Noah's Ark: Efficient Threshold-FHE Using Noise Flooding

Morten Dahl
morten.dahl@zama.ai
Zama
Paris, France

Daniel Demmler
daniel.demmler@zama.ai
Zama
Paris, France

Sarah El Kazdadi
sarah.elkazdadi@zama.ai
Zama
Paris, France

Arthur Meyre
arthur.meyre@zama.ai
Zama
Paris, France

Jean-Baptiste Orfila
jb.orfila@zama.ai
Zama
Paris, France

Dragos Rotaru
dragos.rotaru@zama.ai
Zama
Paris, France

Nigel P. Smart
nigel.smart@kuleuven.be
nigel@zama.ai
COSIC, KU Leuven
Leuven, Belgium
and
Zama
Paris, France

Samuel Tap
samuel.tap@zama.ai
Zama
Paris, France

Michael Walter
michael.walter@zama.ai
Zama
Paris, France

ABSTRACT

We outline a secure and efficient methodology to do threshold distributed decryption for LWE based Fully Homomorphic Encryption schemes. Due to the smaller parameters used in some FHE schemes, such as Torus-FHE (TFHE), the standard technique of "noise flooding" seems not to apply. We show that noise flooding can also be used with schemes with such small parameters, by utilizing a switch to a scheme with slightly higher parameters and then utilizing the efficient bootstrapping operations which TFHE offers. Our protocol is proved secure via a simulation argument, making its integration in bigger protocols easier to manage.

CCS CONCEPTS

• **Theory of computation** → **Cryptographic protocols**; **Cryptographic primitives**; • **Security and privacy** → **Security protocols**.

KEYWORDS

Threshold Decryption, Fully Homomorphic Encryption

ACM Reference Format:
Morten Dahl, Daniel Demmler, Sarah El Kazdadi, Arthur Meyre, Jean-Baptiste Orfila, Dragos Rotaru, Nigel P. Smart, Samuel Tap, and Michael Walter. 2023. Noah's Ark: Efficient Threshold-FHE Using Noise Flooding. In *Proceedings of the 11th Workshop on Encrypted Computing & Applied Homomorphic Cryptography (WAHC '23), November 26, 2023, Copenhagen, Denmark.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3605759.3625259

WAHC '23, November 26, 2023, Copenhagen, Denmark

ACM ISBN 979-8-4007-0255-6/23/11.
https://doi.org/10.1145/3605759.3625259

1 INTRODUCTION

The problem of threshold decryption for Fully Homomorphic Encryption (FHE) schemes, called threshold-FHE from henceforth, is as old as FHE itself. The problem is for a set of n parties to have a secret sharing of the underlying FHE secret key so that they can between them decrypt a given FHE ciphertext correctly, in the case where at most t of the parties are corrupt. Indeed, Gentry's original thesis [22] mentioned threshold-FHE as a way of utilizing FHE to perform a very low round complexity semi-honest MPC protocol.

To understand the technical problem with threshold-FHE it is worth considering the "format" of a simple FHE - either public or private key - scheme To explain we utilize the format of BFV/TFHE ciphertexts, but a similar discussion can be provided for other FHE schemes such as BGV. Consider encrypting an element $m \in \mathbb{Z}_p$, using a standard Learning-With-Errors (LWE) ciphertext of the form $(\mathbf{a}, b)$ with ciphertext modulus q, where $\mathbf{a} \in \mathbb{Z}_q^\ell$ and $b \in \mathbb{Z}_q$, using the equation

$$b = \mathbf{a} \cdot \mathbf{s} + e + \Delta \cdot m \pmod{q}$$

where $\Delta = \lfloor q/p \rfloor$, e is some "noise" term and $\mathbf{s} \in \mathbb{Z}_q^\ell$ is the secret key. Usually, in the FHE setting, $\mathbf{s}$ is chosen to be a vector of small norm, for example $\mathbf{s} \in \{0,1\}^\ell$.

To enable threshold-FHE we first secret share the secret key $\mathbf{s}$ among n parties, a process which we shall denote by $[\mathbf{s}]^{\langle t,q \rangle}$ to signal a sharing modulo q with respect to a threshold $t < n$ linear secret sharing scheme. On input of the ciphertext $(\mathbf{a}, b)$ we can then produce trivially a secret sharing of the value $e + \Delta \cdot m$ by computing

$$[t]^{\langle t,q \rangle} = b - \mathbf{a} \cdot [\mathbf{s}]^{\langle t,q \rangle} = [e + \Delta \cdot m]^{\langle t,q \rangle}.$$

By opening the value of $[t]^{\langle t,q \rangle}$ all parties can then perform rounding to obtain m. However, this reveals the value of e, which combined with the ciphertext and the message, will reveal information about the secret key $\mathbf{s}$.

The way around this is to add some additional noise into the secret sharing before the opening. Thus the decrypting parties somehow generate an additional secret shared noise term $[E]^{\langle t,q \rangle}$, and the value which is opened is now

$$[t]^{\langle t,q \rangle} = b - \mathbf{a} \cdot [\mathbf{s}]^{\langle t,q \rangle} + [E]^{\langle t,q \rangle} = [e + E + \Delta \cdot m]^{\langle t,q \rangle}.$$

The key concern is then that E should introduce enough randomness to mask the e value after the shared value $[t]^{\langle t,q \rangle}$ is opened. If E is too small then too much information about e is revealed, if E is too big then the final rounding will not reveal the correct value of m. Diagrammatically we can consider this process as approximated by the diagram in Figure 1.

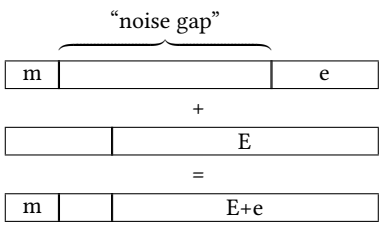

Figure 1: Representation of the noise addition for threshold decryption

To mask, statistically, all information in e we would (naively) require E to be chosen uniformly from a range which is 2^{stat} larger than e. Thus if we can bound the ciphertext noise by $|e| < B$, then we would require E to be chosen uniformly in the range $[-2^{\text{stat}} \cdot B, \ldots, 2^{\text{stat}} \cdot B]$. This process is often dubbed "noise flooding" in the literature. However, this would mean we require $\Delta > 2^{\text{stat}} \cdot B$, which in turn means that the ciphertext modulus q needs to be "large".

We show that by adding an E term on, which is itself the addition of at least two uniform distributions in the range $[-2^{\text{stat}} \cdot B, \ldots, 2^{\text{stat}} \cdot B]$, we are able to obtain a statistical distance of actually $2^{-2 \cdot \text{stat}}$. We can, hence, obtain enough security by selecting stat ≈ 40, and so reduce the need for very much larger q values.

In theory it may be possible to select E from a smaller range, and rely on game based security assumptions. This approach is taken in two recent papers, [13] and [18], via the Renyi divergence. This methodology enables parameters to be chosen in which q is much smaller than the above analysis would require. As we discuss below this approach leads to additional problems in the larger protocols in which we embed our threshold decryption. Thus the use of Renyi divergence is not without problems in this situation.

Before proceeding we note that in many situations there is no problem with the increased size of q that the noise flooding approach requires. The key observation is that FHE enables the use of a bootstrapping operation. The purpose of this operation is to reduce the size of the noise e in the ciphertext $(\mathbf{a}, b)$ to be as small as possible. Thus if bootstrapping is performed, and the FHE scheme is such that the noise gap between e and m in Figure 1 is large enough, then the noise flooding methodology will work "out-of-the-box". Thus for BFV/BGV implementations which enable bootstrapping there is no problem to solve, as noise flooding is enough.

When using BFV/BGV in an SHE leveled mode then the problem also does not occur. In such schemes each level essentially adds an extra 14-24 bits (depending on the implementation) into the noise gap [7, 25, 26]. Thus by simply increasing the number of levels by a small constant (say, two or three) one can obtain a noise gap which is enough to apply the flooding technique. Thus in such schemes our methodology in Section 4 can be applied, without any need for prior pre-processing.

Thus the only place where noise flooding is in practice a problem is when the FHE parameters are such that the noise gap is tiny, even after a bootstrapping operation is performed. This is exactly the situation in TFHE where one (usually) selects a relatively small q value (for example $q = 2^{64}$). This small q value, and associated small LWE dimension ℓ, requires the size of the noise even after bootstrapping to be around 2^{30} in order to ensure security. This means the noise gap is too small, but only by tens of bits. In this work we solve this problem for TFHE, by utilizing the fast bootstrapping enabled by TFHE. In some sense we protect the initial FHE ciphertext from the flooding operation, by placing the underlying message in a larger ciphertext (a kind of protective Noah's Ark).

1.1 Historical Discussion

At about the time of Gentry's thesis on FHE in 2009 [22], the first threshold key generation and decryption for LWE based ciphertexts was given by Bendlin and Damgård [10]. Their methodology used replicated secret sharing to split the secret key, a method whose complexity scales with $\binom{n}{t}$. The simpler case of full-threshold, i.e. $t = n - 1$, decryption for LWE ciphertexts was combined with SHE and formed the basis of the SPDZ MPC protocol [20]. This utilized the BGV encryption scheme, supporting circuits of multiplicative depth one, and used the noise flooding technique mentioned above.

The same techniques were then used in the context of FHE by Asharov et al [4] in the full threshold setting. To obtain active security in settings with dishonest majority one needs to add zero-knowledge proofs into the mix, see [3] which gives a practical instantiation using noise flooding for BGV (in the context of a voting application). A similar application of noise flooding for BGV was given in [17], which considered the threshold setting of $t < n/3$ via Shamir sharing. This enabled active security, without needing to resort to zero-knowledge proofs. In our work we shall adopt the methodology of [17] for our main threshold decryption protocol.

A generic thresholdizer for arbitrary protocols was given by Boneh et al. in [12] using threshold-FHE. The construction of Boneh et al. utilizes a special form of secret sharing called $\{0, 1\}$-LSSS, which is closely related to replicated sharing.

All of these prior works utilized noise flooding as a methodology. As remarked above this requires a super-polynomial gap between the bound on the noise term e and the ciphertext modulus q. Such super-polynomial blow-ups in other areas of cryptography based on LWE have recently been avoided by utilizing the Renyi divergence [6]. This, as an approach to threshold-FHE, was recently examined by [13] and [18].

The problem with using the Renyi divergence in the context of distributed decryption is that the general technique of Renyi divergence is hard to apply to security problems which are inherently about distinguishing one distribution from another. In [18]

and [13] a way around this was found by designing special security games for threshold-FHE usage, which enabled the use of the Renyi divergence. The problem is that these games need to cope with the homomorphic nature of the underlying encryption scheme, and thus cannot be adaptive. In the applications (such as to MPC) mentioned above we really require a threshold-FHE protocol which is indistinguishable, to an adversary, with a simulation interacting with an ideal functionality. The security games presented in [18] and [13] do not allow such a usage.

Another approach is to apply generic MPC to the problem of threshold decryption of FHE ciphertexts. Here one avoids the noise flooding operation, and one executes the rounding operation inherent in decryption via a generic MPC protocol. This is relatively straight forward to implement using modern LSSS-based MPC systems, however the round complexity is very high. This can be a problem when entities are separated by large distances.

Thus we are led back to considering noise flooding. However, as detailed above, for FHE schemes such as BGV and BFV this is not a problem. The only issue comes with schemes such as TFHE, which utilize small parameters in order to achieve very fast bootstrapping operations. However, perhaps the very fast bootstrapping operation itself can be used to solve the problem?

1.2 Our Contribution

We present a simple method for threshold decryption for TFHE ciphertexts in the presence of $t < n/3$ actively (but statically) corrupted adversarial parties. Our methodology produces a threshold decryption functionality which is in the simulation paradigm, this makes it more amenable to being used as a black box in larger protocols than the game-based approaches based on Renyi divergence.

Our approach works for arbitrary prime power values of q, including the important case of $q = 2^{64}$. Adapting it to the case of non-prime power values of q is immediate via the Chinese-Remainder-Theorem. In doing so we utilize the (relatively standard) trick of applying Shamir secret sharing over Galois rings [1], thus we do not need to go via a replicated style secret sharing. In Shamir secret sharing the share sizes do not grow exponentially with the value of $\binom{n}{t}$.

When $\binom{n}{t}$ is "small" we apply a trick, which first appeared in [17], to enable threshold-FHE using a modified Pseudo-Random Secret Sharing (PRSS). In such a situation our protocol is a simple one round protocol, which is robust[1] and works over asynchronous networks when $t < n/3$. When $t < n/2$ we note that we obtain a non-robust protocol, but one which has active-with-abort security. The proof of security in [17] has a number of minor bugs/missing details in it, and is overly complex, thus we also re-prove the main threshold-FHE result from this paper. It turns out that adding two PRSS values for small values of $\binom{n}{t}$, automatically means we are adding a sum of at least two uniform distributions in the flooding term; and thus we can apply our improved statistical distance analysis in this case.

When $\binom{n}{t}$ is large we require slightly more work. In particular we divide our threshold-FHE protocol into two phases, an online and an offline phase. In the offline phase a "generic" MPC protocol is used to generate random shares of bits, which are used to produce two uniformly random noise flooding terms of the correct size. Thus again, we are able to apply our improved statistical distance analysis in this case. In the online phase we consume these random shares of bits to perform the threshold-FHE operation. The online phase is again robust and works over asynchronous networks, when $t < n/3$; and is only active-with-abort secure when $t < n/2$. The security properties of the offline phase are inherited from the underlying MPC protocol used to generate the shares of random bits; if the underlying MPC protocol is robust over asynchronous networks then so is the offline phase of our threshold-FHE protocol; if it only provides active-with-abort security over synchronous networks then they are the properties of our offline phase.

Our methodology for threshold-FHE follows in two conceptually simple steps:

(1) We take the input ciphertext with LWE parameters (ℓ, q) and then transform this into a ciphertext with LWE parameters with slightly larger parameters (L, Q) which encrypts the same message, where Q is a prime power with $q|Q$, and with relatively small noise. This switching to larger parameters is performed during a bootstrapping operation, which enables us to simultaneously reduce the noise, so that the noise gap is sufficiently large. We call this operation Switch-n-Squash, as it both switches the (ℓ, q) values, and also squashes the noise.
(2) We apply the traditional noise flooding operation, followed by a robust opening procedure on the secret shared value.

In practice the value q will be 2^{64}, and we will only need to boost the modulus to a value of $Q = 2^{128}$ in order to have a sufficient noise gap to perform threshold decryption. With such a value of Q it turns out that TFHE bootstrapping is still efficient, and thus the entire threshold decryption process is efficient. In particular it is very low round (requiring only one round in the online phase), thus it is also preferable to techniques based on generic MPC.

Note, the noise-to-modulus ratio after our Switch-n-Squash operation is much smaller. This is the key fact which enables our threshold decryption operation to proceed. That such a smaller ratio still maintains security is because the dimension has increased from ℓ to L.

In the special case when $\binom{n}{t}$ is small (say less than 100) we in addition obtain a one-round, threshold decryption protocol which is robustly secure when $t < n/3$, with no offline phase, and which assumes only asynchronous, as opposed to synchronous, networks.

2 PRELIMINARIES

2.1 Notation

Our basic input ciphertexts will come with ciphertext modulus q, and plaintext modulus p. For the underlying bootstrapping keys for TFHE we will utilize a cyclotomic ring of two-power degree N. The ring we define as

$$\mathcal{R} = \mathbb{Z}[X]/(X^N + 1),$$

with the reduction modulo the ciphertext (resp. plaintext) modulus q (resp. p) being given by

$$\mathcal{R}_q = (\mathbb{Z}/q\mathbb{Z})[X]/(X^N + 1) \quad (\text{ resp. } \mathcal{R}_p = (\mathbb{Z}/p\mathbb{Z})[X]/(X^N + 1)\ ).$$

[1] i.e. it outputs the correct decryption even in the presence of malicious parties.

We fix the global Δ as $\Delta = \lfloor q/p \rfloor$. This is the ratio between the ciphertext modulus q, and the *application* plaintext modulus p.

Elements in $\mathcal{R}$ (resp. $\mathcal{R}_q$, $\mathcal{R}_p$, etc) will be considered as vectors $\mathbf{A}$, $\mathbf{B}$, etc where we apply the component-wise addition operation. Multiplication, however, is performed with respect to the ring multiplication operation. Normal vectors, i.e. non-ring elements, will be written with lower case boldface, $\mathbf{a}$, $\mathbf{b}$, etc.

We let $\mathbf{a}[i]$ denote the i-th component of the vector $\mathbf{a}$, and $\mathbf{A}[i]$ denote the i-th coefficient of the ring element $\mathbf{A}$ when considered in the polynomial embedding. We assume the underlying ring is obvious from the context.

Multiplication of vectors $\mathbf{a} \cdot \mathbf{b}$ is assumed to be the normal dot-product, which results in a scalar value. We abuse notation by allowing $\mathbf{A} \leftarrow \mathbf{a}$ to denote a ring element is defined from a vector $\mathbf{a}$ of the same size. Thus, if $\mathbf{a} = (a_0, \ldots, a_{N-1})$ then we have

$$\mathbf{A} = a_0 + a_1 \cdot X + \cdots + a_{N-1} \cdot X^{N-1}.$$

2.2 Statistical Distance

Let $U(-B, B)$ denote the uniform distribution on the integer interval $(-B, \ldots, B]$ and $U(-B, B)^m$ be m samples from the respective distribution. Define $\Delta_{SD}(D_1, D_2)$ as the standard statistical distance between two distributions D_1 and D_2 which are defined over a common domain X, i.e.

$$\Delta_{SD}(D_1, D_2) = \frac{1}{2} \sum_{x \in X} |D_1(x) - D_2(x)|.$$

Security for our threshold-FHE protocol when $\binom{n}{t}$ is small will rely on the following Lemmas, all of which are variants of the standard Smudging Lemma (see for example Lemma 2.1 of [5])

Lemma 2.1 (Standard Smudging Lemma). *Let $e \in \mathbb{Z}$ and $B, m \in \mathbb{N}$ denote fixed integers, then we have*

$$\Delta_{SD}\left((e + U(-B, B))^m \ , \ U(-B, B)^m\right) \leq \frac{m \cdot |e|}{B},$$

From the data processing inequality, which says that the statistical distance between two distributions cannot increase by applying any (possibly randomized) function to them, one can immediately deduce

Lemma 2.2. *Let $e \in \mathbb{Z}$ and $B, m, v \in \mathbb{N}$ denote fixed integers, then we have*

$$\Delta_{SD}\left((e + \sum_{i=1}^{v} U(-B, B))^m \ , \ \sum_{i=1}^{v} U(-B, B)^m\right) \leq \frac{m \cdot |e|}{B},$$

However, a more accurate estimation, when $v \geq 2$, can be given by Lemma 2.4, which follows, via the data processing inequality, from the following Lemma, whose proof is given in the full version.

Lemma 2.3. *Let $e \in \mathbb{Z}$ and $B, m \in \mathbb{N}$ denote fixed integers, and let $\mathcal{P} = U(-B, B) + U(-B, B)$. Then*

$$\Delta_{SD}(\mathcal{P}^m, (e + \mathcal{P})^m) \leq \frac{m \cdot |e|}{B^2} + \sqrt{m \cdot \frac{|e|^2 \cdot \log B + 2}{2 \cdot (B^2 + B)}}.$$

Lemma 2.4. *Let $e \in \mathbb{Z}$ and $B, m, v \in \mathbb{N}$ denote fixed integers with $v \geq 2$, then we have,*

$$\Delta_{SD}\left((e + \sum_{i=1}^{v} U(-B, B))^m \ , \ \sum_{i=1}^{v} U(-B, B)^m\right)$$

$$\leq \frac{m \cdot |e|}{B^2} + \sqrt{m \cdot \frac{|e|^2 \cdot \log B + 2}{2 \cdot (B^2 + B)}},$$

In our application we always utilize $v \geq 2$, in which case we apply Lemma 2.4. When we apply this for m distributed decryption queries we are actually sampling a different value of e per query. On each application, the specific e value used is the output noise term from a bootstrapping operation for a given input ciphertext. Thus the above distances are simplified, upper bounds in our application scenario of the actual statistical distances between the various distributions we analyze.

In our application we will set $B = 2^{\mathsf{stat}} \cdot |e|$, where $\mathsf{stat} = 40$, since Lemma 2.4 tells us that distinguishing the two distributions (for fixed e) requires around

$$\frac{B^2}{|e|^2 \cdot \log B} = \frac{2^{2 \cdot \mathsf{stat}} \cdot |e|^2}{|e|^2 \cdot (\mathsf{stat} + \log |e|)} = \frac{2^{2 \cdot \mathsf{stat}}}{\mathsf{stat} + \log |e|} \approx 2^{2 \cdot \mathsf{stat}}$$

samples.

2.3 Learning-With-Errors (LWE)

The (decision) LWE problem is to distinguish between samples drawn from the two distributions

$$D_1 = \{ \ (\mathbf{a}, b) \ : \ \mathbf{a} \leftarrow \mathbb{Z}_q^\ell, \ b \leftarrow \mathbb{Z}_q \ \},$$
$$D_2 = \{ \ (\mathbf{a}, b) \ : \ \mathbf{a} \leftarrow \mathbb{Z}_q^\ell, \ e \leftarrow \mathcal{D}, \ b = \mathbf{a} \cdot \mathbf{s} + e \ \},$$

where $\mathbf{s} \in \mathbb{Z}_q^\ell$ is a fixed (secret) value, and $\mathcal{D}$ is the LWE-error distribution. In practice $\mathcal{D}$ is usually a discrete form of the Gaussian distribution with "small" standard deviation. For appropriate values of the parameters (q, ℓ) the problem is believed to be hard.

The Ring-LWE problem we define as trying to distinguish the two distributions

$$D_1 = \{ \ (A, B) \ : \ A, B \leftarrow \mathcal{R}_q \ \},$$
$$D_2 = \{ \ (A, B) \ : \ A \leftarrow \mathcal{R}_q, \ E \leftarrow \mathcal{D}_{\mathcal{R}}, \ B = A \cdot S + E \ \},$$

where $S \in \mathcal{R}_q$ is a fixed (secret) value, and $\mathcal{D}_{\mathcal{R}}$ is the Ring-LWE-error distribution on elements of $\mathcal{R}$. We can think of the Ring-LWE problem as being a special version of the LWE problem in which N LWE samples of dimension N are obtained on every iteration.

To enable easier selection of such parameters we approximate the required standard deviation for the distribution $\mathcal{D}$ for given values of q and ℓ, and a given security level. In this work we select an LWE security level of 128, namely distinguishing D_1 from D_2 should require a work effort of 2^{128}. We represent this function of the standard deviation for 128-bit security, as a function of q and ℓ, as the function $\Sigma_{\mathsf{LWE}}(q, \ell)$. This function can be approximated by fitting curves to the output of the LWE-estimator [2], when $\mathcal{D}$ is a discrete Gaussian distribution. One should strictly speaking give a separate approximation for each value of q, but it turns out for the two values of q which are important to us, namely $q = 2^{64}$ and $q = 2^{128}$, the same approximation can be used. For $\ell \geq 450$ we have

the approximation[2]

$$\alpha = -0.02659946234310527,$$
$$\beta = 2.98154318414599,$$
$$\Sigma_{\mathsf{LWE}}(q, \ell) = \max(q \cdot 2^{\alpha \cdot \ell + \beta}, 4).$$

We insert a minimum standard deviation of four into the approximation function to avoid problems when ℓ is very, very large.

2.4 TFHE

Our basic input TFHE ciphertext will be of the form $(\mathbf{a}, b)$ where $\mathbf{a} \in \mathbb{Z}_q^\ell$ and $b \in \mathbb{Z}_q$ such that

$$b = \mathbf{a} \cdot \mathbf{s} + e + \Delta \cdot m$$

for the message $m \in \mathbb{Z}_p$ and a noise value e. We assume the plaintext space $p = 2^{\varrho+1}$, where ϱ is the number of bits of plaintext and we add one bit to enable efficient non-negacylic operations. For our purposes we will not require specific details of the operations on TFHE ciphertexts, however we will require a detailed understanding of the associated noise growths in each operation. For the reader interested in the specific algorithm details we refer to [16] (also [14] and [15]) as well as the details of how the following noise formulae are derived.

The operations we will perform (*modulus switch*, *keyswitch* and *bootstrap*) may require additional encryptions of the secret keys with respect to different LWE-style encryption schemes. Thus we have to also keep track of the different types of ciphertexts which each operation is performed on. For our purposes we can focus on just the basic LWE ciphertexts as above, plus a so-called "flattened-GLWE" ciphertext, or F-GLWE, which one can think of as a normal LWE ciphertext but with dimension $w \cdot N$, for the ring-LWE dimension N used in the GLWE ciphertexts and w an associated parameter. GLWE ciphertexts are a generalization of the RLWE ciphertexts introduced above.

For simplicity in this paper we present the noise formulae only for the case of q a power of two. This is the main application area of our work; small changes are needed for other prime power values of q. Note that with q and p both powers of two we have that p exactly divides q, which is what makes the noise formulae slightly easier to describe. We note that all the following operations are deterministic in nature; thus every party executing these operations will produce the same output; this assumes that the parties execute the same Fast Fourier Transform (FFT) algorithms internally to multiply polynomials and are working on identical hardware. This requirement of operating the same FFT algorithm on identical hardware can be relaxed, see Section 5.7 of [11].

2.4.1 Modulus Switch. This operation takes an LWE ciphertext, with ciphertext modulus q, and switches it to an LWE ciphertext with modulus $2 \cdot N$. This algorithm is never *explicitly* called by our algorithms, however it is the first stage of bootstrapping and thus we do need to take into account the noise added by this operation in our analysis. This algorithm will be correct (with probability pr_{MS}) if we have that

$$c_{MS} \cdot \sqrt{\sigma^2 + \sigma_{MS}^2} < \frac{\Delta}{2} \tag{1}$$

where

$$\mathrm{pr}_{MS} = 1 - \mathrm{erfc}\left(\frac{c_{MS}}{\sqrt{2}}\right) = \mathrm{erf}\left(\frac{c_{MS}}{\sqrt{2}}\right)$$

and

$$\sigma_{MS}^2 = \frac{q^2}{48 \cdot N^2} - \frac{1}{12} + \ell \cdot \left(\frac{q^2}{96 \cdot N^2} + \frac{1}{48}\right),$$

where erf (resp. erfc) is the error (resp. complementary error) function. The function $\mathrm{erfc}(x)$ measures the chance of a Gaussian variable with zero mean and variance $\sigma = 0.5$ (or standard normal distribution) to fall outside the bounds $[-x, x]$. We will take $c_{MS} \approx 7.2$ in our analysis, leading to an error probability of $\mathrm{erfc}(7.2/\sqrt{2}) = 2^{-40}$ on homomorphic operations.

2.4.2 Key Switch. The KeySwitch operation takes a F-GLWE ciphertext and returns a normal LWE ciphertext. We require this operation as the bootstrapping operation below produces an F-GLWE ciphertext, and we need to translate it back to a standard LWE ciphertext for further processing by our algorithm. The output noise variance is $\sigma^2 + \sigma_{KS}^2$, where

$$\begin{aligned}
\sigma_{KS}^2 &= w \cdot N \cdot \left(\frac{q^2}{12 \cdot \beta_{ksk}^{2 \cdot v_{ksk}}} - \frac{1}{12}\right) \cdot (\mathrm{Var}(s_i) + \mathbb{E}^2(s_i)) \\
&\quad + \frac{w \cdot N}{4} \cdot \mathrm{Var}(s_i) + w \cdot N \cdot v_{ksk} \cdot \sigma_{ksk}^2 \cdot \left(\frac{\beta_{ksk}^2 + 2}{12}\right) \\
&= \frac{w \cdot N}{2} \cdot \left(\frac{q^2}{12 \cdot \beta_{ksk}^{2 \cdot v_{ksk}}} - \frac{1}{12}\right) \\
&\quad + w \cdot N \cdot \left(\frac{1}{16} + v_{ksk} \cdot \sigma_{ksk}^2 \cdot \left(\frac{\beta_{ksk}^2 + 2}{12}\right)\right) \\
&= w \cdot N \cdot \left(\frac{q^2}{24 \cdot \beta_{ksk}^{2 \cdot v_{ksk}}} + \frac{1}{48} + v_{ksk} \cdot \sigma_{ksk}^2 \cdot \left(\frac{\beta_{ksk}^2 + 2}{12}\right)\right)
\end{aligned}$$

since for a binary secret key we have $Var[s_i] = 1/4$ and $\mathbb{E}[s_i] = 1/2$. The values v_{ksk} and β_{ksk} are parameters associated with the key-switching keys, in particular how the decomposition gadget is formed. The value σ_{ksk} is the standard deviation used to generate the noise term in the key-switching keys. The latter is selected such that an LWE problem with dimension ℓ, modulus q and standard deviation for the noise term σ_{ksk} is hard to solve, i.e. $\sigma_{ksk} = \Sigma_{\mathsf{LWE}}(q, \ell)$.

2.4.3 Bootstrap. Bootstrapping takes an LWE ciphertext and outputs a F-GLWE ciphertext but with (potentially) smaller noise. The first thing a bootstrap operation performs is a modulus switch, therefore the input to the bootstrap operation (for it to be correct with a given probability) must satisfy equation (1). The specific details of how a bootstrap is performed is outside the scope of this paper, here we just describe its behavior. The noise output from bootstrap has variance σ_{BR}^2 where

$$\sigma_{BR}^2 = \ell \cdot \left(v_{bk} \cdot (w+1) \cdot N \cdot \left(\frac{\beta_{bk}^2 + 2}{12}\right) \cdot \sigma_{bk}^2 \right.$$

[2]The coefficients α and β were estimated with the commit made on January 5, 2023: https://github.com/malb/lattice-estimator/tree/f9f4b3c69d5be6df2c16243e8b1faa80703f020c

$$+\left(\frac{q^2-\beta_{bk}^{2\cdot v_{bk}}}{24\cdot\beta_{bk}^{2\cdot v_{bk}}}\right)\cdot\left(1+\frac{w\cdot N}{2}\right)$$
$$+\frac{w\cdot N}{32}+\frac{1}{16}\cdot\left(1-\frac{w\cdot N}{2}\right)^2\Bigg)$$

Again, the values v_{bk} and β_{bk} are parameters associated with the decomposition gadget associated to the bootstrapping keys, and the value σ_{bk} is the standard deviation used to generate the noise term in the bootstrapping keys. The latter is selected such that an LWE problem with dimension $w \cdot N$, modulus q and standard deviation for the noise term σ_{bk} is hard to solve, $\sigma_{bk} = \Sigma_{\mathsf{LWE}}(q, w \cdot N)$.

2.4.4 Refresh. Refresh is the key operation behind our method for threshold-FHE. It is the combination of bootstrap and keyswitch. We shall refer to this operation by the notation $(\mathbf{a}, b) \leftarrow \mathsf{Refresh}((\mathbf{a}', b'), \mathfrak{pk})$, where $\mathfrak{pk}$ is the public key. As such the operation will be correct (with a given probability) only if the input noise satisfies equation (1), with the output noise being $\sigma_{BR}^2 + \sigma_{KS}^2$.

2.5 Secret Sharing

We want to utilize basic Shamir secret sharing over the ring $\mathbb{Z}_Q$ for the larger ciphertext modulus Q. For ease of exposition we will assume that Q is a prime power, with the most challenging case being $Q = 2^K$. To cope with Q being a power of two we need to use Shamir sharing over Galois rings.

2.5.1 Galois Ring Structures. The use of Galois rings for Shamir sharing has a long history, going back to (at least) Serge Fehr's masters thesis [21]. For more modern usage see [1, 23]. We first need to fix a Galois ring extension, and write $Q = \mathfrak{p}^K$, where in our case of interest $\mathfrak{p} = 2$. We then define

$$d = \lceil \log_{\mathfrak{p}}(n+1) \rceil$$

this means that the finite field $\mathbb{F}_{\mathfrak{p}^d}$ contains at least $n+1$ values where n is the number of parties. Fix an irreducible polynomial $F(Y)$, of degree d, for this finite field

$$\mathbb{K} = \mathbb{F}_{\mathfrak{p}^d} = \mathbb{F}_{\mathfrak{p}}[Y]/F(Y).$$

Elements in $\mathbb{F}_{\mathfrak{p}^d}$ will be represented by polynomials of degree less than d in a formal root θ of $F(Y)$, i.e. we write $\gamma = c_0 + c_1 \cdot \theta + \cdots + c_{d-1} \cdot \theta^{d-1} \in \mathbb{F}_{\mathfrak{p}^d}$ with $c_i \in \mathbb{F}_{\mathfrak{p}}$. We shall use *the same* polynomial to define the Galois ring extension

$$\mathcal{G} = \mathbb{Z}_Q[\theta] = \mathbb{Z}_Q[Y]/F(Y).$$

We assume that $\mathbb{F}_{\mathfrak{p}^d}$ is embedded into $\mathcal{G}$ in the obvious way, and so can freely talk about elements in $\mathbb{F}_{\mathfrak{p}^d}$ as if they are also in $\mathcal{G}$. Note, in the case where $Q = \mathfrak{p}$ (i.e. Q is a "large" prime) we have that $\mathcal{G} = \mathbb{F}_Q$.

We enumerate the non-zero elements in $\mathbb{F}_{\mathfrak{p}^d}$ as $\{\gamma_1, \ldots, \gamma_{\mathfrak{p}^d-1}\}$, and so for every player $\mathcal{P}_i$ we can refer to "their" element γ_i. Note, in the case where $Q = \mathfrak{p}$ we have $\gamma_i = i$ for $i \in [1, \ldots, n]$. Note that when thinking of the γ_i as elements of $\mathcal{G}$ we have that $\gamma_i - \gamma_j$ is invertible for every distinct pair (i, j). This allows us to define the following polynomials in $\mathcal{G}[X]$, for $i \in \{1, \ldots, n\}$.

$$\delta_i(X) = \prod_{j \neq i} \frac{X - \gamma_j}{\gamma_i - \gamma_j}.$$

Note that

(1) $\delta_i(\gamma_i) = 1$.
(2) $\delta_i(\gamma_j) = 0$, if $i \neq j$.
(3) $\deg \delta_i(X) = n - 1$.

More generally, $\delta_i(X)$ can be defined for any subset of at least $t+1$ players.

2.5.2 Shamir Sharing over $\mathbb{Z}_Q$. We now define a secret sharing scheme for elements $a \in \mathbb{Z}_Q$, given in Figure 2, which has threshold t out of n players. This means that the scheme perfectly hides a value if at most t parties combine their share, however if $t+1$ parties come together then the share value can be perfectly reconstructed (if no party deliberately introduces an error into their share value). We write $[a]^{\langle t,Q\rangle}$ to denote that a value $a \in \mathbb{Z}_Q$ is secret shared according to the sharing, and we write $[a]_i^{\langle t,Q\rangle} \in \mathcal{G}$ to denote player $\mathcal{P}_i$'s share. Note, that our sharing can also share elements in $\mathcal{G}$ and not just elements in $\mathbb{Z}_Q$, in which case upon opening such an element the opening procedure will abort.

The Secret Sharing Scheme $[x]^{\langle t,Q\rangle}$

Share(a): Given $a \in \mathbb{Z}_Q$ this produces a sharing, i.e. values $[a]_i^{\langle t,Q\rangle} \in \mathcal{G}$

(1) Generate a polynomial $g_a(X) \in \mathcal{G}[X]$ of degree at most t such that $g_a(0) = a$.
(2) Define $[a]_i^{\langle t,Q\rangle} = g_a(\gamma_i)$.

Open$([a]_1^{\langle t,Q\rangle}, \ldots, [a]_n^{\langle t,Q\rangle})$:

(1) Compute the polynomial
$$g_a(X) \leftarrow \sum_i [a]_i^{\langle t,Q\rangle} \cdot \delta_i(X).$$
(2) If $\deg g_a(X) > t$ then abort.
(3) If $g_a(0) \notin \mathbb{Z}_Q$ then abort.
(4) Return $g_a(0)$.

Figure 2: The Secret Sharing Scheme $[x]^{\langle t,Q\rangle}$.

Notice, that the opening algorithm, given in Figure 2, will abort if any of the n-parties send in a share value which is inconsistent. In addition it will abort if the shared value is not in $\mathbb{Z}_Q$, but in $\mathcal{G} \setminus \mathbb{Z}_Q$.

The secret sharing scheme is linear, namely given secret sharings $[a]^{\langle t,Q\rangle}$ and $[b]^{\langle t,Q\rangle}$ we can produce a secret sharing of the value $\alpha \cdot a + \beta \cdot b + \gamma$ for any values $\alpha, \beta, \gamma \in \mathbb{Z}_Q$ with no interaction. This is done by each party $\mathcal{P}_i$ computing

$$[\alpha \cdot a + \beta \cdot b + \gamma]_i^{\langle t,Q\rangle} \leftarrow \alpha \cdot [a]_i^{\langle t,Q\rangle} + \beta \cdot [b]_i^{\langle t,Q\rangle} + \gamma.$$

We shall write this as global operation in the notation

$$[\alpha \cdot a + \beta \cdot b + \gamma]^{\langle t,Q\rangle} \leftarrow \alpha \cdot [a]^{\langle t,Q\rangle} + \beta \cdot [b]^{\langle t,Q\rangle} + \gamma.$$

2.5.3 Error Correction Over Galois Rings. In this section we explain how to do Reed-Solomon error correction over the Galois ring $\mathcal{G}$ when $t < n/3$. The methodology is taken from [1, Figure 1].

The standard Berlekamp–Welch or Gao algorithms for error correcting Reed-Solomon codes over $\mathbb{K}$ take as input $(x_1, \ldots, x_n)$

where $x_i \in \mathbb{K}$. We denote this by RS-Decode$_{\mathbb{K}}(x_1, \ldots, x_n)$. It is assumed on input that $x_i = f(\gamma_i)$, for all except at most t values, and for a polynomial $f \in \mathbb{F}_{\mathfrak{p}^d}[X]$ of degree at most t. The "error" values x_i can either be incorrect values x_i or the $\perp$ symbol. One could think of $\perp$ as zero, but sometimes in decoding algorithms it is faster to keep data around which we know to be a definite error. The output of RS-Decode$_{\mathbb{K}}(x_1, \ldots, x_n)$ is the polynomial $f(X)$.

The Berlekamp–Welch and Gao algorithms can take an additional parameter r which specifies the maximum expected number of errors, with the algorithm returning $\perp$ if more than r errors are detected. In this context we write RS-Decode$^r_{\mathbb{K}}(x_1, \ldots, x_n)$, with $r = \perp$ denoting the usual operation of no assumption on the errors.

In our Galois ring we have a similar decoding problem but now we have $x_i = f(\gamma_i)$, where f is a polynomial in $\mathcal{G}[X]$ of degree at most t, and the γ_i have been (trivially) lifted from $\mathbb{F}_{\mathfrak{p}^d}$ to $\mathcal{G}$. Note that every element $\alpha \in \mathcal{G}$ can be written as

$$\alpha = a_0 + a_1 \cdot \mathfrak{p} + \cdots + a_{K-1} \cdot \mathfrak{p}^{K-1}$$

where $a_i \in \mathbb{F}_{\mathfrak{p}^d}$. We will write the polynomial f in a similar manner as

$$f(X) = f_0(X) + f_1(X) \cdot \mathfrak{p} + \cdots + f_{K-1}(X) \cdot \mathfrak{p}^{K-1}$$

and we will recover the f_i values recursively using the standard algorithm RS-Decode$_{\mathbb{K}}(x_1, \ldots, x_n)$ as a subroutine.

2.5.4 Robust Opening. We can now define an opening procedure called RobustOpen, which will robustly open the shared value, depending on the relationship between t and n, and the underlying network properties. When $d = t$ robust opening is only available when $t < n/3$. Note, that for asynchronous networks and $t = d < n/4$ we can execute less computational steps than for the case $t = d < n/3$ by simply waiting for more data to arrive. The method for asynchronous networks and $t = d < n/3$ is called "online error correction", and was first presented in [9].

Assuming the input sharing is of an element in $\mathbb{Z}_Q$ then robust open protocol will output the value in $\mathbb{Z}_Q$, even if adversarial parties introduce errors. This is despite the shares themselves, and the Lagrange interpolation coefficients, being defined by elements in $\mathcal{G}$.

3 THE Switch-n-Squash OPERATION

The first step in our threshold decryption operation is to take an LWE ciphertext $(\mathbf{a}, b)$ defined with parameters (q, ℓ), with respect to a secret key $\mathbf{s} \in \{0,1\}^\ell$, and with noise variance σ^2. Then we switch it to a ciphertext $(\mathbf{a}', b')$ defined for parameters (Q, L) with $q|Q$, $L > \ell$, and for a secret key $\mathbf{s}' \in \{0,1\}^L$, and with a new noise variance σ'^2_{BR}, for a suitably small noise variance. Thus we increase both the ciphertext modulus and the LWE dimension, but we also increase the noise gap. We perform this switch by performing a bootstrapping operation which outputs a ciphertext with an LWE dimension $L = \omega \cdot N$ and ciphertext modulus Q. Indeed, one can see the entire method a just bootstrapping, with specially designed bootstrapping keys in order to result in a ciphertext with output parameters (L, Q).

The reason for moving a ciphertext from parameter set (q, ℓ) to (Q, L) is to enable us to have a lot more room between the noise bound and the value of $\Delta' = \frac{Q}{p}$. In particular the noise-gap should be big enough to enable noise flooding for threshold decryption. Thus, we need to select large enough cryptographic parameters to enable this refresh operation to output a suitably small noise value.

If our input ciphertext with parameter set (q, ℓ) has noise variance σ^2, then, after the modulus switch inside the bootstrap, we obtain a ciphertext with parameter set $(2 \cdot N, \ell)$ with noise variance

$$\sigma'^2 = \sigma^2 + \sigma^2_{MS}$$

To guarantee correctness up to a probability of failure pr_{MS}, we need the condition in equation (1) to be met. After the bootstrapping, we end up with a ciphertext with parameter set (Q, L) with noise variance σ'^2_{BR}, with σ'_{BR} a function of L, Q, N', w', σ'_{bk} etc as described earlier in the case of (q, l). We use N', w' etc to differentiate these values from the "normal" values used in standard FHE operations.

In the next section we will require the following equations to be satisfied, for some integer parameter pow. The parameter pow denotes the extra factor of noise we will add during flooding, i.e. it is approximately $\log_2 |E/e|$. We make it slightly larger than stat (by an extra additive term of $\log_2 100$) in order to cope with a non-uniform value of E which will be used in our procedure when $\binom{n}{t}$ is small (see later for a further discussion of this case).

$$\begin{aligned} \mathrm{Bd} &= c_{Dec} \cdot \sigma'_{BR}, \\ 2^{\mathrm{pow}+1} \cdot \mathrm{Bd} &\leq \frac{\Delta'}{2}, \\ \mathrm{pow} &\geq \mathrm{stat} + \log_2 100, \end{aligned}$$

where $c_{Dec} \approx 7.2$. Hence, combining these all together we have that

$$\mathrm{stat} + \log_2 100 \leq \mathrm{pow} \leq \log_2\left(\frac{\Delta'}{2}\right) - \log_2\left(c_{Dec} \cdot \sigma'_{BR}\right).$$

In particular this means that we must have

$$\log_2\left(c_{Dec} \cdot \sigma'_{BR}\right) \leq \log_2\left(\frac{\Delta'}{2}\right) - \mathrm{stat} - \log_2 100 - 1. \tag{2}$$

Given stat ≈ 40, we thus need to select parameters so that the noise after bootstrapping for these large parameters is at least stat bits smaller than the decryption correctness bound of $\Delta'/2$.

To find cryptographic parameters that guarantee the correctness, the efficiency and the security, we used the optimization method introduced in [11]. In a nutshell, it consists into solving the following optimization problem. We aim to minimize the function $\left(\mathrm{Cost}\,(BS)\right)$ which is a surrogate of the execution time of the bootstrapping as defined in [11], subject to the two constraints

$$\begin{aligned} c_{MS} \cdot \sqrt{\sigma^2 + \sigma^2_{MS}} &< \frac{\Delta}{2}, \\ \log_2\left(c_{Dec} \cdot \sigma'_{BR}\right) &\leq \log_2\left(\frac{\Delta'}{2}\right) - \mathrm{stat} - \log_2 100 - 1, \end{aligned}$$

where σ^2 is the variance of the input ciphertext, $\Delta = \frac{q}{p}$, $\Delta' = \frac{Q}{p}$, and $c_{MS} = c_{Dec} \approx 7.2$.

A summary of four potential parameter sets are given in Table 1. We give four sets of parameters; two for each plaintext size of $\varrho = 1$ and $\varrho = 4$, and for each plaintext size we give a variant with ℓ a non-power of two and ℓ a power of two. The former for use with the "traditional" methodology of giving out many encryptions of zero,

and the latter for use with the more compact public key encryption methodology given in [24].

Table 1: Parameters for switching up operations with the four sets of basic parameters.

	$\varrho = 1$	$\varrho = 4$	$\varrho = 1$	$\varrho = 4$
(q, ℓ)	$(2^{64}, 777)$	$(2^{64}, 870)$	$(2^{64}, 1024)$	$(2^{64}, 1024)$
(Q, L)	$(2^{128}, 4096)$	$(2^{128}, 4096)$	$(2^{128}, 4096)$	$(2^{128}, 4096)$
pow	47	47	47	47
N'	1024	2048	1024	2048
w'	4	2	4	2
β'_{bk}	2^{32}	2^{32}	2^{32}	2^{32}
v'_{bk}	2	2	2	2
$\log_2 \sigma'_{bk}$	22.0	22.0	22.0	22.0
$\log_2 \sigma'_{BR}$	72.0	72.1	72.2	72.2

We see that the standard deviation of the output noise after bootstrapping is around 2^{72}, which is gives us around 50 bits of noise gap for a ciphertext modulus of 2^{128}. Which is enough to fit in our flooding by a value of approximately $72 + 40 = 112$ bits. Note that for the input ciphertext, with parameters (q, ℓ), the noise gap is with overwhelming probability much smaller than 2^{50}, indeed it is less than 2^{10}.

4 THRESHOLD DECRYPTION OPERATION

After applying the methods from the previous sections we now have a ciphertext

$$(\mathbf{a}, b) = (\mathbf{a}, \mathbf{a} \cdot \mathbf{s}' + e + \Delta' \cdot m)$$

where $\mathbf{a} \in \mathbb{Z}_Q^L$, $\mathbf{s}' \in \{0, 1\}^L$, the message m lies in $\mathbb{Z}_p$, and $\Delta' = Q/p$, and e is a noise term. The noise term is assumed to have variance ${\sigma'_{BR}}^2$, i.e. the LWE ciphertext instance is an output of the Refresh operation from the previous section. In what follows we shall assume $|e| \leq \mathsf{Bd}$, where we assume (with overwhelming probability) that

$$\mathsf{Bd} = c_{Dec} \cdot \sigma'_{BR},$$

where $c_{Dec} \approx 7.2$. To fix ideas think of $Q = 2^{128}$ and the variance being of size roughly 2^{140}, and so $\mathsf{Bd} \approx 2^{70}$. Thus we have a noise gap of around 50 bits (assuming a plaintext space of at most 10 bits).

We assume the secret key $\mathbf{s}'$ has been secret shared with respect to our secret sharing scheme, i.e. we have a sharing $[\mathbf{s}']^{\langle t, Q \rangle}$. Formally we define the threshold decryption for the parameters (Q, L) via two ideal functionalities. The first $\mathcal{F}_{\mathsf{KeyGen}}$, in Figure 3, acts as a set-up assumption for our protocol, needed for the UC proof we provide. It generates a key pair, and secret shares the secret key among the players using the secret sharing scheme. One can realize this functionality using a generic MPC protocol, see the full version for an outline. Note, despite wanting active security we do not "complete" adversarially input shares into a complete sharing (as is often done in such situations), as the implementing actively protocol for $\mathcal{F}_{\mathsf{KeyGen}}$ does not actually need to do this.

The key functionality we want to implement is $\mathcal{F}_{\mathsf{KeyGenDec}}$ given in Figure 4. Note, that this functionality always returns the correct result, irrespective of what the adversary does.

Our threshold decryption protocol comes in two flavours, one where $\binom{n}{t}$ is "small" and one where $\binom{n}{t}$ is "large"[3]. When $\binom{n}{t}$ is small our threshold decryption protocol requires only one round of interaction, whilst when $\binom{n}{t}$ is large the online phase of our threshold decryption still requires only one round, however there is a (slightly) complex, ciphertext independent, offline phase which needs to be completed first.

In both cases we assume $t < n/3$, as we wish to have a robust asynchronous threshold decryption protocol; at least in the online phase of our protocol. We also recall we require that the protocol's security should come via a simulation, as opposed to a game based argument. This is to enable composition of the threshold decryption protocol easily within other larger protocols.

[3]Think of the small/large regime being divided at a value such as 100

$\mathcal{F}_{\mathsf{KeyGen}}$

Init():
(1) Execute $(\mathfrak{pk}, \mathbf{s}') \leftarrow \mathsf{KeyGen}(1^\kappa)$ for the underlying TFHE encryption scheme with parameters (Q, L).
(2) Generate a secret sharing $[\mathbf{s}']^{\langle t, Q \rangle}$ of the secret key.
(3) Send $\mathfrak{pk}$ to all players (including the adversary), and send $[\mathbf{s}']_i^{\langle t, Q \rangle}$ to player $\mathcal{P}_i$ (including adversarially controlled players).

Figure 3: The ideal functionality for distributed key generation

$\mathcal{F}_{\mathsf{KeyGenDec}}$

Init():
(1) Execute $(\mathfrak{pk}, \mathbf{s}') \leftarrow \mathsf{KeyGen}(1^\kappa)$ for the underlying TFHE encryption scheme with parameters (Q, L).
(2) Send $\mathfrak{pk}$ to all players, including the adversary and store the value $\mathbf{s}'$.

DistDecrypt($\mathfrak{ct}, \mathcal{U}$): For a ciphertext $\mathfrak{ct}$ with error e such that $|e| < \mathsf{Bd}$.
(1) Compute $m \leftarrow \mathsf{Dec}(\mathfrak{ct}, \mathbf{s}')$.
(2) If $\mathcal{U}$ is adversarially controlled then send $(\mathfrak{ct}, m)$ to the adversary.
(3) Otherwise send m to player $\mathcal{U}$ and $\mathfrak{ct}$ to the adversary.

Figure 4: The ideal functionality for distributed key generation and decryption

4.1 Threshold Decryption for "Small" $\binom{n}{t}$

We start with the case of $\binom{n}{t}$ being small, where we can utilize a variant of the standard Pseudo-Random Secret Sharing (PRSS),

originally introduced in [19]. The problem is that the complexity of a PRSS depends on $\binom{n}{t}$, which can become exponentially big as n increases. Thus this method can only be used when $\binom{n}{t}$ is small. We use a slightly modified form of PRSS in that we do not output sharings of uniformly random values from $\mathbb{Z}_Q$, but from a different range. This form of PRSS was originally used in [17], for exactly the purpose of threshold-FHE.

The algorithms for a non-interactive PRSS are defined in Figure 5. The algorithm PRSS.Init() iterates over all sets A of size $n-t$. Thus the complexity of PRSS, Init(), i.e. the number of sets A we need to deal with, depends on $\binom{n}{t}$, which can become very large for large n and t.

The PRSS makes use of a PRF ψ of the form

$$\psi : \begin{cases} \{0,1\}^{\mathsf{sec}} \times S & \longrightarrow & \mathbb{Z} \\ (\kappa, \mathsf{cnt}) & \longmapsto & \psi(\kappa, \mathsf{cnt}) \end{cases}$$

where $\{0,1\}^{\mathsf{sec}}$ is the keyspace and S is a set of counters. The output of the function ψ is assumed to be bounded in absolute value by

$$\mathsf{Bd}_1 = \frac{(2^{\mathsf{pow}} - 1) \cdot \mathsf{Bd}}{\binom{n}{t}},$$

recall that pow is roughly speaking $\log_2 |E/e|$.

One can implement ψ using AES in an obvious counter mode, e.g. as $\log_2 \mathsf{Bd}_1 < 256$ one can set

$$\psi(\kappa, \mathsf{cnt}) = \left(\mathsf{AES}_\kappa(0\|\mathsf{cnt}) + 2^{128} \cdot \mathsf{AES}_\kappa(1\|\mathsf{cnt})\right) \pmod{\mathsf{Bd}_1},$$

where we treat the output block of the AES cipher as an integer in $[0, \ldots, 2^{128} - 1]$. Note, the output of ψ is only statistically uniform in the required range here if $\log_2 \mathsf{Bd}_1 < 256 - \mathsf{stat}$, which will be true in our usage. Since the output of ψ is bounded as above, we have that the value E is bound by $(2^{\mathsf{pow}} - 1) \cdot \mathsf{Bd}$, as the sum used in the PRSS has at most $\binom{n}{t}$ terms. Note, the shared value which is output by the PRSS invocation is the sharing of the value

$$E \leftarrow \sum_A \psi(r_A, \mathsf{cnt}).$$

Given this PRSS we can define our threshold decryption protocol, which we give in Figure 6, where we assume a dedicated player $\mathcal{U}$ (possibly not one of the threshold decryption parties) will receive the final output. If all threshold decryption parties are to receive the output of the threshold decryption, or the output is to be public and not just to player $\mathcal{U}$, then the communication in step 3 does not need to be done securely.

For correctness we require that

$$2^{\mathsf{pow}+1} \cdot \mathsf{Bd} \leq \frac{\Delta'}{2},$$

since then the PRSS addition will not effect the correctness of the final result as $E + \mathsf{Bd} \leq 2 \cdot (2^{\mathsf{pow}} - 1) \cdot \mathsf{Bd} + \mathsf{Bd} = 2^{\mathsf{pow}+1} \cdot \mathsf{Bd} < \Delta'/2$.

On the other hand (see below) for security we require that

$$\mathsf{pow} \geq \mathsf{stat} + \log_2 \binom{n}{t},$$

where stat is the security parameter related to statistical distance. Thus this method is only applicable when we have a large Δ' in comparison to the noise bound Bd. This is why we needed to boost

PRSS

PRSS.Init(): For every set $A \subseteq \{1, \ldots, n\}$ of size $n - t$:

(1) $S \leftarrow \{\mathcal{P}_i\}_{i \in A}$.
(2) Players $\mathcal{P}_i$ with $i \in A$ execute $r_A \leftarrow$ AgreeRandom(S, sec), from the full version.
(3) Define $f_A(X) \in \mathbb{Z}_q[X] = \mathbb{Z}_{2^k}[X]$ to be the polynomial of degree t such that $f_A(0) = 1$ and $f_A(\gamma_i) = 0$ for all $i \notin A$. Each party $\mathcal{P}_i$ only needs store $f_A(\gamma_i)$ though.
(4) $\mathsf{cnt}_{\mathsf{PRSS}} \leftarrow 0$.

PRSS.Next():

(1) Party $\mathcal{P}_i$ computes, where the sum is over every set A containing i,
$$[E]_i^{\langle t,Q \rangle} \leftarrow \sum_{A: i \in A} \psi(r_A, \mathsf{cnt}_{\mathsf{PRSS}}) \cdot f_A(\gamma_i).$$
(2) $\mathsf{cnt}_{\mathsf{PRSS}} \leftarrow \mathsf{cnt}_{\mathsf{PRSS}} + 1$.
(3) Return $[E]^{\langle t,Q \rangle}$.

Figure 5: Pseudo-Random Secret Sharing PRSS

Threshold Decryption - Protocol 1

Init():

(1) The parties $\mathcal{P}_1, \ldots, \mathcal{P}_n$ execute PRSS.Init().
(2) The parties obtain $[\mathbf{s}']^{\langle t,Q \rangle}$ via a threshold key generation protocol, the full version.

DistDecrypt($\mathfrak{ct}$, $[\mathbf{s}']^{\langle t,Q \rangle}$, $\mathcal{U}$): On input of $\mathfrak{ct} = (\mathbf{a}, b) \in \mathbb{Z}_Q^{L+1}$ this executes the following steps:

(1) The parties $\mathcal{P}_i$ execute $[E]^{\langle t,Q \rangle} \leftarrow$ PRSS.Next() + PRSS.Next().
(2) The parties $\mathcal{P}_i$ compute $[v]^{\langle t,Q \rangle} \leftarrow b - \mathbf{a} \cdot [\mathbf{s}']^{\langle t,Q \rangle} + [E]^{\langle t,Q \rangle}$.
(3) Party $\mathcal{P}_i$ sends the value $[v]_i^{\langle t,Q \rangle}$ *securely* to the player $\mathcal{U}$.
(4) Player $\mathcal{U}$ applies algorithm RobustOpen to robustly reconstruct the value $b - \mathbf{a} \cdot \mathbf{s}' + E$, and hence m.

Figure 6: Threshold Decryption - Protocol 1

the ciphertext from one with parameters (q, ℓ) to one with parameters (Q, L) in the previous sections. Since we are assuming the small regime for $\binom{n}{t}$ is when $\binom{n}{t} \leq 100$, and we used the inequality

$$\mathsf{pow} \geq \mathsf{stat} + \log_2 100$$

in the previous section to derive the bounds on the noise after refreshing to ensure that

$$2^{\mathsf{pow}+1} \cdot \mathsf{Bd} \leq \frac{\Delta'}{2}$$

we are assured that the conditions of the following theorem are satisfied for our refresh parameters.

Simulator Threshold Decryption

On input of

(1) A ciphertext $\mathfrak{ct} = (\mathbf{a}, b)$ and a public key $\mathfrak{pk}$.
(2) The underlying message m encrypted by $\mathfrak{ct}$.
(3) A set of adversarial parties I with $|I| \leq t$.
(4) The share values $[\mathbf{s}']_i^{\langle t,Q \rangle}$ for $i \in I$.
(5) The PRSS secret keys r_A for all sets A such that $A \cap I \neq \emptyset$.

this algorithm outputs the simulated shares $\{[v]_j^{\langle t,Q \rangle}\}_{j \notin I}$.

Sim − DistDecrypt:

(1) The simulator computes, for $i \in I$,

$$[\hat{v}]_i^{\langle t,Q \rangle} = b - \mathbf{a} \cdot [\mathbf{s}']_i^{\langle t,Q \rangle} + \sum_{A: i \in A} (\psi(r_A, \mathsf{cnt}_{\mathsf{PRSS}}) + \psi(r_A, \mathsf{cnt}_{\mathsf{PRSS}} + 1)) \cdot f_A(i).$$

(2) The simulator computes

$$E' = \sum_{A: A \cap I \neq \emptyset} (\psi(r_A, \mathsf{cnt}_{\mathsf{PRSS}}) + \psi(r_A, \mathsf{cnt}_{\mathsf{PRSS}} + 1)) \quad + \sum_{B: B \cap I = \emptyset} (r_B + r'_B)$$

where r_B and r'_B are chosen uniformly at random so that $|r_B|, |r'_B| \leq \mathsf{Bd}_1$.

(3) The simulator computes $v = \Delta' \cdot m + E'$.
(4) The simulator generates the decryption shares $\{[\hat{v}]_j^{\langle t,Q \rangle}\}_{j \notin I}$ via Lagrange interpolation (and possibly generating random shares if $|I| < t$) from v and the values $\{[\hat{v}]_i^{\langle t,Q \rangle}\}_{i \in I}$.
(5) The simulator outputs $\{[v]_j^{\langle t,Q \rangle}\}_{j \notin I}$.

Figure 7: Simulator for $\mathsf{DistDecrypt}(\mathfrak{ct}, [\mathbf{s}']^{\langle t,Q \rangle}, \mathcal{U})$

Theorem 4.1. *Assuming*

$$\mathsf{pow} \geq \mathsf{stat} + \log_2 \binom{n}{t},$$

in the $\mathcal{F}_{\mathsf{KeyGen}}$*-hybrid model the protocol in Figure 6 implements* $\mathcal{F}_{\mathsf{KeyGenDec}}$ *with statistical security against any static active adversary corrupting I parties, with* $|I| \leq t$*, making at most* $2^{2 \cdot \mathsf{stat}}$ *threshold decryption queries.*

Assuming

$$2^{\mathsf{pow}+1} \cdot \mathsf{Bd} \leq \frac{\Delta'}{2},$$

the protocol is correct.

Proof. Correctness follows, even in the presence of $t < n/3$ fully malicious parties, on noticing that the bounds on the noise, described above, imply that the value v does encode the original message correctly when it is robustly opened.

Security of the protocol follows by showing that the output of simulator in Figure 7 is statistically indistinguishable, from the output of an adversary controlling I parties, with $|I| \leq t$, in a real execution of the protocol.

The proof of this security claim follows essentially the argument in Section 7.5 of the full version of [17]; where we have to switch from a BGV style of looking at ciphertexts to one of BFV. However, the proof in [17] is overly complex and has a few minor bugs, which we correct here.

First note, the values $\{[v]_j^{\langle t,Q \rangle}\}_{j \in I}$ produced by the simulator are the true decryption share values which the adversary should broadcast (even if he does not) if they acted honestly. The bounds on the noise described above then imply that the value v does encode the original message correctly. This means that the Lagrange interpolation in the simulation will recover the shares for the honest players $\{[v]_j^{\langle t,Q \rangle}\}_{j \notin I}$ as required.

Now, let e denote the value of $b - \mathbf{a} \cdot \mathbf{s}' - \Delta' \cdot m \pmod{Q}$.

In a real execution of the protocol the shares output by the honest players are consistent and are enough to allow the honest parties to decrypt correctly, since $t < n/3$. The simulation has exactly the same properties.

The value E is the value output by summing two executions of the PRSS in the real execution of the protocol, and the value E' is the value simulated for the sum of the two executions of the PRSS in the simulated protocol.

In the real protocol the adversary sees the value

$$\Delta' \cdot m + e + E$$

whereas in the simulated protocol he sees the value

$$\Delta' \cdot m + E'.$$

By the security of the PRSS the value of $e + E$ and $e + E'$ are indistinguishable. Thus we only need to show that $e + E'$ and E' are indistinguishable.

However, by Lemmas 2.4 and 2.3 (applied with $v = \binom{n}{t} \geq n$, $B = \mathsf{Bd}_1$, and $|e| = \mathsf{Bd}$) we have that, when executing at most d distributed decryption operations,

$$\Delta_{SD}\Big((e + \sum_{i=1}^{v} U(-B,B))^d \,,\, \sum_{i=1}^{v} U(-B,B)^d\Big) \leq c \cdot \sqrt{d} \cdot 2^{-\mathsf{stat}},$$

for some relatively 'small' constant c. Lemma 2.4/2.3 applies, since the number of uniform random variables $U(-B,B)$ added by the honest players is lower bounded by two (one from each PRSS evaluation). Thus to distinguish the two distributions the adversary would need to sample $d > 2^{2 \cdot stat}$ operations. □

4.2 Threshold Decryption for "Large" $\binom{n}{t}$

When $\binom{n}{t}$ is large we can no longer relay on a non-interactive PRSS. We can also not rely on "standard" interactive PRSS's, as our PRSS was used to create a small-ish element above and not a uniformly random one. Thus when $\binom{n}{t}$ is large we generate the masking value $[E]^{\langle t,Q \rangle}$ above, as a sum of two uniformly random values, using random bits provided by an "offline" phase. This offline phase is abstracted in the ideal functionality $\mathcal{F}_{\mathsf{Offline}}$ in Figure 8. We discuss how this can be implemented in the full version.

In Protocol 1 the value E was a sum of $\binom{n}{t}$ uniform random variables in $[-2^{\mathsf{pow}-1} \cdot \mathsf{Bd}, \ldots, 2^{\mathsf{pow}-1} \cdot \mathsf{Bd})$, only two of which had to be truly random to ensure security. In Protocol 2 the value E is selected by adding two values obtained uniformly from the range $[-2^B, \ldots, 2^B))$ where $B = \lceil \log_2 \mathsf{Bd} \rceil + \mathsf{pow}$. The full procedure

The Functionality $\mathcal{F}_{\text{Offline}}$

We describe this functionality as a robust ideal functionality, the modifications to make a functionality which is only secure in an active-with-abort setting are easily made.

$\mathcal{F}_{\text{Offline}}.\text{Bits}(b)$:

(1) The functionality samples uniformly random bits $b_i \in \{0,1\}$ for $i = 1, \ldots, b$.
(2) The functionality creates random sharings $[b_i]^{\langle t,Q \rangle}$ of these bits.
(3) The functionality distributes the shares $[b_i]_j^{\langle t,Q \rangle}$ to each player $\mathcal{P}_j$.

Figure 8: The Offline Functionality $\mathcal{F}_{\text{Offline}}$

is given in Figure 9; the call to the Offline procedure in lines 1 and 2, indicate that this is where data is obtained from the offline procedure. This "offline" operation can either be executed in place (in which case it is not offline but online) or it is the place where the data is fetched from the prior offline execution. The correctness and security of the protocol follows from similar (but simpler) arguments to those presented above.

Threshold Decryption - Protocol 2

Init():

(1) The parties obtain $[\mathbf{s}']^{\langle t,Q \rangle}$ via a threshold key generation protocol.

DistDecrypt($\mathfrak{ct}$, $[\mathbf{s}']^{\langle t,Q \rangle}, \mathcal{U}$): On input of $\mathfrak{ct} = (\mathbf{a}, b)$ this executes the following steps:

(1) $([b_i]^{\langle t,Q \rangle})_{i=0}^{B} \leftarrow \mathcal{F}_{\text{Offline}}.\text{Bits}(B+1)$.
(2) $([b'_i]^{\langle t,Q \rangle})_{i=0}^{B} \leftarrow \mathcal{F}_{\text{Offline}}.\text{Bits}(B+1)$.
(3) The parties $\mathcal{Q}_i$ compute $[E]^{\langle t,Q \rangle} \leftarrow (-2^B + \sum_{i=0}^{B} [b_i]^{\langle t,Q \rangle} \cdot 2^i) + (-2^B + \sum_{i=0}^{B} [b'_i]^{\langle t,Q \rangle} \cdot 2^i)$.
(4) The parties $\mathcal{Q}_i$ compute $[v]^{\langle t,Q \rangle} \leftarrow b - \mathbf{a} \cdot [\mathbf{s}']^{\langle t,Q \rangle} + [E]^{\langle t,Q \rangle}$.
(5) Party $\mathcal{Q}_i$ sends the value $[v]_i^{\langle t,Q \rangle}$ *securely* to the player $\mathcal{U}$.
(6) Player $\mathcal{U}$ applies algorithm RobustOpen to robustly reconstruct the value $b - \mathbf{a} \cdot \mathbf{s} + E$, and hence m.

Figure 9: Threshold Decryption - Protocol 2

5 EXPERIMENTS

We can now present our threshold decryption procedure for TFHE ciphertexts, which we give in Figure 10. Recall, from the introduction, for BGV or BFV ciphertexts by selecting parameters suitably or by bootstrapping, one can proceed directly to step 2 in Figure 10.

Complete Threshold Decryption

FullDistDecrypt($\mathfrak{ct}$, $[\mathbf{s}']^{\langle t,Q \rangle}, \mathcal{U}$): On input of $\mathfrak{ct} = (\mathbf{a}, b) \in \mathbb{Z}_q^{\ell+1}$ this executes the following steps:

(1) Execute $\mathfrak{ct}' \leftarrow$ Switch-n-Squash($\mathfrak{ct}$) to obtain $\mathfrak{ct}' \in \mathbb{Z}_Q^{L+1}$ encrypting the same value under the key $\mathbf{s}' \in \{0,1\}^L$, with noise with variance ${\sigma'_{BR}}^2$.
(2) Execute $m \leftarrow$ DistDecrypt($\hat{\mathfrak{ct}}$, $[\mathbf{s}']^{\langle t,Q \rangle}, \mathcal{U}$) to obtain m.

Figure 10: The complete threshold decryption protocol for TFHE ciphertexts

We note that line 1 of Figure 10 does not require interaction, whereas line 2 does. We thus first present experimental times for the evaluation of line 1 (Switch-n-Squash) for our four parameter sets. These we present in Table 2 of our Rust implementation. These results were obtained on an AWS `m6i.metal` instance with 128 Intel Xeon Gen 3 vCPUs and 512 GiB RAM, taking an average execution time over 100 runs of the relevant algorithms.

Table 2: Execution times (in milliseconds) for line 1 (Switch-n-Squash) of Figure 10

Parameters	Switch-n-Squash
$(2^{64}, 777) \rightarrow (2^{128}, 4096)$	241.01
$(2^{64}, 870) \rightarrow (2^{128}, 4096)$	265.80
$(2^{64}, 1024) \rightarrow (2^{128}, 4096)$	316.77
$(2^{64}, 1024) \rightarrow (2^{128}, 4096)$	314.19

Recall these timings are for a part of the computation which does not requite interaction, and which are amenable to acceleration by the FHE accelerators currently being developed. For example, the paper [8] shows a three orders of magnitude acceleration using only FPGA acceleration (as opposed to ASIC acceleration).

To time line 2 of Figure 10 we need to consider various other factors; the number of parties n performing the distributed decryption, the threshold t, the type of network, the number of active corruptions. For each of our four parameter sets we utilized three different sets of (n, t) values; namely $(n, t) = (4, 1), (10, 3)$ and $(40, 13)$. For the first of these one can utilize the PRSS-based distributed decryption method, for the other two one needs to utilize the methodology requiring an offline phase. In our experiments we only timed the online phase for the latter two cases. In all cases we present the average run-time over 1000 iterations for a single honest party. Recall this party will terminate as soon as it has received enough shares to robustly reconstruct the underlying encrypted value.

We also investigated the effect of a LAN-like setting (1 Gbit/s with small ping times of ≈1 ms) versus a WAN-like setting (100 Mbit/s with high ping times of ≈100 ms), and whether we are optimistic or pessimistic in terms of the number of errors introduced by the adversary during the distributed decryption. If there are no errors then the online-error correction method underlying RobustOpen will execute faster than if there are maximal, i.e. t, adversarial errors. The asynchronous channels are implemented using gRPC with `tokio` and `tonic` Rust crates.

We measured our experiments on a single AWS `m6i.metal` instance as above. We ran the n protocol parties as individual docker containers and simulated the LAN/WAN connection between them. Our full results are given in the full version.

In the most favorable situation, namely four parties where we can tolerate one dishonest party over a LAN, we obtain execution times for line 2 of Figure 10 of under 2 milliseconds. In the least favorable situation we investigated, namely 40 parties of which thirteen are malicious (and send invalid share values), and over a WAN, we are able to execute line 2 of Figure 10 in under 100 milliseconds.

6 ACKNOWLEDGEMENTS

The work of Nigel Smart on this work was supported by CyberSecurity Research Flanders (VR20192203), and by the FWO under an Odysseus project (GOH9718N). We would like to thank Fre Vercauteren for conversations during the course of this work.

REFERENCES

[1] Mark Abspoel, Ronald Cramer, Ivan Damgård, Daniel Escudero, and Chen Yuan. 2019. Efficient Information-Theoretic Secure Multiparty Computation over $\mathbb{Z}/p^k\mathbb{Z}$ via Galois Rings. In *TCC 2019: 17th Theory of Cryptography Conference, Part I (Lecture Notes in Computer Science, Vol. 11891)*, Dennis Hofheinz and Alon Rosen (Eds.). Springer, Heidelberg, Germany, Nuremberg, Germany, 471–501. https://doi.org/10.1007/978-3-030-36030-6_19

[2] Martin R. Albrecht, Rachel Player, and Sam Scott. 2015. On the concrete hardness of Learning with Errors. *J. Math. Cryptol.* 9, 3 (2015), 169–203. http://www.degruyter.com/view/j/jmc.2015.9.issue-3/jmc-2015-0016/jmc-2015-0016.xml

[3] Diego F. Aranha, Carsten Baum, Kristian Gjøsteen, and Tjerand Silde. 2022. Verifiable Mix-Nets and Distributed Decryption for Voting from Lattice-Based Assumptions. Cryptology ePrint Archive, Report 2022/422. https://eprint.iacr.org/2022/422.

[4] Gilad Asharov, Abhishek Jain, Adriana López-Alt, Eran Tromer, Vinod Vaikuntanathan, and Daniel Wichs. 2012. Multiparty Computation with Low Communication, Computation and Interaction via Threshold FHE. In *Advances in Cryptology – EUROCRYPT 2012 (Lecture Notes in Computer Science, Vol. 7237)*, David Pointcheval and Thomas Johansson (Eds.). Springer, Heidelberg, Germany, Cambridge, UK, 483–501. https://doi.org/10.1007/978-3-642-29011-4_29

[5] Gilad Asharov, Abhishek Jain, and Daniel Wichs. 2011. Multiparty Computation with Low Communication, Computation and Interaction via Threshold FHE. Cryptology ePrint Archive, Report 2011/613. https://eprint.iacr.org/2011/613.

[6] Shi Bai, Tancrède Lepoint, Adeline Roux-Langlois, Amin Sakzad, Damien Stehlé, and Ron Steinfeld. 2018. Improved Security Proofs in Lattice-Based Cryptography: Using the Rényi Divergence Rather than the Statistical Distance. *Journal of Cryptology* 31, 2 (April 2018), 610–640. https://doi.org/10.1007/s00145-017-9265-9

[7] Carsten Baum, Daniele Cozzo, and Nigel P. Smart. 2019. Using TopGear in Overdrive: A More Efficient ZKPoK for SPDZ. In *SAC 2019: 26th Annual International Workshop on Selected Areas in Cryptography (Lecture Notes in Computer Science, Vol. 11959)*, Kenneth G. Paterson and Douglas Stebila (Eds.). Springer, Heidelberg, Germany, Waterloo, ON, Canada, 274–302. https://doi.org/10.1007/978-3-030-38471-5_12

[8] Michiel Van Beirendonck, Jan-Pieter D'Anvers, and Ingrid Verbauwhede. 2022. FPT: a Fixed-Point Accelerator for Torus Fully Homomorphic Encryption. Cryptology ePrint Archive, Report 2022/1635. https://eprint.iacr.org/2022/1635.

[9] Michael Ben-Or, Ran Canetti, and Oded Goldreich. 1993. Asynchronous secure computation. In *25th Annual ACM Symposium on Theory of Computing*. ACM Press, San Diego, CA, USA, 52–61. https://doi.org/10.1145/167088.167109

[10] Rikke Bendlin and Ivan Damgård. 2010. Threshold Decryption and Zero-Knowledge Proofs for Lattice-Based Cryptosystems. In *TCC 2010: 7th Theory of Cryptography Conference (Lecture Notes in Computer Science, Vol. 5978)*, Daniele Micciancio (Ed.). Springer, Heidelberg, Germany, Zurich, Switzerland, 201–218. https://doi.org/10.1007/978-3-642-11799-2_13

[11] Loris Bergerat, Anas Boudi, Quentin Bourgerie, Ilaria Chillotti, Damien Ligier, Jean-Baptiste Orfila, and Samuel Tap. 2023. Parameter Optimization and Larger Precision for (T)FHE. *Journal of Cryptology* 36, 3 (2023), 28. https://doi.org/10.1007/s00145-023-09463-5

[12] Dan Boneh, Rosario Gennaro, Steven Goldfeder, Aayush Jain, Sam Kim, Peter M. R. Rasmussen, and Amit Sahai. 2018. Threshold Cryptosystems from Threshold Fully Homomorphic Encryption. In *Advances in Cryptology – CRYPTO 2018, Part I (Lecture Notes in Computer Science, Vol. 10991)*, Hovav Shacham and Alexandra Boldyreva (Eds.). Springer, Heidelberg, Germany, Santa Barbara, CA, USA, 565–596. https://doi.org/10.1007/978-3-319-96884-1_19

[13] Katharina Boudgoust and Peter Scholl. 2023. Simple Threshold (Fully Homomorphic) Encryption From LWE With Polynomial Modulus. Cryptology ePrint Archive, Report 2023/016. https://eprint.iacr.org/2023/016.

[14] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. 2020. TFHE: Fast Fully Homomorphic Encryption Over the Torus. *Journal of Cryptology* 33, 1 (Jan. 2020), 34–91. https://doi.org/10.1007/s00145-019-09319-x

[15] Ilaria Chillotti, Marc Joye, and Pascal Paillier. 2021. Programmable Bootstrapping Enables Efficient Homomorphic Inference of Deep Neural Networks. In *Cyber Security Cryptography and Machine Learning - 5th International Symposium, CSCML 2021, Be'er Sheva, Israel, July 8-9, 2021, Proceedings (Lecture Notes in Computer Science, Vol. 12716)*, Shlomi Dolev, Oded Margalit, Benny Pinkas, and Alexander A. Schwarzmann (Eds.). Springer, Be'er Sheva, Isreal, 1–19. https://doi.org/10.1007/978-3-030-78086-9_1

[16] Ilaria Chillotti, Damien Ligier, Jean-Baptiste Orfila, and Samuel Tap. 2021. Improved Programmable Bootstrapping with Larger Precision and Efficient Arithmetic Circuits for TFHE. In *Advances in Cryptology – ASIACRYPT 2021, Part III (Lecture Notes in Computer Science, Vol. 13092)*, Mehdi Tibouchi and Huaxiong Wang (Eds.). Springer, Heidelberg, Germany, Singapore, 670–699. https://doi.org/10.1007/978-3-030-92078-4_23

[17] Ashish Choudhury, Jake Loftus, Emmanuela Orsini, Arpita Patra, and Nigel P. Smart. 2013. Between a Rock and a Hard Place: Interpolating between MPC and FHE. In *Advances in Cryptology – ASIACRYPT 2013, Part II (Lecture Notes in Computer Science, Vol. 8270)*, Kazue Sako and Palash Sarkar (Eds.). Springer, Heidelberg, Germany, Bengalore, India, 221–240. https://doi.org/10.1007/978-3-642-42045-0_12

[18] Siddhartha Chowdhury, Sayani Sinha, Animesh Singh, Shubham Mishra, Chandan Chaudhary, Sikhar Patranabis, Pratyay Mukherjee, Ayantika Chatterjee, and Debdeep Mukhopadhyay. 2022. Efficient Threshold FHE with Application to Real-Time Systems. Cryptology ePrint Archive, Report 2022/1625. https://eprint.iacr.org/2022/1625.

[19] Ronald Cramer, Ivan Damgård, and Yuval Ishai. 2005. Share Conversion, Pseudorandom Secret-Sharing and Applications to Secure Computation. In *TCC 2005: 2nd Theory of Cryptography Conference (Lecture Notes in Computer Science, Vol. 3378)*, Joe Kilian (Ed.). Springer, Heidelberg, Germany, Cambridge, MA, USA, 342–362. https://doi.org/10.1007/978-3-540-30576-7_19

[20] Ivan Damgård, Valerio Pastro, Nigel P. Smart, and Sarah Zakarias. 2012. Multiparty Computation from Somewhat Homomorphic Encryption. In *Advances in Cryptology – CRYPTO 2012 (Lecture Notes in Computer Science, Vol. 7417)*, Reihaneh Safavi-Naini and Ran Canetti (Eds.). Springer, Heidelberg, Germany, Santa Barbara, CA, USA, 643–662. https://doi.org/10.1007/978-3-642-32009-5_38

[21] Serge Fehr. 1993. Span Programs over Rings and How to Share a Secret from a Module. Masters Thesis, ETH Zurich. https://crypto.ethz.ch/publications/Fehr98.html.

[22] Craig Gentry. 2009. *A fully homomorphic encryption scheme*. Ph. D. Dissertation. Stanford University. crypto.stanford.edu/craig.

[23] Robin Jadoul, Nigel P. Smart, and Barry Van Leeuwen. 2022. MPC for Q_2 Access Structures over Rings and Fields. In *SAC 2021: 28th Annual International Workshop on Selected Areas in Cryptography (Lecture Notes in Computer Science, Vol. 13203)*, Riham AlTawy and Andreas Hülsing (Eds.). Springer, Heidelberg, Germany, Virtual Event, 131–151. https://doi.org/10.1007/978-3-030-99277-4_7

[24] Marc Joye. 2023. TFHE Public-Key Encryption Revisited. Cryptology ePrint Archive, Paper 2023/603. https://eprint.iacr.org/2023/603 https://eprint.iacr.org/2023/603.

[25] Marcel Keller, Valerio Pastro, and Dragos Rotaru. 2018. Overdrive: Making SPDZ Great Again. In *Advances in Cryptology – EUROCRYPT 2018, Part III (Lecture Notes in Computer Science, Vol. 10822)*, Jesper Buus Nielsen and Vincent Rijmen (Eds.). Springer, Heidelberg, Germany, Tel Aviv, Israel, 158–189. https://doi.org/10.1007/978-3-319-78372-7_6

[26] Emmanuela Orsini, Nigel P. Smart, and Frederik Vercauteren. 2020. Overdrive2k: Efficient Secure MPC over $\mathbb{Z}_{2^k}$ from Somewhat Homomorphic Encryption. In *Topics in Cryptology – CT-RSA 2020 (Lecture Notes in Computer Science, Vol. 12006)*, Stanislaw Jarecki (Ed.). Springer, Heidelberg, Germany, San Francisco, CA, USA, 254–283. https://doi.org/10.1007/978-3-030-40186-3_12