration between investigators from different agencies and countries. There is then a conflict between safety and privacy and the fundamental paradox that arises from this picture is that protecting citizens' rights makes it more difficult to protect citizens' rights.

The cross-border collaboration between law enforcement agencies is already regulated by a EU framework that was introduced in 2008 by EU Council Decision 2008/615/JHA [1] and EU Council Decision 2008/616/JHA [2], which applied the previously established Prum Convention [57] to all the member states. These decisions describe a framework in which member states can grant one another access rights to their automated DNA analysis files, automated dactyloscopic identification systems or vehicle registration data via a two-step process: a hit/no-hit system (whose result should be available in less than 15 minutes) followed by a request for specific related personal data.

The current framework still presents issues regarding the protection of citizens' privacy; for example the country that receives a query can learn the query itself. At the time the decisions were produced, several cryptographic primitives like Fully Homomorphic Encryption (FHE) were not known to be possible, so the results that can be achieved now are more sophisticated than what was imaginable in 2008 and the privacy model can be considerably strengthened.

As in [21], we assume a party, say France (FR), wants to query a database held by another party, say Germany (DE); furthermore, we assume that both countries recognize the authority of a trusted party, say a Judge (JU). Then the solution that we propose achieves the following privacy/security goals:

- FR learns no more than whether there is a hit/no-hit on the data XXX and, if authorized by JU, which records match XXX;
- DE does not learn anything (not even XXX), but the possible data it would legally be obliged to provide after a successful match;
- JU learns the query XXX but not the associated data, even if he allows the query.

In our architecture, DE is the only one who controls its database, and sends an (encrypted) inverted index of the database to an untrusted party (e.g. the Cloud). This step only happens once at the beginning, and each time the database needs to be updated. If we denote by PRF a pseudo random function, in a first step DE sends to the Cloud a database

containing tuples of the form

$$\left(\quad \mathsf{PRF}(\text{“DNA}_{\mathsf{D19S433}} = 7, 8\text{”}), \quad \mathsf{Enc}_{k_{\mathsf{DE}}}(\mathcal{I}) \quad \right)$$

where $\mathcal{I}$ is the set of indices in the database of the individuals for which the DNA marker D19S433 is equal to $7, 8$, encrypted with a key k_{DE} only known to DE.

Then in the second step, FR and DE can obliviously compute the values of the form

$$\mathsf{PRF}(\text{“DNA}_{\mathsf{D19S433}} = 7, 8\text{”})$$

corresponding to the search query FR wants to perform, similarly to what was shown in [21].

In a third step, if JU accepts FR's query; he gets the encrypted set of indices corresponding to the tokens FR and DE computed; homomorphically performs the required computation if needed – an intersection of the sets of indices if the query is conjunctive, a union of the sets if it is a disjunctive query –; signs the resulting encrypted set of indices with his secret key k_{JU} sends;

$$\left(\qquad \mathsf{Enc}_{k_{\mathsf{DE}}}(\mathcal{I}), \qquad \mathsf{Sgn}_{k_{\mathsf{JU}}}\left(\mathsf{Enc}_{k_{\mathsf{DE}}}(\mathcal{I})\right) \qquad \right)$$

to FR.

In a final step, FR forwards the data to DE, who first checks JU's signature and, if it is correct, decrypts the set of indices $\mathcal{I}$ and sends the corresponding data to FR.

In order to allow FR to perform conjunctive and disjunctive queries on DE's database, we need to use an encryption scheme with homomorphic properties.

We now turn to specifying the homomorphic encryption scheme $\mathsf{Enc}_{k_{\mathsf{DE}}}(\mathcal{I})$ used to encrypt a set of indices $\mathcal{I}$ under the secret key k_{DE}. In order to encrypt the set of indices $\mathcal{I} \subset \mathbb{Z}$, we first encode the indices in a polynomial

$$p = \prod_{i \in \mathcal{I}} (x - i)\,,$$

of degree $\mathrm{card}(\mathcal{I})$, the cardinality of $\mathcal{I}$. Then, we encrypt independently the coefficients p_j of $p = \sum p_j x^j$, and we define

$$\mathsf{Enc}_{k_{\mathsf{DE}}}(\mathcal{I}) = \left(\quad \mathsf{Enc}_{k_{\mathsf{DE}}}(p_0), \quad \ldots, \quad \mathsf{Enc}_{k_{\mathsf{DE}}}(p_{\mathrm{card}(\mathcal{I})-1}) \quad \right).$$

Note that the integer roots of the polynomial p match exactly the set $\mathcal{I}$ and the encrypted polynomials do not leak anything thanks to the

semantic security of the encryption scheme. A solution to ensure secrecy for several queries is defined in [6].

When performing a disjunctive search query (union), one must then compute homomorphically the coefficients of the product polynomial

$$p(x) = p^{(1)}(x) \cdot p^{(2)}(x)$$

where $\mathcal{I}_1$ and $\mathcal{I}_2$ are the roots of $p^{(1)}(x)$ and $p^{(2)}(x)$ respectively; then the roots of $p(x)$ are exactly $\mathcal{I}_1 \bigcup \mathcal{I}_2$ as required. Since the coefficients of $p(x)$ are quadratic in the coefficients of $p^{(1)}(x)$ and $p^{(2)}(x)$, the homomorphic encryption scheme must allow for at least one multiplication.

On the other hand, when performing a conjunctive search query (intersection), one chooses two random polynomials $r_1(x)$ of degree $d_2 - 1$ and $r_2(x)$ of degree $d_1 - 1$ where d_i is the degree of $p^{(i)}(x)$. Then the party homomorphically computes the coefficients of

$$p(x) = p^{(1)}(x) \cdot r_1(x) + p^{(2)}(x) \cdot r_2(x)$$

We stress that this operation can introduce "parasitic" roots (i.e. false positives) but the probability of such an event can be made arbitrarily small, e.g. by setting the parameters to make it negligible or, more simply, by repeating the operation multiple times with different random polynomials and then taking the intersection of the resulting sets of roots for $p(x)$. We also note that the coefficients of $r_1(x)$ and $r_2(x)$ need not be encrypted, so a linearly homomorphic encryption scheme is sufficient to support conjunctive queries. Finally, we stress that, in order to support more complex (ideally, arbitrary) queries, we need an encryption scheme that supports arbitrarily many additions and multiplications, i.e. a fully homomorphic encryption scheme.

For our implementation, we extended the DGHV scheme [26], as described below.

The DGHV scheme with message space $\mathbb{Z}_q$. Given the security parameter λ, all the other parameters are chosen as a function of λ, i.e. $\eta = \eta(\lambda), \gamma = \gamma(\lambda), \rho = \rho(\lambda)$.

$\mathsf{KeyGen}(1^\lambda)$. Generate a random prime integer p of size η bits. Generate a random prime s_0 of size $\gamma - \eta$ bits and let $x_0 = s_0 \cdot p$. Let $pk = x_0$ and $sk = p$.

$\mathsf{Encrypt}(sk, m \in \mathbb{Z}_q)$. Generate a random postive integer s of $\gamma - \eta$ bits, a random integer r in $(-2^\rho, 2^\rho)$ and output the ciphertext:

$$c = s \cdot p + r \cdot q + m$$

$\mathsf{Add}(c_1, c_2, pk)$. Return $c \leftarrow c_1 + c_2 \bmod x_0$.

$\mathsf{Mult}(c_1, c_2, pk)$. Return $c \leftarrow c_1 \cdot c_2 \bmod x_0$.

$\mathsf{Decrypt}(sk, c)$. Output $m \leftarrow (c \bmod p) \bmod q$.

We also implemented the proposed solution on realistically-sized databases and we ran it on an average machine. The timing performance is well within the limit set by the EU Council Decisions [1, 2] and we thus obtain the first usable implementation of the Automated Detection of Organized Crime (ADOC) framework with an enhanced privacy model.

References

1. Council Decision 2008/615/JHA of 23 June 2008 on the stepping up of cross-border cooperation, particularly in combating terrorism and cross-border crime. `http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex:32008D0615`, 2008.
2. Council Decision 2008/616/JHA of 23 June 2008 on the implementation of Decision 2008/615/JHA on the stepping up of cross-border cooperation, particularly in combating terrorism and cross-border crime. `http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex:32008D0616`, 2008.
3. M. R. Albrecht, S. Bai, and L. Ducas. A subfield lattice attack on overstretched NTRU assumptions - cryptanalysis of some FHE and graded encoding schemes. In M. Robshaw and J. Katz, editors, *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part I*, volume 9814 of *Lecture Notes in Computer Science*, pages 153–178. Springer, 2016.
4. T. Bianchi, A. Piva, and M. Barni. Comparison of different FFT implementations in the encrypted domain. In *2008 16th European Signal Processing Conference, EUSIPCO 2008, Lausanne, Switzerland, August 25-29, 2008*, pages 1–5. IEEE, 2008.
5. T. Bianchi, A. Piva, and M. Barni. On the implementation of the discrete fourier transform in the encrypted domain. *IEEE Transactions on Information Forensics and Security*, 4(1):86–97, 2009.
6. D. Boneh, C. Gentry, S. Halevi, F. Wang, and D. J. Wu. Private database queries using somewhat homomorphic encryption. In M. J. J. Jr., M. E. Locasto, P. Mohassel, and R. Safavi-Naini, editors, *Applied Cryptography and Network Security - 11th International Conference, ACNS 2013, Banff, AB, Canada, June 25-28, 2013. Proceedings*, volume 7954 of *Lecture Notes in Computer Science*, pages 102–118. Springer, 2013.
7. C. Bonte, C. Bootland, J. W. Bos, W. Castryck, I. Iliashenko, and F. Vercauteren. Faster homomorphic function evaluation using non-integral base encoding. In *CHES 2017*, LNCS, page (to appear), 2017.
8. J. W. Bos, W. Castryck, I. Iliashenko, and F. Vercauteren. Privacy-friendly forecasting for the smart grid using homomorphic encryption and the group method of data handling. In M. Joye and A. Nitaj, editors, *AFRICACRYPT 2017*, volume 10239 of *LNCS*, pages 184–201, 2017.

9. J. W. Bos, K. E. Lauter, J. Loftus, and M. Naehrig. Improved security for a ring-based fully homomorphic encryption scheme. In M. Stam, editor, *Cryptography and Coding - 14th IMA International Conference, IMACC 2013, Oxford, UK, December 17-19, 2013. Proceedings*, volume 8308 of *Lecture Notes in Computer Science*, pages 45–64. Springer, 2013.
10. Z. Brakerski, C. Gentry, and V. Vaikuntanathan. Fully homomorphic encryption without bootstrapping. In *Innovations in Theoretical Computer Science (ITCS'12)*, 2012.
11. Z. Brakerski and V. Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In *FOCS'11*. IEEE Computer Society, 2011.
12. Z. Brakerski and V. Vaikuntanathan. Fully homomorphic encryption from ring-LWE and security for key dependent messages. In *Advances in Cryptology - CRYPTO 2011*, volume 6841 of *Lecture Notes in Computer Science*, pages 505–524. Springer, 2011.
13. P. Campbell, M. Groves, and D. Shepherd. SOLILOQUY: A cautionary tale. ETSI 2nd Quantum-Safe Crypto Workshop, 2014.
14. J. H. Cheon, A. Kim, M. Kim, and Y. S. Song. Floating-point homomorphic encryption. *IACR Cryptology ePrint Archive*, 2016:421, 2016.
15. Commission for Energy Regulation. Electricity smart metering customer behaviour trials (CBT) findings report. Technical Report CER11080a, 2011. `http://www.cer.ie/docs/000340/cer11080(a)(i).pdf`.
16. J.-S. Coron, A. Mandal, D. Naccache, and M. Tibouchi. Fully homomorphic encryption over the integers with shorter public keys. In *Advances in Cryptology - CRYPTO 2011*, volume 6841 of *Lecture Notes in Computer Science*, pages 487–504. Springer, 2011.
17. J.-S. Coron, D. Naccache, and M. Tibouchi. Public key compression and modulus switching for fully homomorphic encryption over the integers. In *Advances in Cryptology - EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 446–464. Springer, 2012.
18. A. Costache, N. P. Smart, and S. Vivek. Faster homomorphic evaluation of discrete fourier transforms. *IACR Cryptology ePrint Archive*, 2016:1019, 2016. To appear in Financial Cryptography 2017.
19. A. Costache, N. P. Smart, S. Vivek, and A. Waller. Fixed-point arithmetic in SHE scheme. In *Selected Areas in Cryptography - SAC*, 2016. Full version available at `http://eprint.iacr.org/2016/250`.
20. R. Cramer, L. Ducas, C. Peikert, and O. Regev. Recovering short generators of principal ideals in cyclotomic rings. In M. Fischlin and J. Coron, editors, *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part II*, volume 9666 of *Lecture Notes in Computer Science*, pages 559–585. Springer, 2016.
21. E. D. Cristofaro, Y. Lu, and G. Tsudik. Efficient techniques for privacy-preserving sharing of sensitive information. In J. M. McCune, B. Balacheff, A. Perrig, A. Sadeghi, M. A. Sasse, and Y. Beres, editors, *Trust and Trustworthy Computing - 4th International Conference, TRUST 2011, Pittsburgh, PA, USA, June 22-24, 2011. Proceedings*, volume 6740 of *Lecture Notes in Computer Science*, pages 239–253. Springer, 2011.
22. CryptoExperts. FV-NFLlib. `https://github.com/CryptoExperts/FV-NFLlib`, 2016.

23. CryptoExperts, INP ENSEEIHT, and Quarkslab. NFLlib. https://github.com/quarkslab/NFLlib, 2016.
24. Department of Energy & Climate Change. Smart metering implementation programme. Technical Report Third Annual Report on the Roll-out of Smart Meters, 2014. https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/384190/smip_smart_metering_annual_report_2014.pdf.
25. Department of Energy and Climate Change. Smart metering implementation programme – data access and privacy. https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/43043/4933-data-access-privacy-con-doc-smart-meter.pdf.
26. M. v. Dijk, C. Gentry, S. Halevi, and V. Vaikuntanathan. Fully homomorphic encryption over the integers. In H. Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *Lecture Notes in Computer Science*, pages 24–43. Springer, 2010.
27. N. Dowlin, R. Gilad-Bachrach, K. Laine, K. E. Lauter, M. Naehrig, and J. Wernsing. Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy. In M. Balcan and K. Q. Weinberger, editors, *International Conference on Machine Learning*, volume 48, pages 201–210. JMLR.org, 2016.
28. L. Ducas and D. Micciancio. FHEW: bootstrapping homomorphic encryption in less than a second. In E. Oswald and M. Fischlin, editors, *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part I*, volume 9056 of *Lecture Notes in Computer Science*, pages 617–640. Springer, 2015.
29. Z. Erkin and G. Tsudik. Private computation of spatial and temporal power consumption with smart meters. In F. Bao, P. Samarati, and J. Zhou, editors, *ACNS*, volume 7341 of *LNCS*, pages 561–577. Springer, 2012.
30. European Commission. Commission recommendation of 9 March 2012 on preparations for the roll-out of smart metering systems. Official Journal of the European Union http://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX:32012H0148, March 2012.
31. European Commission. Benchmarking smart metering deployment in the EU-27 with a focus on electricity. Technical Report 365, June 2014. http://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52014DC0356&from=EN.
32. J. Fan and F. Vercauteren. Somewhat practical fully homomorphic encryption. *IACR Cryptology ePrint Archive*, 2012:144, 2012.
33. F. D. Garcia and B. Jacobs. Privacy-friendly energy-metering via homomorphic encryption. In J. Cuéllar, J. Lopez, G. Barthe, and A. Pretschner, editors, *STM*, volume 6710 of *LNCS*, pages 226–238. Springer, 2011.
34. C. Gentry. *A fully homomorphic encryption scheme*. PhD thesis, Stanford University, 2009. http://crypto.stanford.edu/craig.
35. C. Gentry. Fully homomorphic encryption using ideal lattices. In M. Mitzenmacher, editor, *STOC*, pages 169–178. ACM, 2009.
36. C. Gentry and S. Halevi. Implementing gentry's fully-homomorphic encryption scheme. In *EUROCRYPT*, volume 6632 of *Lecture Notes in Computer Science*, pages 129–148. Springer, 2011.
37. C. Gentry, S. Halevi, C. Peikert, and N. P. Smart. Ring switching in bgv-style homomorphic encryption. In I. Visconti and R. D. Prisco, editors, *Security and Cryptography for Networks - 8th International Conference, SCN 2012, Amalfi, Italy, September 5-7, 2012. Proceedings*, volume 7485 of *Lecture Notes in Computer Science*, pages 19–37. Springer, 2012.

38. C. Gentry, S. Halevi, and N. Smart. Fully homomorphic encryption with polylog overhead. In *EUROCRYPT*, volume 7237 of *Lecture Notes in Computer Science*, pages 465–482. Springer, 2012.
39. C. Gentry, S. Halevi, and N. P. Smart. Better bootstrapping in fully homomorphic encryption. In M. Fischlin, J. A. Buchmann, and M. Manulis, editors, *Public Key Cryptography - PKC 2012 - 15th International Conference on Practice and Theory in Public Key Cryptography, Darmstadt, Germany, May 21-23, 2012. Proceedings*, volume 7293 of *Lecture Notes in Computer Science*, pages 1–16. Springer, 2012.
40. C. Gentry, S. Halevi, and N. P. Smart. Homomorphic evaluation of the AES circuit. In R. Safavi-Naini and R. Canetti, editors, *Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings*, volume 7417 of *Lecture Notes in Computer Science*, pages 850–867. Springer, 2012.
41. C. Gentry, A. Sahai, and B. Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In R. Canetti and J. A. Garay, editors, *Advances in Cryptology - CRYPTO 2013, Part I*, pages 75–92. Springer, 2013.
42. S. Halevi and V. Shoup. Design and implementation of a homomorphic-encryption library. Manuscript, available at `http://people.csail.mit.edu/shaih/pubs/he-library.pdf`, Accessed January 2015.
43. A. Ivakhnenko. Heuristic self-organization in problems of engineering cybernetics. *Automatica*, 6(2):207 – 219, 1970.
44. M. Jawurek, F. Kerschbaum, and G. Danezis. Privacy technologies for smart grids - a survey of options. Technical Report MSR-TR-2012-119, November 2012. `http://research.microsoft.com/apps/pubs/default.aspx?id=178055`.
45. K. Kursawe, G. Danezis, and M. Kohlweiss. Privacy-friendly aggregation for the smart-grid. In S. Fischer-Hübner and N. Hopper, editors, *Privacy Enhancing Technologies – PETS*, volume 6794 of *LNCS*, pages 175–191. Springer, 2011.
46. F. Li, B. Luo, and P. Liu. Secure information aggregation for smart grids using homomorphic encryption. In *Smart Grid Comm.*, pages 327–332. IEEE, 2010.
47. R. Livni, S. Shalev-Shwartz, and O. Shamir. On the computational efficiency of training neural networks. In *Advances in Neural Information Processing Systems*, pages 855–863, 2014.
48. A. López-Alt, E. Tromer, and V. Vaikuntanathan. On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In H. J. Karloff and T. Pitassi, editors, *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 1219–1234. ACM, 2012.
49. C. A. Melchor, J. Barrier, S. Guelton, A. Guinet, M. Killijian, and T. Lepoint. NFLlib: NTT-based fast lattice library. In K. Sako, editor, *CT-RSA 2016*, volume 9610 of *LNCS*, pages 341–356. Springer, 2016.
50. A. Molina-Markham, P. J. Shenoy, K. Fu, E. Cecchet, and D. E. Irwin. Private memoirs of a smart meter. In A. G. Ruzzelli, editor, *Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, pages 61–66. ACM, 2010.
51. P. Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *EUROCRYPT 99*, pages 223–238. Springer, 1999.
52. Recommendation to the European Commission. Essential regulatory requirements and recommendations for data handling, data safety, and consumer protection. Technical Report version 1.0, 2011. `https://ec.europa.eu/energy/sites/ener/files/documents/Recommendations%20regulatory%20requirements%20v1.pdf`.

53. A. Rial and G. Danezis. Privacy-preserving smart metering. In H. Reimer, N. Pohlmann, and W. Schneider, editors, *Securing Electronic Business Processes, Highlights of the Information Security Solutions Europe 2012 Conference*, pages 105–115. Springer, 2012.
54. R. Rivest, L. Adleman, and M. L. Dertouzos. On data banks and privacy homomorphisms. In *Foundations of Secure Computation*, pages 169–180, 1978.
55. N. P. Smart and F. Vercauteren. Fully homomorphic encryption with relatively small key and ciphertext sizes. In *Public Key Cryptography - PKC'10*, volume 6056 of *Lecture Notes in Computer Science*, pages 420–443. Springer, 2010.
56. The Smart Grid Interoperability Panel – Smart Grid Cybersecurity Committee. Guidelines for smart grid cybersecurity: Volume 1 - smart grid cybersecurity strategy, architecture, and high-level requirements. Technical Report NISTIR 7628 Revision 1, September 2014. `http://nvlpubs.nist.gov/nistpubs/ir/2014/NIST.IR.7628r1.pdf`.
57. E. Union. Prum convention. `http://register.consilium.europa.eu/doc/srv?l=EN&f=ST%2010900%202005%20INIT`, 2005.
58. P. Xie, M. Bilenko, T. Finley, R. Gilad-Bachrach, K. E. Lauter, and M. Naehrig. Crypto-nets: Neural networks over encrypted data. *CoRR*, abs/1412.6181, 2014.