

High-Throughput Universally Composable Threshold FHE Decryption

Guy Zyskind
Fhenix
Tel Aviv, Israel
guy@fhenix.io

Max Leibovich
Fhenix
Tel Aviv, Israel
max@fhenix.io

Doron Zarchy
Fhenix
Tel Aviv, Israel
doron@fhenix.io

Chris Peikert
Fhenix
University of Michigan
Ann Arbor, MI, USA
chris@fhenix.io, cpeikert@umich.edu

Abstract

Threshold Fully Homomorphic Encryption (FHE) enables arbitrary computation on encrypted data, while distributing the decryption capability across multiple parties. A primary application of interest is low-communication multi-party computation (MPC), which benefits from a fast and secure threshold FHE decryption protocol.

Several works have addressed this problem, but all existing solutions rely on “noise flooding” for security. This incurs significant overhead and necessitates large parameters in practice, making it unsuitable for many real-world deployments. Some constructions have somewhat better efficiency, but at the cost of weaker, non-simulation-based security definitions, which limits their usability and composability.

In this work, we propose a novel threshold FHE decryption protocol that avoids “noise flooding” altogether, and provides simulation-based security. Rather than masking the underlying ciphertext noise, our technique securely removes it via an efficient MPC rounding procedure. The cost of this MPC is mitigated by an offline/online design that preprocesses special gates for secure comparisons in the offline phase, and has low communication and computation in the online phase. This approach is of independent interest, and should also benefit other MPC protocols (e.g., secure machine learning) that make heavy use of non-linear comparison operations.

We prove our protocol secure in the Universal Composability (UC) framework, and it can be generally instantiated for a variety of adversary models (e.g., security-with-abort against a dishonest majority, or guaranteed output delivery with honest majority). Compared to the state of the art, our protocol offers significant gains both in the adversary model (i.e., dishonest vs. honest majority) and practical performance: empirically, our online phase obtains approximately 20,000× better throughput, and up to a 37× improvement in latency.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CCS '25, Taipei, Taiwan

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-1525-9/2025/10
<https://doi.org/10.1145/3719027.3744884>

CCS Concepts

• **Security and privacy** → **Cryptography**; *Distributed systems security*; *Security protocols*; **Key management**.

Keywords

Fully Homomorphic Encryption, Threshold FHE, Universal Composability, Secure Comparison

ACM Reference Format:

Guy Zyskind, Doron Zarchy, Max Leibovich, and Chris Peikert. 2025. High-Throughput Universally Composable Threshold FHE Decryption. In *Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security (CCS '25)*, October 13–17, 2025, Taipei, Taiwan. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3719027.3744884>

1 Introduction

Since the introduction of the first Fully Homomorphic Encryption (FHE) scheme [30] about 16 years ago, FHE has received intense sustained interest from both academia and industry. This is of no surprise, since FHE enables a wide array of powerful secure applications. In its simplest form, FHE enables a client to delegate arbitrary computations on its private data to an untrusted server, without revealing anything about the data. More specifically: the client first *encrypts* its data using an FHE scheme, yielding a ciphertext (which hides the data) that is sent to the server. The server then *homomorphically evaluates* any desired function on the ciphertext, yielding a new ciphertext that encrypts the value of the function on the data, which is returned to the client. Finally, the client uses its secret key to *decrypt* the derived ciphertext, thus recovering the function value.

Another very important application of FHE is to *secure multi-party computation* (MPC), which enables several parties to compute a function on their individual secret inputs, without revealing anything but the final result. The standard approach here relies on *threshold FHE*, which securely distributes the decryption key across all the parties (typically, using some form of secret sharing). Each party encrypts its own input, then the desired function is homomorphically evaluated on these ciphertexts, then a qualified subset of the parties jointly runs a *threshold decryption* protocol (using their shares of the secret key) to recover the final answer. This threshold decryption is the most delicate step, since it must

reveal the underlying plaintext and nothing more (about the secret key or the like).

Every well studied FHE scheme is based on some variant of the Learning With Errors (LWE) problem [44]. LWE is believed to resist quantum attacks, and is the foundation for countless other cryptographic constructions, from basic public-key encryption to countless powerful and versatile concepts (see [43] for a survey). In general, most of these constructions can benefit from being “thresholdized,” i.e., distributing their secret keys and decryption operations across several entities, to eliminate single points of failure.

Given the importance of threshold decryption for FHE and LWE-based encryption more broadly, many works have addressed this problem in recent years (e.g., [2, 4, 21, 35, 37], to name just a few). However, all of these protocols are either too inefficient for most applications [2, 4, 21, 37]; rely on new operational models (involving additional trusted parties) or computational assumptions (for practical efficiency) [35, 42, 51]; or settle for less than simulation security, instead resorting to weaker security definitions (e.g., [6, 19, 22]) that can be insufficient for the desired applications.

Noise flooding and its discontent. While all these works differ in some respects, every one that applies to FHE uses some form of *noise flooding*. The basic reason for this arises from the structure of “standalone” FHE/LWE decryption, which works in two steps:

- (1) First, we compute the inner product $\langle c, s \rangle$ of the ciphertext c and secret-key vector s . This yields a “noisy encoding” of the plaintext μ , of the form $\mu + e$ for some small enough error (or noise) term e .
- (2) Then, we “round away” the noise e to recover the plaintext μ itself.

In other words, the overall decryption function is the *rounded inner product* of the ciphertext and the secret key.

In the threshold setting, each decrypting party typically holds a share of the secret key, under some linear secret-sharing scheme. Thanks to the linearity, the parties can locally compute a sharing of the inner product $\langle s, c \rangle = \mu + e$, i.e., the noisy plaintext. At this point, the parties could simply reveal their shares to obtain the noisy plaintext, and round away the error. Alas, this would be *fatal* for security, since the noise (combined with some knowledge of the encryption randomness) leaks secret information that can even be used to recover the secret key. In particular, there is no hope of simulating this threshold decryption protocol using knowledge of the plaintext alone.

Noise flooding addresses this issue by having the decrypting parties add a large amount of extra noise before revealing their decryption shares. This statistically masks, or “floods,” the original noise e , enabling a simulation-based security proof. However, this comes at a high cost: it requires the original noise rate to be very small—asymptotically negligible—so to compensate, the other LWE parameters must be made much larger to ensure security. This results in significantly larger ciphertexts and slower homomorphic operations, compared to the standalone setting.

Indeed, the current state-of-the-art proposed by Dahl *et al.* [21], which is the only scheme having an adequate form of (simulation-based) security and reasonable performance in practice, walks this fine line by using *bootstrapping* to switch to small noise rate and

large LWE parameters *just prior to* threshold decryption (and not for regular FHE ciphertexts and operations). Yet this still imposes a high computational cost: a single decryption takes hundreds of milliseconds on a well-equipped 128-core machine. More importantly, being compute-bounded like this limits the *throughput* of the protocol, i.e., the number of possible decryptions per server per unit of time—regardless of the speed of the network.

Given the growing interest in practical threshold FHE systems (e.g., [38, 48, 50, 53]), there is a clear need for approaches that are more efficient and scalable. In particular, a *low-computation* and *low-communication* protocol would have much higher throughput than existing ones, even if it required slightly more rounds of interaction. In fact, if the latency of the extra rounds was less than the savings in computation time, such a protocol could even have better latency as well. This is the approach we pursue in this work.

1.1 Our Contribution

We propose a novel approach to threshold decryption for widely used FHE schemes, and LWE-based encryption more generally. Our protocol uses small LWE parameters—matching those for the standalone setting, in fact—and requires low computation and communication, at the cost of a slight increase in the rounds of interaction. As shown by our evaluation, it thereby achieves substantially higher throughput than the state of the art, and, in most realistic network settings, even lower latency as well. Crucially, we accomplish all this without compromising on security, obtaining strong simulation-based security in the universal composability (UC) framework [10] for a variety of adversary models.

In contrast to all prior work on threshold FHE, our protocol completely avoids the costly “noise flooding” technique—in fact, it adds no additional noise at all (thus preserving the LWE parameters of standalone schemes). Instead, we devise a novel *MPC-based approach* for removing the ciphertext error during decryption. High efficiency is obtained from an offline-online structure, where the offline phase can be executed at any time before the ciphertext to be decrypted is available, and the online phase takes a small number of rounds and uses very little computation and communication. We also highly optimize the offline phase and show empirically that it is efficient enough to support real-time usage in practice.

At a more technical level, some highlights of our results include the following.

- (1) Our protocol is designed in the abstract Arithmetic Black Box (ABB) model (see, e.g., [20, 26]), making it simple to express and analyze, and very generally instantiable for a variety of adversary models (e.g., dishonest majority, honest majority, semi-honest or robust, static or adaptive corruptions, etc.).
- (2) Empirically, our evaluation shows that our protocol is orders of magnitude more efficient than the state of the art [21]. In an instantiation of our protocol with the SPDZ_{2k} protocol for dishonest majority [33], we achieve up to 20,000× better throughput, and up to 37× lower latency for decrypting a single ciphertext.
- (3) As a main ingredient, we develop new efficient MPC protocols for comparison-type functions (less-than, mod, etc.) for moderate-size inputs (e.g., 64–256 bits). These protocols require very low communication and round complexity in

their online phase, at the cost of a moderate but very practical offline phase. This contribution should be of independent interest for many other MPC protocols and applications, such as secure Machine Learning (e.g., [36, 47, 52]).

1.2 Technical Overview

The basic approach behind our protocol is to entirely avoid noise flooding, and instead use MPC to *completely remove* the noise before anything is revealed to the parties. This sidesteps the complexities of managing noise, and dramatically simplifies the security analysis—instead of subtle arguments about noise distributions, we just prove that our MPC protocol is UC-secure. Therefore, the parties effectively have access only to an idealized decryption function, and in particular they learn nothing except the decrypted plaintext. We do not need to resort to novel security models or special game-based definitions.

Of course, securely removing noise requires secure comparisons, which typically involves bit-level operations that can be expensive if implemented naively in MPC [12, 23, 24, 40]. To avoid such a cost, we introduce a novel approach that trades a somewhat larger offline phase for a much cheaper online phase. More concretely, in the offline phase we build special preprocessed (i.e., input-independent) “gates” of moderate size, so that in the online phase the parties can securely perform the rounding step using just two “layers” of these gates. Our empirical evaluation shows that this protocol achieves much higher throughput and even significantly lower latency than prior noise-flooding approaches, delivering fast threshold FHE decryption in practice.

We next give a high-level technical overview of our threshold decryption protocol. For simplicity of presentation and generality of instantiation, it operates in the abstract Arithmetic Black-Box (ABB) model for the ring \mathbb{Z}_{2^k} . Essentially, the ABB provides an idealized trusted party that: stores secret values in \mathbb{Z}_{2^k} , including random bits that it generates; computes desired products and linear combinations of stored values; and reveals values upon request by the honest parties. We stress that powers of two are popular choices of moduli in FHE and LWE-based cryptography, but we can also trivially support any other choice as well, by cheaply “modulus-switching” to a power of two just before decryption.

Let $q = 2^k$ be the ciphertext modulus and $p = 2^m \ll q$ be the plaintext modulus, and let $L = q/p$. The secret key is a vector \mathbf{s} over $\mathbb{Z}_q = \mathbb{Z}_{2^k}$, which is stored in the ABB. When the honest parties wish to decrypt a ciphertext c , they proceed as follows:

- (1) **Inner product.** They instruct the ABB to internally compute the “noisy message” $z = \langle c, \mathbf{s} \rangle = L \cdot \mu + e$, where $\mu \in \mathbb{Z}_p$ is the message and $e \in [0, L)$ is the noise. This is done by taking a linear combination, which is purely local computation in a typical instantiation of the ABB.
- (2) **Error extraction.** They engage in a lightweight MPC protocol, in which the ABB computes and stores the error, as $e = z \bmod L$ (still stored as an element of \mathbb{Z}_{2^k}).
- (3) **Message recovery.** They instruct the ABB to cancel out the noise value by computing $z - e = L \cdot \mu$ (again using linearity), then have ABB open this value, which reveals μ itself.

The only nontrivial step is the secure computation (by ABB) of $z \bmod L$ from z . Doing so naively requires concretely too many

rounds and/or too much communication for our purposes [23]. Instead, we take a novel approach of computing the mod operation using just two layers of moderate-size preprocessed “gates,” for very specially designed functions; see Section 3.3 for details.

An interesting side note is that these gates are inspired by works on Function Secret Sharing (FSS), which allows to perform MPC with very little or even no interaction [7, 8, 55]—but usually at high computational and/or storage cost, and often requiring novel hardness assumptions. While precomputing our gates is somewhat (but not very) costly, it allows us to compute secure comparisons using a small constant number of rounds and low communication in the online phase. For example, for $k = 64$, which is a value of interest in this work and in many others (given the prevalence of 64-bit architectures), the online phase requires only three sequential openings from the ABB (two of which are randomly “masked,” and hence can be handled more efficiently). Finally, since the online phase has very low communication and computation, it allows us to scale up the number of parallel decryptions to obtain high throughput.

2 Preliminaries

2.1 Notation

For a positive integer n , define $[n] = \{0, 1, \dots, n-1\}$. (We caution that this notation almost overlaps with the one for stored values in \mathcal{F}_{ABB} , but the intended meaning should always be clear from context.) Let \mathbb{Z}_A be the quotient ring of integers modulo a positive integer A ; in this work, A is always a power of two.

2.2 Arithmetic Black Box (ABB) Model

In the MPC literature, the Arithmetic Black Box (ABB) model is a widely used abstraction for secure low-level arithmetic operations [20, 25, 28, 34]. The ABB model supports basic operations, including addition and multiplication, on elements of some specific ring or field; in this work, we exclusively use rings of the form \mathbb{Z}_{2^k} . The ABB greatly simplifies the design, presentation, and modularity of protocols that use it, by abstracting away the underlying cryptographic mechanisms like secret-sharing schemes and MACs.

The functionality \mathcal{F}_{ABB} (Figure 1) formalizes the Arithmetic Black Box model as an idealized trusted party in the UC framework, defining an abstraction for secure, reactive arithmetic computations. It allows parties to input, perform arithmetic operations on, and reveal secret values that are stored within the functionality. We present just \mathcal{F}_{ABB} ’s “core” definition, which can be paired with any standard “shell” corresponding to the adversary model, e.g., honest majority with fairness, dishonest majority with abort, adaptive security, etc.

Formally, we present our decryption protocol in the \mathcal{F}_{ABB} -hybrid model using the UC framework, and show that it perfectly realizes an abstract decryption functionality having the same shell as \mathcal{F}_{ABB} . This makes our protocol highly general and applicable, since it automatically inherits the specific efficiency and security features of any particular realization of \mathcal{F}_{ABB} with a given shell.

For convenience of usage, we have slightly modified the interface of \mathcal{F}_{ABB} as follows:

- We have elided the Input command, because our protocols do not use it.

Functionality \mathcal{F}_{ABB} for ring $R = \mathbb{Z}_{2^k}$

- **Random Bit Generation:** On input (RandBit, sid, id) from all honest parties (where id is fresh), send that tuple to the adversary, sample uniform $b \leftarrow \{0, 1\}$, and store (id, b).
- **Linear Combination:** On input (LinComb, sid, id, $(\text{id}_j)_j, c \in R, (c_j \in R)_j$) from all honest parties (where each id_j is stored in memory and id is fresh), send that tuple to the adversary, retrieve each (id_j, x_j) from memory, compute $y = c + \sum_j c_j \cdot x_j \in R$, and store (id, y).
- **Multiplication:** On input (Mult, sid, id, id_1, id_2) from all honest parties (where id_1, id_2 are stored in memory and id is fresh), send that tuple to the adversary, retrieve $(\text{id}_1, x), (\text{id}_2, y)$ from memory, compute $z = x \cdot y \in R$, and store (id, z).
- **Opening:** On input (Open, sid, id, l) from all honest parties, where $l \leq k$ (and id is present in memory), retrieve (id, y) from memory, let $y' = y \bmod 2^l$, and output (Open, sid, id, l, y') to the adversary and to all the parties.

Figure 1: Arithmetic black box functionality

- We have added the RandBit command, which generates and stores a uniformly random bit with some specified identifier. This has known implementations [23, 28, 46] that are compatible with existing realizations of \mathcal{F}_{ABB} for \mathbb{Z}_{2^k} .
- We have made the Open command a “truncated opening,” where only the *least-significant* l bits of the stored value are revealed, for some desired $l \leq k$. This enhancement can be implemented generically (in terms of the standard full opening) by using RandBit and linear operations to randomize the most-significant $k - l$ bits of the revealed value. Or, in some realizations of \mathcal{F}_{ABB} this can be implemented more directly and efficiently; e.g., with Shamir sharing over Galois rings [20, 27], revealing just the low l bits of each share opens just the low l bits of the shared value.

Notation. When defining protocols that use \mathcal{F}_{ABB} , it is cumbersome to use its formal interface, so for convenience we instead use some more natural notation.

We write $[x]_k$ to denote a value $x \in \mathbb{Z}_{2^k}$ that has been stored in the functionality \mathcal{F}_{ABB} ; this value is implicitly identified by some unique id associated with x .¹ The ids for ephemeral variables are *distinct* (not reused) across multiple calls to the same session of a protocol. Note that a public value can be stored in \mathcal{F}_{ABB} by calling LinComb with empty sets of identifiers id_j and coefficients c_j , and c equal to the public value.

¹Note that this matches the widespread notation for secret-shared values, and this is by intent: the reader may think of values stored in \mathcal{F}_{ABB} as being secret-shared among the parties, in a typical realization of \mathcal{F}_{ABB} . However, we do not use any secret-sharing scheme explicitly, but instead abstract it away with \mathcal{F}_{ABB} .

We overload the operators $+$ and \cdot instead of explicitly calling \mathcal{F}_{ABB} ’s LinComb and Mult commands. For example, $[z]_k = [x]_k \cdot [y]_k + c$ denotes that the values x, y stored in \mathcal{F}_{ABB} are multiplied using the Mult command, then the constant value c is added using LinComb, and the resulting value z is stored in \mathcal{F}_{ABB} ’s memory. This overloading of notation is commonly used in prior works, e.g., [34].

We also extend \mathcal{F}_{ABB} for the ring \mathbb{Z}_{2^k} to store and operate on values in \mathbb{Z}_{2^l} , for any $l < k$. We denote such stored values by $[x]_l$, later opening at most l of their bits, and implicitly “down-cast” stored values modulo smaller powers of two as needed. All this is well defined via the natural $(\bmod 2^l)$ ring homomorphism from \mathbb{Z}_{2^k} to \mathbb{Z}_{2^l} . In other words, we can generically obtain this enhancement by using, in place of $x \in \mathbb{Z}_{2^l}$ itself, any $\bar{x} \in \mathbb{Z}_{2^k}$ for which $\bar{x} \equiv x \pmod{2^l}$. Moreover, in typical realizations of \mathcal{F}_{ABB} like [20, 29], this can be implemented more directly and efficiently, simply by reducing each share by a correspondingly smaller modulus.

2.2.1 Realizing \mathcal{F}_{ABB} . Protocols that realize \mathcal{F}_{ABB} for various adversary models has been extensively studied in the literature (though more often in the prime-field case, whereas we work with power-of-two modulus).

- **Dishonest majority.** Most commonly, the setting of dishonest majority (with abort) is covered by SPDZ line of work (over \mathbb{Z}_{2^k}), e.g., [20, 23, 28, 29]. In this setting, additive secret-sharing is used, and all shares are authenticated via information-theoretic message authentication codes (MACs). This technique is very efficient, as it adds very little overhead compared to a realization for semi-honest adversaries.
- **Honest (super-)majority.** In these settings, it is common to use either replicated secret-sharing (for a small number of parties, due to the size of the shares) or Shamir secret sharing over Galois rings. These techniques provide both semi-honest and active security, and have been explored in several works [21, 29]. In the case of an honest super-majority (i.e., $t < n/3$ corrupt parties), using standard error-correction techniques on the parties’ revealed shares can ensure robustness and fairness; see, e.g., [29] for details.

Essentially all realizations of \mathcal{F}_{ABB} work via some linear secret-sharing scheme, which means that LinComb can be performed locally (with no communication between parties) using linear operations on shares. By contrast, Mult typically requires some communication, and is often implemented by preprocessing Beaver multiplication triples [3, 15, 26].

2.3 Preprocessed Gates

Our protocols obtain their online efficiency from the offline preparation and careful use of *preprocessed gates*. Let $F: A \rightarrow B$ be an arbitrary function, where $A = \mathbb{Z}_{2^a}$ and $B = \mathbb{Z}_{2^b}$. A preprocessed gate for F consists of:

- (1) a stored value $[r]_a$ in \mathcal{F}_{ABB} of some secret, uniformly random *masking term* $r \in A$, and
- (2) a *lookup table* $([y_{x'}]_b)_{x' \in A}$, stored in \mathcal{F}_{ABB} , of all the outputs $y_{x'} = F(x' - r)$ of the shifted function.

Conceptually, the gate can be seen as a kind of naïve function secret sharing of F .

To *apply* such a gate to a stored $[x]_a$, yielding the stored output $[F(x)]_b$, the parties simply:

- (1) Let $x' = \mathcal{F}_{\text{ABB}}.\text{Open}([x]_a + [r]_a)$. (Here $x' = x + r \in A$ is the “masked” input.)
- (2) Output $[y_{x'}]_b = [F(x' - r)]_b = [F(x)]_b$. (This just identifies the appropriate stored value in \mathcal{F}_{ABB} corresponding to x' .)

This uses one call to $\mathcal{F}_{\text{ABB}}.\text{Open}$; the remainder is one invocation of LinComb (recall that this is typically realized locally, with no communication) and a local computation of the stored output’s identifier. Observe that if $r \in A$ is uniformly random and independent of everything else, then this procedure is information-theoretically secure, because the only revealed value is x' , which perfectly hides x .

More generally, we also consider gates for *parts* of functions whose domains A are *infinite* additive groups, like $A = \mathbb{Z}$. In this case, the mask value r is from some suitable finite subset of the domain, and the lookup table is $([F(x' - r)]_b)_{x' \in X'}$ for some suitable finite $X' \subset A$. Applying such a gate to a stored $[x]_a$ is an ad-hoc process, because opening $x' = x + r \in A$ may reveal information about x , and also it might be that $x' \notin X'$. Instead, we will take care to open only partial values that fully mask the stored input, and to guarantee that the needed output is stored in the lookup table. Our Sign gates, as constructed in Section 4.2 and used in Section 3.3, are the primary example of this.

3 Distributed LWE Decryption via MPC Rounding

In this section we define an abstract ideal functionality for decrypting LWE ciphertexts—in particular, decryption in all widely used FHE schemes—and give a protocol that uses \mathcal{F}_{ABB} to perfectly realize this functionality. The functionality is defined in Section 3.1, the protocol and its main subroutine are defined in Sections 3.2 and 3.3, and suggested parameterizations are given in Section 3.4.

3.1 Functionality $\mathcal{F}_{\text{Decrypt}}$

Here we define and discuss our abstract ideal functionality $\mathcal{F}_{\text{Decrypt}}$ (Figure 2) for decrypting LWE ciphertexts. It is parameterized by a power-of-two ciphertext modulus $q = 2^k$ and plaintext modulus $p = 2^m$ for some positive integers $k > m$. These parameters are used globally throughout this work.

The functionality $\mathcal{F}_{\text{Decrypt}}$ has two commands: *Init*, which generates a public/secret key pair, and *Decrypt*, which may be called many times to decrypt given ciphertexts. The *Decrypt* command simply outputs the rounded (from \mathbb{Z}_q to \mathbb{Z}_p) inner product of the given ciphertext and the secret key.

Discussion. The functionality is parameterized by a *KeyGen* algorithm for the underlying LWE-based cryptosystem. The *Init* command simply runs this algorithm, stores the secret key s for later use, and outputs the public key to all parties. In our realization we assume a secure \mathcal{F}_{ABB} -aided procedure for *KeyGen*, which can be obtained generically by standard MPC techniques, or more efficiently using the scheme’s specific structure. For example, we can generate public LWE samples for a secret s by sampling secret random errors and using \mathcal{F}_{ABB} ’s linearity features.

We strongly emphasize that our $\mathcal{F}_{\text{Decrypt}}$ is not a complete functionality for threshold encryption; it is just a key ingredient in a

Functionality $\mathcal{F}_{\text{Decrypt}}$

- **Initialize:** Once $(\text{Init}, \text{sid})$ is received from each honest party, run $(pk, s) \leftarrow \text{KeyGen}$, store s , and send $(\text{Init}, \text{sid}, pk)$ to the adversary and all the parties. Do not act on any further *Init* commands for this session.
- **Decrypt:** On input $(\text{Decrypt}, \text{sid}, c)$ from all the honest parties, where c is an LWE ciphertext over \mathbb{Z}_q , send $(\text{Decrypt}, \text{sid}, c, \lfloor \langle c, s \rangle \rfloor_p)$ to the adversary and to all the parties, where $\lfloor x \rfloor_p := \lfloor \frac{p}{q} \cdot x \rfloor \in \mathbb{Z}_p$.

Figure 2: Functionality for decrypting LWE ciphertexts

protocol to realize such a functionality. In order to be meaningfully secure, this protocol would need to be designed so that honest parties *Decrypt* only ciphertexts that are known to be suitably “well formed.” Otherwise, the *Decrypt* command would act as an unrestricted decryption oracle, from which it is easy to learn the secret key using standard techniques. Specific restrictions on decrypted ciphertexts in the threshold setting have been considered in many prior works (e.g., [21, 35]), and are outside the scope of this work.

3.2 Protocol Π_{Decrypt}

Here we define our main decryption protocol (Figure 3), which perfectly realizes $\mathcal{F}_{\text{Decrypt}}$ in the \mathcal{F}_{ABB} -hybrid model. (See Section 2.2 for the details of \mathcal{F}_{ABB} .) Recalling that the ciphertext modulus is $q = 2^k$ and the plaintext modulus is $p = 2^m < q$, we let $l = k - m \geq 1$ and $L = 2^l = q/p$.

For *Decrypt* on a given ciphertext c , the parties use the stored secret key to linearly compute (inside \mathcal{F}_{ABB}) the “noisy decryption value” $[z]_k = [\mu \cdot L + e]_k$, where the message $\mu \in \mathbb{Z}_{2^m}$ and the noise $e \in [L]$ (for convenience, we shift the noise term to be non-negative). In other words, the message occupies the high m bits, and the noise occupies the low l bits. Then, using the $\text{Mod}_{L,k}$ subroutine (Procedure 1), the parties securely compute and cancel out the noise term, i.e., they compute $[e]_k = [z \bmod L]_k$ and $[z]_k - [e]_k = [\mu \cdot L]_k$. Finally, they *Open* the latter value to get μ .²

The main challenge, therefore, is to securely compute the noise term $[e]_k = [z \bmod L]_k$ from $[z]_k$, i.e., mod- L reduction on a k -bit secret input. We give a subroutine for this in Section 3.3 below.

The following is our main security theorem, which follows straightforwardly from the correctness of our subroutines and the fact that every opened intermediate value is *masked* by a fresh uniformly random value. The formal proof is given in Section A.

Theorem 1. *Protocol Π_{Decrypt} (Figure 3) perfectly realizes $\mathcal{F}_{\text{Decrypt}}$ in the \mathcal{F}_{ABB} -hybrid model, under the same (adaptive corruption) shell.*

Recall that the shell of a protocol captures the underlying adversarial model—e.g., static or more general adaptive corruptions, semi-honest or malicious adversarial behavior, security with abort or guaranteed output delivery—as well as the “access structure.” For the access structure, a prevalent choice is the threshold model,

²We remark that opening the original value z itself would reveal the noise in the ciphertext, which cannot be simulated from the decrypted message alone, and typically can lead to a complete break of the scheme.

Protocol Π_{Decrypt}

- On input (Init, sid), the parties run a secure \mathcal{F}_{ABB} -hybrid protocol for KeyGen that results in a stored secret key $[s]_k$ and a public key pk .
- On input (Decrypt, sid, c), the parties do the following:
 - (1) Use \mathcal{F}_{ABB} to compute $[z]_k = \langle c, [s]_k \rangle + 2^{l-1}$.
 - (2) Run $\text{Mod}_{l,k}(\text{sid}, [z]_k)$ to produce $[e]_k$.
 - (3) Call $\mathcal{F}_{\text{ABB}}.\text{Open}$ on $[z]_k - [e]_k$ and receive $\mu' \in \mathbb{Z}_q$ (which is a multiple of L).
 - (4) Output $\mu = \mu'/L \in \mathbb{Z}_p$.

Figure 3: Protocol for decrypting LWE ciphertext

wherein the functionality reveals all its (past and future) internal random choices to the adversary once a certain number of parties have been corrupted.

3.3 Secure Mod and Comparison via $\text{LTRand}_{l,k}$

Here we securely implement the $\text{Mod}_{l,k}$ subroutine and similar comparison functions, like (modular) less-than-zero ModLTZ_k . We implement these as “thin wrappers” around a new core abstraction and efficient procedure we call $\text{LTRand}_{l,k}$, which is one of our main technical contributions.

Essentially, $\text{LTRand}_{l,k}$ takes a stored input in \mathbb{Z}_{2^k} (or $[2^k]$), randomly masks it modulo 2^l , and returns the (opened) masked value, the stored mask, and a stored bit indicating whether the former is less than the latter. Its precise specification is as follows, and our efficient implementation is given below in Procedure 3.

Input: stored value $[z]_k$ for some $z \in \mathbb{Z}_{2^k}$.

Output: • $z' = z + r \bmod 2^l$ for fresh uniformly random $r \in [L]$,

- stored value $[r]_k$, and
- stored bit $[u]_k$, where $u = (z' \stackrel{?}{<} r) = (z' - r \stackrel{?}{<} 0) \in \{0, 1\}$.

Observe that the output reveals nothing, because $z' \in [2^l]$ is uniformly random for any input z .

Although the interface of $\text{LTRand}_{l,k}$ may seem somewhat ad-hoc, it is rich enough to directly implement several fundamental comparison procedures of wide applicability, even beyond our immediate purposes, e.g., secure machine learning [36, 47, 52, 54]. We next give two useful examples of this.

Secure mod and comparison. $\text{Mod}_{l,k}$ (Procedure 1) securely implements modular reduction, mapping stored input $[z]_k$ to stored output $[z \bmod 2^l]_k$. Importantly, the output is also a k -bit value in \mathbb{Z}_{2^k} , with zeros in its most significant $k - l$ bits.³ (Recall that in the decryption protocol Π_{Decrypt} , this is needed to extract and cancel out the ciphertext noise, leaving the message unaffected.)

As an optimization, with a typical realization of \mathcal{F}_{ABB} it suffices for $\text{LTRand}_{l,k}$ to produce $[u]_{k-l}$ instead of $[u]_k$, which can save significant storage in our typical setting where $k - l \ll k$. This

is because $2^l \cdot [u]_{k-l}$ can naturally be treated as $[2^l \cdot u]_k$ in the underlying secret-sharing schemes.

Procedure 1: $\text{Mod}_{l,k}(\text{sid}, [z]_k)$

- (1) Run $\text{LTRand}_{l,k}(\text{sid}, [z]_k)$ to get $z', [r]_k, [u]_k$.
- (2) Output $[e]_k = z' - [r]_k + 2^l \cdot [u]_k$.

Lemma 2. *Procedure $\text{Mod}_{l,k}$ is correct.*

PROOF. The output $e = z' - r + 2^l \cdot u$ satisfies $e \equiv z' - r \equiv z \pmod{2^l}$ and $e \in [2^l]$, because $z' - r \in (-2^l, 2^l)$ (since $z', r \in [2^l]$) and $u = (z' - r \stackrel{?}{<} 0)$. So, $e = z \bmod 2^l$, as claimed. \square

As another example, for any integer $t \geq 1$ defining $T = 2^t$, define the “modular less-than-zero” function $\text{ModLTZ}_t: \mathbb{Z}_T \rightarrow \{0, 1\}$ as

$$\text{ModLTZ}_t(x) := \begin{cases} 0 & \text{if } x \in [0, T/2) \pmod{T} \\ 1 & \text{if } x \in [-T/2, 0) \equiv [T/2, T) \pmod{T}. \end{cases} \quad (1)$$

In other words, this corresponds to whether the input’s *signed* representative in $[-T/2, T/2)$ is less than zero; equivalently, it is the most-significant bit of the input’s *unsigned* representative in $[0, T)$. Procedure 2 gives a secure procedure for this function; it is closely inspired by the MSB procedure in [23]. Observe that it reveals nothing because the only opened value is $T/2$ times the bit $h' = h \oplus b$, which is uniformly random for any input z .

Procedure 2: $\text{ModLTZ}_t(\text{sid}, [z]_t)$

- (1) Run $\text{Mod}_{t-1,t}(\text{sid}, [z]_t)$ to get $[e]_t = [z \bmod T/2]_t$.
- (2) Call $\mathcal{F}_{\text{ABB}}.\text{RandBit}$ to get $[b]_t$.
- (3) Call $\mathcal{F}_{\text{ABB}}.\text{Open}$ on $[H']_t = [e]_t - [e]_t + (T/2) \cdot [b]_t$, yielding $H' = (T/2) \cdot h' \in \mathbb{Z}_T$ for some $h' \in \{0, 1\}$.
- (4) Output $h' + [b]_t - 2h' \cdot [b]_t$.

Lemma 3. *Procedure ModLTZ_t is correct.*

PROOF. By Lemma 2, $e = z \bmod T/2$, so $z - e \equiv (T/2) \cdot h \pmod{T}$, where $h = \text{ModLTZ}_t(z) \in \{0, 1\}$. Therefore, the opened bit is $h' = h \oplus b \in \{0, 1\}$. Finally, $h + b - 2h'b = h' \oplus b = h = \text{ModLTZ}_t(z)$, as desired. \square

Implementation of $\text{LTRand}_{l,k}$. The $\text{LTRand}_{l,k}$ procedure has an offline-online structure with a very efficient online phase, thanks to its usage of rich preprocessed gates, as prepared in the offline phase. It is parameterized by an input bit length b for (most of) these gates, which determines the following additional parameters:

- For $L = 2^l$, values in $[L]$ are represented in base $B = 2^b$, i.e., using “digits” of bit length b .
- The number of digits is therefore $d = \lceil l/b \rceil$, where we let $D = 2^d$, and the bit length of the most-significant digit is $b' = l - (d - 1)b \leq b$, where we let $B' = 2^{b'}$.

In the offline phase, $\text{LTRand}_{l,k}$ prepares (using the procedures given in Section 4) several Sign gates for b -bit inputs, and one ModLTZ gate for a $(d + 1)$ -bit input. Each of these is used just once in a later call to the online phase.

³Recall from Section 2.2 that reducing from modulus 2^k to 2^l for $l < k$ is trivially supported by \mathcal{F}_{ABB} , but the result has just l bits of precision (it is in \mathbb{Z}_{2^l}), whereas here we need to retain k bits of precision, with zeros in the most-significant bits.

The online phase first opens $z' = z + r \bmod L$, where $r \in [L]$ is a uniformly random mask constructed from the masks r_i of the individual Sign gates—specifically, the r_i are the base- B digits of r . (Observe that r perfectly masks the low l bits of z , so z' reveals nothing.) It then uses the preprocessed gates to securely compute (inside \mathcal{F}_{ABB}) the bit

$$u = (z' \stackrel{?}{<} r) \in \{0, 1\}.$$

Finally, it outputs z' and the stored values r, u .

The main challenge lies with securely computing the bit u . The remainder of this subsection describes how we do this using the preprocessed Sign and ModLTZ gates.

Define the function $\text{Sign}: \mathbb{Z} \rightarrow \{-1, 0, 1\}$ as

$$\text{Sign}(x) := \begin{cases} -1 & \text{if } x < 0 \\ 0 & \text{if } x = 0 \\ +1 & \text{if } x > 0. \end{cases} \quad (2)$$

Our approach is based on the following central recurrence for the Sign function.

Lemma 4. *Let $x, r \in [B^d] = [2^{bd}]$ have base- B representations $x = \sum_{i \in [d]} x_i \cdot B^i$ with each $x_i \in [B]$, and similarly for r . Then the sign of $x - r$ is the “sign of weighted signs” of the digit differences:*

$$\text{Sign}(x - r) = \text{Sign}\left(\sum_{i \in [d]} (x_i - r_i) \cdot B^i\right) = \text{Sign}\left(\sum_{i \in [d]} \text{Sign}(x_i - r_i) \cdot 2^i\right). \quad (3)$$

PROOF. If $x_{d-1} - r_{d-1} \neq 0$, then $\text{Sign}(x - r) = \text{Sign}(x_{d-1} - r_{d-1})$, since

$$|x_{d-1} - r_{d-1}| \cdot B^{d-1} > \left| \sum_{i \in [d-1]} (x_i - r_i) \cdot B^i \right|$$

by the triangle inequality, the fact that each $|x_i - r_i| \leq B - 1$, and geometric sums. For similar reasons, $\text{Sign}(x_{d-1} - r_{d-1})$ equals the right-hand side of Equation (3), because each $|\text{Sign}(x_i - r_i)| \leq 1$. Otherwise, $x_{d-1} - r_{d-1} = 0$, so its contribution to each side of Equation (3) is zero, and the claim holds by induction. \square

Similarly to Lemma 4, for any $x, r \in [B^d]$ with their d -digit base- B representations,

$$(x \stackrel{?}{<} r) = (x - r \stackrel{?}{<} 0) = \text{ModLTZ}_{d+1}\left(\sum_{i \in [d]} \text{Sign}(x_i - r_i) \cdot 2^i \bmod 2D\right). \quad (4)$$

We emphasize that the weighted sum is in the interval $(-D, D)$, so the reduction modulo $2D$ does not “wrap around,” hence the equation holds true. We need this form of the equation because in our procedure, the weighted sum is computed by \mathcal{F}_{ABB} , which is limited to modular arithmetic with a certain amount of precision.

Altogether, Equation (4) gives a very efficient way to securely compute the bit $u = (z' \stackrel{?}{<} r)$ using the preprocessed Sign gates and ModLTZ gate. Specifically, using the digits z'_i of z' , we first look up the stored digit-difference signs $[\text{Sign}(z'_i - r_i)]$ from the Sign gates, then compute their weighted sum, then apply the ModLTZ gate to the result. Finally, we remark that this approach generalizes to more “layers” of signs of weighted signs, by applying Lemma 4 inductively. This yields a tradeoff between the preprocessing computation and

storage, and number of sequential calls to \mathcal{F}_{ABB} .Open, and can be helpful for handling larger input moduli (i.e., values of k).

Procedure 3: LTRand _{l,k} (sid, $[z]_k$)

Preprocessing phase (before $[z]_k$ is known):

- (1) Call $\text{PrepModLTZ}_{d+1,m}(\text{sid})$ (Procedure 7) to prepare a ModLTZ gate.
- (2) For each $i \in [d-1]$, call $\text{PrepSign}_{b,d+1}(\text{sid})$ (Procedure 5) to prepare a Sign gate

$$[r_i]_k, ([\text{Sign}(x_i - r_i)]_{d+1})_{x_i \in [B]}.$$

For $i = d-1$, do the same but with b', B' in place of b, B (respectively).

- (3) Compute $[r]_k = \sum_{i \in [d]} [r_i]_k \cdot B^i$.
 $\triangleright r = (r_{d-1} \cdots r_1 r_0)_B$ is the base- B representation of a uniformly random $r \in [L]$.

Online phase (once $[z]_k$ is known):

- (1) Call \mathcal{F}_{ABB} .Open on $[z]_l + [r]_l$ and receive $z' \in [L]$.
 $\triangleright z' = z + r \bmod L$
 - (2) Express z' in base B , as $z' = \sum_{i \in [d]} z'_i \cdot B^i$ where $z'_i \in [B]$, and $z'_{d-1} \in [B']$.
 - (3) For each $i \in [d]$, let $[y_i]_{d+1} = [\text{Sign}(z'_i - r_i)]_{d+1}$ from the corresponding Sign gate.
 - (4) Compute $[y]_{d+1} = \sum_{i \in [d]} [y_i]_{d+1} \cdot 2^i$.
 - (5) Apply (as defined in Section 2.3) the ModLTZ gate to $[y]_{d+1}$, yielding $[u]_k$.
 $\triangleright u = (z' \stackrel{?}{<} r)$, by Equation (4)
 - (6) Output $z', [r]_k, [u]_k$.
-

Lemma 5. *Procedure LTRand _{l,k} is correct.*

PROOF. Adopting the procedure’s notation, we analyze the values that are computed and opened by \mathcal{F}_{ABB} . First, on Line (1), \mathcal{F}_{ABB} opens $z' = z + r \bmod L$, where $r \in [L]$ is uniformly random. Then, by the correctness of the Sign gates, Lines (3) and (4) compute $y = \sum_{i \in [d]} \text{Sign}(z'_i - r_i) \cdot 2^i \bmod 2D$. Then, by Equation (4) and the correctness of the ModLTZ gate, Line (5) computes $u = (z' \stackrel{?}{<} r)$. \square

3.4 Efficiency and Parameters

Communication and rounds. The online phase of the protocol Π_{Decrypt} has very low communication and computation, and few rounds. Specifically, each invocation of Decrypt has three sequential calls to \mathcal{F}_{ABB} .Open:

- (1) at the start of LTRand, to reveal a masked l -bit value;
- (2) when applying the ModLTZ gate in LTRand, to reveal a masked $(d+1)$ -bit value; and
- (3) when opening the decrypted message, to reveal an unmasked k -bit value.

So, the total number of opened bits is $l + (d+1) + k$. The rest of the online phase is either local computation or calls to \mathcal{F}_{ABB} .LinComb, which in a typical realization are also just local computation. There are no calls to \mathcal{F}_{ABB} .Mult in the online phase.

Gate storage and offline efficiency. Each call to Decrypt consumes, and hence requires producing, several preprocessed gates in \mathcal{F}_{ABB} . Specifically, it uses (ignoring the masking values, which are insignificant):

- One ModLTZ gate, represented by a table of 2^{d+1} stored m -bit values (using the optimization mentioned in Section 3.3).
- $d = \lceil l/b \rceil$ Sign gates, each represented by a table of $B = 2^b$ (or $B' = 2^{b'}$, for the most-significant digit) stored $(d+1)$ -bit values.

Therefore, the total bit length of the stored gates is

$$(d+1) \cdot ((d-1) \cdot 2^b + 2^{b'}) + m \cdot 2^{d+1}. \quad (5)$$

This is approximately optimized for $b \approx d \approx \sqrt{l} = \sqrt{k-m}$ (because the 2^b and 2^{d+1} factors dominate their respective terms), though finer-grained optimization tends to produce somewhat larger d and somewhat smaller b .

The offline cost is roughly proportional to the total number of entries in the tables, because the preprocessing makes roughly that many calls to $\mathcal{F}_{\text{ABB}}.\text{Mult}$ (plus local computation); see Section 4. The total number of table entries is roughly $d \cdot 2^b + 2^{d+1}$, which again is optimized for $b \approx d \approx \sqrt{l}$.

Realization with SPDZ_{2k}. We give some concrete communication and storage numbers when \mathcal{F}_{ABB} is instantiated using the SPDZ_{2k} protocol [20], for the realistic example parameters

$$k = 64, m = 1, b = 8 \text{ and hence } d = 8, b' = 7.$$

We use $s = 64$ for the MAC's statistical security parameter, which is also the statistical security level of each call to $\mathcal{F}_{\text{ABB}}.\text{Open}$.

Naïvely, using the SPDZ_{2k} protocol exactly as written, the three sequential calls to $\mathcal{F}_{\text{ABB}}.\text{Open}$ each take three broadcast rounds, for a total of nine rounds. More specifically, each call to $\mathcal{F}_{\text{ABB}}.\text{Open}$ runs the SPDZ_{2k} SingleCheck procedure⁴, which first broadcasts shares of the value, then commitments to shifted MAC shares, then those shares themselves, and checks their validity (i.e., that the committed shares were opened correctly and sum to zero).

Fortunately, due to the specific structure of our protocol, we believe it suffices to use just *five* broadcast rounds. This is because the first two calls to $\mathcal{F}_{\text{ABB}}.\text{Open}$ are for *randomly masked* values, and a similar kind of (committed) masking can be done for the final output. Because these values are masked, they can be revealed before checking the authenticity (MACs) of the prior values they depend on, without violating privacy. This allows us to “pipeline” the rounds associated with the runs of SingleCheck. Specifically:

- the first round commits to shares of a random mask w for the final output, and reveals the shares of the first (masked) value;
- the second round commits to the shifted MAC shares for the first value, and reveals the shares of the second (masked) value;
- the third round announces and checks those shifted MAC shares (aborting if they are invalid), commits to the shifted MAC shares of the second value, and reveals the shares of the final output *masked by w* ; etc.

⁴There are no calls to $\mathcal{F}_{\text{ABB}}.\text{Mult}$ during the online phase, so the BatchCheck procedure is not needed there.

In the fifth round, if the protocol has not aborted, the committed shares of the mask w are also announced and checked; if they are valid, w is reconstructed and used to compute the unmasked final output. Informally, this pipelined approach should be secure because the only non-random shared value is not revealed until all of the values that contribute to it have been authenticated. Formalizing this seems to require a refined definition of \mathcal{F}_{ABB} and a ground-up proof for the SPDZ_{2k} protocol, which we leave to future work.

The communication per party for opening a t -bit value is as follows. First, each party reveals its $(t+s)$ -bit share, then it commits to its $(t+s)$ -bit (shifted) MAC share, then it reveals that share. So, the total communication per party for opening a t -bit value is just $2(t+s)$ bits, plus the commitments. Our protocol opens values of $t = l, d+1, k$ bits, for a total of just $2(l + (d+1) + k + 3s) = 656$ bits, plus the commitments.

For the gate storage per invocation of Decrypt (again, ignoring the random mask values), by Equation (5), \mathcal{F}_{ABB} stores and consumes a total of 17 792 bits. In the SPDZ_{2k} realization of \mathcal{F}_{ABB} , each stored t -bit value has a share size and MAC share size of $t+s$ bits each. So, adjusting Equation (5), each party stores a total of $2 \cdot ((d+1+s) \cdot ((d-1) \cdot 2^b + 2^{b'}) + (m+s) \cdot 2^{d+1}) = 346\,880$ bits (for our example parameters).

4 Preprocessing Procedures

Here we show how to efficiently prepare gates—i.e., random mask values r and corresponding lookup tables—for the Sign and ModLTZ functions. A key commonality is that, as we show, the lookup tables for both functions can be seen as “structured” linear transforms of the *subset-products* of the bits of r . These structured transforms can be evaluated by fast algorithms, roughly analogous to the Fast Fourier Transform. In Section 4.1 we first give a simple protocol for computing a random mask value along with the subset-products of its bits. Then, in Sections 4.2 and 4.3 we derive the fast linear transforms that map from the subset-product vectors to the lookup tables.

4.1 Procedure PrepSubsetProds

Let $a \geq 1$ be an input length (in bits) with $A = 2^a$, and $c \leq k$ be an output length. Here we give a procedure, in the \mathcal{F}_{ABB} -hybrid model, that prepares a uniformly random $r \in [A]$ and all the subset-products of its bits, as integers modulo 2^c .

In what follows, it is convenient to index vectors by subsets. For this purpose, the subset $S \subseteq [a]$ corresponds to the index

$$\sum_{s \in S} 2^s = \sum_{i \in [a]} \delta_{S,i} \cdot 2^i \in [A],$$

where $\delta_{S,i}$ is 1 if $i \in S$, and 0 otherwise.

Procedure 4: PrepSubsetProds_{a,c}(sid)

Output: $[r]_k$ for a uniformly random $r = (r_{a-1} \cdots r_0)_2 \in [A]$, along with the vector $([p_S]_c)_{S \subseteq [a]}$ of all the subset-products $p_S = \prod_{i \in S} r_i \in \{0, 1\}$ of the bits r_i of r .

- (1) If $a = 1$, call \mathcal{F}_{ABB} on (RandBit, sid) to generate $[r]_k$, and output $[r]_k, ([1]_c, [r]_c)$.
- (2) Write $a = a_0 + a_1$ for some $a_0, a_1 \geq 1$ (typically, $a_0 = \lceil a/2 \rceil$). In parallel, recursively call PrepSubsetProds_{a₀,c}(sid) and PrepSubsetProds_{a₁,c}(sid) to prepare (respectively)

$$[r_0]_k, ([p_{S_0}]_c)_{S_0 \subseteq [a_0]} \quad \text{and} \quad [r_1]_k, ([p'_{S_1}]_c)_{S_1 \subseteq [a_1]}.$$

- (3) Let $[r]_k = [r_0]_k + 2^{a_0} \cdot [r_1]_k$.
 - (4) For each $S_0 \subseteq [a_0]$ and nonempty $S_1 \subseteq [a_1]$, use \mathcal{F}_{ABB} to compute $[p_{S_0 \cup (a_0 + S_1)}]_c = [p_{S_0}]_c \cdot [p'_{S_1}]_c$.
 - (5) Output $[r]_k, ([p_S]_c)_{S \subseteq [a]}$.
-

Lemma 6. Procedure PrepSubsetProds_{a,c} is correct.

PROOF. The base case is correct by inspection. For the recursive case, first note that $r = r_0 + 2^{a_0} \cdot r_1 \in [A] = [2^a]$ is uniformly random because $r_0 \in [2^{a_0}]$, $r_1 \in [2^{a_1}]$ are uniform and independent, by induction. For correctness of the subset-products, observe that every subset $S \subseteq [a]$ is the disjoint union of some $S_0 \subseteq [a_0]$ and the shifted subset $a_0 + S_1$ for some $S_1 \subseteq [a_1]$. By induction, p_{S_0}, p'_{S_1} are the subset-products of the bits of r_0, r_1 indexed by S_0, S_1 (respectively), so $p_{S_0} \cdot p'_{S_1}$ is the subset-product of the bits of $r = r_0 + 2^{a_0} \cdot r_1$ indexed by $S = S_0 \cup (a_0 + S_1)$. \square

Efficiency analysis. The procedure can be implemented using a total of $2^a - a - 1$ calls to \mathcal{F}_{ABB} .Mult. For the base case this is true by inspection. For the recursive case, a call to Mult is needed only for nonempty S_0, S_1 , because the empty set has subset-product 1. So, the number of Mult calls satisfies the recurrence $M(a) = M(a_0) + M(a_1) + (2^{a_0} - 1)(2^{a_1} - 1)$, which solves to $M(a) = 2^a - a - 1$ by induction.

The number of sequential “rounds” of (parallel) Mult calls can be as small as $\lceil \log_2 a \rceil$, by always taking $a_0 = \lceil a/2 \rceil$.

4.2 Procedure PrepSign

The procedure PrepSign_{b,d+1} (Procedure 5) prepares a Sign gate for b -bit inputs, i.e., the domain $[B]$ where $B = 2^b$, with output modulo $2D = 2^{d+1}$. It does this by first preparing the subset-products of the bits of a (secret) uniformly random $r \in [B]$, then converts those subset-products to the vector of $\text{Sign}(x - r)$ values for all $x \in [B]$. The conversion is a *linear* function with a *fast* implementation via a recursive divide-and-conquer procedure, which we give below in Procedure 6.

Procedure 5: PrepSign_{b,d+1}(sid)

- (1) Call PrepSubsetProds_{b,d+1}(sid) to prepare $[r]_k, [\vec{p}]_{d+1}$
 - (2) Output $[r]_k, \text{SubsetProdsToSigns}_b([\vec{p}]_{d+1}) \triangleright$ Procedure 6
-

Fast linear transform from subset-products to signs. Here we describe the fast linear transform that converts the subset-products of

the bits of r to the values of $\text{Sign}(x - r)$, for all $x \in [B]$. The key idea is that the output vector satisfies a simple recurrence relation in terms of the top bits of x and r (respectively), and the output vector for their remaining $b - 1$ bits each; see Equation (6) below. So, there is a fast algorithm that recursively evaluates the transform, analogous to the Fast Fourier Transform.

Because the transform is linear, the algorithm can be implemented straightforwardly on values stored in \mathcal{F}_{ABB} using just its LinComb operation, and hence using just local computation (no communication) in a typical realization of \mathcal{F}_{ABB} . Therefore, for simplicity of presentation we omit the $[\cdot]$ notation around all the values, and present the algorithm as operating on the values themselves.

View $x, r \in [B]$ in binary as $x = (x_{b-1} \cdots x_0)_2$ and $r = (r_{b-1} \cdots r_0)_2$. When $b = 1$, clearly $\text{Sign}(x - r) = x_0 - r_0$. When $b > 1$, the key insight is that for the “truncated” values $x' = (x_{b-2} \cdots x_0)_2$ and $r' = (r_{b-2} \cdots r_0)_2$, we have the recurrence relation

$$\text{Sign}(x - r) = (x_{b-1} - r_{b-1}) + \text{Sign}(x' - r') \cdot \begin{cases} 1 - r_{b-1} & \text{if } x_{b-1} = 0 \\ r_{b-1} & \text{if } x_{b-1} = 1. \end{cases} \quad (6)$$

This can be seen by observing that if $x_{b-1} \neq r_{b-1}$, $\text{Sign}(x - r) = x_{b-1} - r_{b-1}$ (the less-significant bits are irrelevant); otherwise, $x_{b-1} - r_{b-1} = 0$ and hence $\text{Sign}(x - r) = \text{Sign}(x' - r')$. In the former case, $\text{Sign}(x' - r')$ is multiplied by 0 in the above expression, and in the latter case it is multiplied by 1.

Procedure 6: Linear Function SubsetProdsToSigns_b(\vec{p})

Input: $\vec{p} = (p_S)_{S \subseteq [b]}$, where $p_S = \prod_{i \in S} r_i$ for some $r = (r_{b-1} \cdots r_0)_2 \in [B]$

Output: vector $(\text{Sign}(x - r))_{x \in [B]}$

- (1) If $b = 0$ output (0).
- (2) Split \vec{p} into its initial half $\vec{p}^{(0)} = (p_{S'})_{S' \subseteq [b-1]}$ and latter half $\vec{p}^{(1)} = r_{b-1} \cdot \vec{p}^{(0)}$.
- (3) Apply SubsetProdsToSigns_{b-1} to each half, yielding

$$\vec{s}^{(0)} = (\text{Sign}(x' - r'))_{x' \in [B/2]}, \quad \vec{s}^{(1)} = r_{b-1} \cdot \vec{s}^{(0)}$$

respectively, where $r' = (r_{b-2} \cdots r_0)_2$.

- (4) Output the vector, indexed by $[B]$, whose initial half is

$$\begin{aligned} & -p_0^{(1)} \cdot \vec{1}_{B/2} + (\vec{s}^{(0)} - \vec{s}^{(1)}) \\ & = -r_{b-1} \cdot \vec{1}_{B/2} + \vec{s}^{(0)} \cdot (1 - r_{b-1}) \end{aligned}$$

and whose latter half is

$$\begin{aligned} & (p_0^{(0)} - p_0^{(1)}) \cdot \vec{1}_{B/2} + \vec{s}^{(1)} \\ & = (1 - r_{b-1}) \cdot \vec{1}_{B/2} + \vec{s}^{(0)} \cdot r_{b-1}, \end{aligned}$$

where $\vec{1}_{B/2}$ is the all-ones vector of dimension $B/2$.

The algorithm is correct by its correct recursive computation of $\vec{s}^{(0)}$ and $\vec{s}^{(1)}$, which follows from its linearity and the relationship between $\vec{p}^{(0)}$ and $\vec{p}^{(1)}$, and by Equation (6). The running time is $O(B \log B) = O(Bb)$ additions and subtractions, due to the recursive divide-and-conquer nature of the algorithm.

4.3 Procedure PrepModLTZ

For $d + 1 \leq k$, the procedure $\text{PrepModLTZ}_{d+1,m}$ (Procedure 7) prepares a ModLTZ gate for $(d + 1)$ -bit inputs, i.e., the domain \mathbb{Z}_{2D} where $D = 2^d$, with output modulo 2^m . It does this by first preparing the subset-products of the bits of a (secret) uniformly random $r \in \mathbb{Z}_{2D}$, then converts those subset-products to the vector of $\text{ModLTZ}_{d+1}(x - r)$ values for all $x \in \mathbb{Z}_{2D}$. As with PrepSign , the conversion is a *linear* function with a fast divide-and-conquer implementation, which we give below in Procedure 8.

Procedure 7: $\text{PrepModLTZ}_{d+1,m}(\text{sid})$

- (1) Call $\text{PrepSubsetProds}_{d+1,m}(\text{sid})$ to prepare $[r]_{d+1}, [\vec{p}]_m$
 - (2) Output $[r]_{d+1}, \text{SubsetProdsToModLTZ}_{d+1}([\vec{p}]_m)$
- ▷ Procedure 8
-

Fast linear transform from subset-products. Here we derive a fast linear transform that converts the subset-products of the bits of $r \in \mathbb{Z}_{2D}$ to the values of $\text{ModLTZ}_{d+1}(x - r)$ for all $x \in \mathbb{Z}_{2D}$. As above, we give a recurrence relation in terms of the top bits of x and r (respectively) and their remainders; see Equations (7) and (8) below. This in turn yields a fast recursive algorithm that evaluates the transform. As above, because the transform is linear, we express it as operating on vectors of values themselves, omitting the $[\cdot]$ notation.

First, for any $a \geq 1$ and integers $x, \bar{r} \in [2^a]$, define

$$\text{Carry}_a(x, \bar{r}) := (x + \bar{r} + 1 \geq 2^a)$$

to be the “carry” (or “overflow”) bit of $x + \bar{r} + 1$. For convenience, we also define the trivial base case $\text{Carry}_0(\varepsilon, \varepsilon) = 1$. Then by inspection, this function satisfies the recurrence

$$\text{Carry}_a(x, \bar{r}) = \begin{cases} \bar{r}_{a-1} \cdot \text{Carry}_{a-1}(x', \bar{r}') & \text{if } x_{a-1} = 0 \\ \bar{r}_{a-1} + (1 - \bar{r}_{a-1}) \cdot \text{Carry}_{a-1}(x', \bar{r}') & \text{if } x_{a-1} = 1, \end{cases} \quad (7)$$

where $x = x_{a-1} \cdot 2^{a-1} + x'$ for its top bit $x_{a-1} \in \{0, 1\}$ and remainder $x' \in [2^{a-1}]$, and similarly for $\bar{r}, \bar{r}_{a-1}, \bar{r}'$. This recurrence directly yields the function $\text{SubsetProdsToCarries}$ in Procedure 9 below.

Now we turn to $\text{ModLTZ}_{d+1}(x - r)$. Identify $x, r \in \mathbb{Z}_{2D}$ with their $(d + 1)$ -bit representatives in $[2D]$, and recall from Equation (1) that $\text{ModLTZ}_{d+1}(x - r)$ outputs whether $x - r \in [-D, 0) \equiv [D, 2D) \pmod{2D}$. This is equivalent to the d th bit (counting from 0) of $x + \bar{r} + 1$, where $\bar{r} = (2D - 1) - r \in [2D]$ is the integer represented by the bitwise complement of r . So by inspection,

$$\text{ModLTZ}_{d+1}(x - r) = \begin{cases} \bar{r}_d + (1 - 2\bar{r}_d) \cdot \text{Carry}_d(x', \bar{r}') & \text{if } x_d = 0 \\ (1 - \bar{r}_d) + (2\bar{r}_d - 1) \cdot \text{Carry}_d(x', \bar{r}') & \text{if } x_d = 1, \end{cases} \quad (8)$$

where $x = x_d \cdot 2^d + x'$ for its top bit $x_d \in \{0, 1\}$ and remainder $x' \in [2^d]$, and similarly for $\bar{r}, \bar{r}_d, \bar{r}'$. This equation directly yields the function $\text{SubsetProdsToModLTZ}$ in Procedure 8 below. Similarly to the above, its running time is $O(D \log D) = O(Dd)$ additions and subtractions.

Procedure 8: Linear Function $\text{SubsetProdsToModLTZ}_{d+1}(\vec{p})$

Input: $\vec{p} = (p_S)_{S \subseteq [d+1]}$, where $p_S = \prod_{i \in S} r_i$ for some $r = (r_d \cdots r_0)_2 \in \mathbb{Z}_{2D}$ with $D = 2^d$.

Output: vector $(\text{ModLTZ}_{d+1}(x - r))_{x \in \mathbb{Z}_{2D}}$

- (1) Split the input vector into its initial half $\vec{p}^{(0)} = (p_{S'})_{S' \subseteq [d]}$ and latter half $\vec{p}^{(1)} = r_d \cdot \vec{p}^{(0)}$.
- (2) Apply $\text{SubsetProdsToCarries}_d$ (Procedure 9) to each half, yielding

$$\vec{c}^{(0)} = (\text{Carry}(x', \bar{r}'))_{x' \in [D]}, \quad \vec{c}^{(1)} = r_d \cdot \vec{c}^{(0)}$$

respectively, where $\bar{r}' = (D - 1) - (r_{d-1} \cdots r_0)_2 \in [D]$.

- (3) Output the vector, indexed by \mathbb{Z}_{2D} , whose initial half (indexed by $[D]$) is

$$\begin{aligned} & (p_0^{(0)} - p_0^{(1)}) \cdot \vec{1}_D - \vec{c}^{(0)} + 2\vec{c}^{(1)} \\ & = (1 - r_d) \cdot \vec{1}_D + (2r_d - 1) \cdot \vec{c}^{(0)}, \end{aligned}$$

and whose latter half is

$$p_0^{(1)} \cdot \vec{1}_D + \vec{c}^{(0)} - 2\vec{c}^{(1)} = r_d \cdot \vec{1}_D + (1 - 2r_d) \cdot \vec{c}^{(0)}.$$

Procedure 9: Linear Function $\text{SubsetProdsToCarries}_d(\vec{p})$

Input: $\vec{p} = (p_S)_{S \subseteq [d]}$, where $p_S = \prod_{i \in S} r_i$ for some $r = (r_{d-1} \cdots r_0)_2 \in [D]$ with $D = 2^d$.

Output: vector $(\text{Carry}_d(x, \bar{r}))_{x \in [D]}$, where $\bar{r} = (D - 1) - r$.

- (1) If $d = 0$, output $\vec{p} = (p_0) = (1)$.
- (2) Split the input vector into its initial half $\vec{p}^{(0)} = (p_{S'})_{S' \subseteq [d-1]}$ and latter half $\vec{p}^{(1)} = r_{d-1} \cdot \vec{p}^{(0)}$.
- (3) Apply $\text{SubsetProdsToCarries}_{d-1}$ to each half, yielding

$$\vec{c}^{(0)} = ([\text{Carry}(x', \bar{r}')]_{x' \in [D/2]}), \quad \vec{c}^{(1)} = r_{d-1} \cdot \vec{c}^{(0)}$$

respectively, where $\bar{r}' = (D/2 - 1) - (r_{d-2} \cdots r_0)_2$.

- (4) Output the vector of shares, indexed by $[D]$, whose initial half is

$$\vec{c}^{(0)} - \vec{c}^{(1)} = (1 - r_{d-1}) \cdot \vec{c}^{(0)},$$

and whose latter half is

$$\begin{aligned} & (p_0^{(0)} - p_0^{(1)}) \cdot \vec{1}_{D/2} + \vec{c}^{(1)} \\ & = (1 - r_{d-1}) \cdot \vec{1}_{D/2} + r_{d-1} \cdot \vec{c}^{(0)}. \end{aligned}$$

5 Implementation and Evaluation

For an empirical evaluation, we implemented our protocol by realizing \mathcal{F}_{ABB} for a dishonest majority, using the SPDZ_{2k} protocol [20]. As a reminder, this uses authenticated additive secret sharing. We implemented the offline and online parts of our protocols separately:⁵

- (1) **Offline phase.** This is implemented using the well-known MP-SPDZ library [33]. In this phase, all Sign and ModLTZ gates are prepared, as described in Section 4.
- (2) **Online phase.** This is implemented in Rust. Each party executes a program to handle the protocol’s computation and communication. Precomputed Sign and ModLTZ gates are loaded into memory at the start.

⁵See <https://github.com/FhenixProtocol/thresholdthe-paper> for the code.

We ran all benchmarks on a single AWS c7a.32xlarge instance with 128 AMD EPYC 9R14 vCPUs and 256 GB of RAM. For the online phase, the tc utility was used to emulate network conditions and constraints. This utility allows precise control over network behavior, such as introducing configurable delays and bandwidth limitations. The selected network has a fully connected topology of nodes using TCP for pairwise communication.

To compare with the state of the art [21], we used the same LWE parameters $(n, q, p) = (1024, 2^{64}, 2)$ (hence $k = 64$, $m = 1$, $l = 63$), which are also commonly used for TFHE [16] in commercial implementations.⁶ Following Section 3.4, for the input bit length of the Sign gates we use $b = 8 \approx \sqrt{l}$, which implies that $d = \lceil l/b \rceil = 8$.

We note that [21] only reported their results, but did not provide a public implementation we could directly benchmark against. Therefore, we ran our experiments on a similar, though not identical, system. They also did not report direct throughput results, but their protocol is CPU-bound by the “squish-and-squash” (bootstrapping) operation, which works in conjunction with noise flooding. For that reason, we assume they must perform this computation serially on a single machine, and any parallelization would require scaling horizontally across machines (which would not improve throughput). Table 1 shows a benchmark for 4 parties with a network ping time of 1 ms and 1 Gbit bandwidth. We estimate that the online phase of our protocol has approximately 37 times better latency and about 20 000 times better throughput in this setting.

Protocol	Latency ms	Throughput (Dec/sec)
[21]	315.62	3.18
Our online	8.48	64 319

Table 1: Comparison of latency and throughput for 4 parties with 1 ms ping time and 1 Gbit bandwidth (online phase only).

5.1 Different Network Conditions

Table 2 shows how the end-to-end latency of the online phase changes according to the number of parties, or under increased network latency. Note that bandwidth does not affect the latency of a single decryption, since throughout the protocol the parties perform only a small amount of communication. By contrast, the network latency does have an effect, because the online phase does three sequential public openings, and in this implementation, two more rounds of authenticity checks.

	4 Parties	8 Parties	16 Parties
Ping Time 1 ms	8.483	10.086	24.230
Ping Time 10 ms	55.616	55.490	56.979

Table 2: Time (in ms) of the online decryption protocol, for a single ciphertext.

Our protocol achieves very high throughput. Since the online phase has very little communication and computation (in contrast

⁶For example, see <https://docs.zama.ai/tfhe-rs>.

to previous works), the throughput can scale to thousands and even tens of thousands of decryptions per second on a single server, regardless of network latency. This is captured in Table 3, which shows throughput under various conditions. It appears that the CPU is the bottleneck for 1 ms and 10 ms ping times (because there is little difference in throughput for these settings), but the network is the bottleneck for 100 ms ping times.

5.2 Preprocessing

To empirically benchmark our preprocessing protocols that construct Sign and ModLTZ gates, we implemented them using the MP-SPDZ [33] implementation of SPDZ_{2^k} . To provide good estimates for *our own protocols*, we separate out the costs to generate shared random bits and multiplication triples (see below for these), since there are many different ways to do this (e.g., [23, 28, 46]). All tests were run with the default MP-SPDZ value of two parties, 1 ms ping time and 1 Gbit/sec bandwidth. Table 4 shows the costs to produce a gate, as a function of its input bit length. (Both kinds of gates take roughly the same time to produce for the same input length, since the bulk of the work is done in PrepSubsetProds.) Table 5 shows that it takes only about 1.35 sec and 79 MB of online communication (per party) to preprocess enough material for 1000 runs of Π_{Decrypt} .

Triple-generation communication. Our preprocessing protocols use secure multiplications and random bit generation via \mathcal{F}_{ABB} . For each call to Π_{Decrypt} , the preprocessing uses only about bd random bits, but requires significantly more multiplications, so these are the main cost. Following the analysis in Section 4.1, the preprocessing does $2478 = d(2^b - b - 1) + (2^{d+1} - d - 2)$ multiplications (recall that $b = d = 8$). In SPDZ_{2^k} (and SPDZ more generally), each multiplication consumes one *Beaver triple* [3], a form of correlated randomness.

In our setting, Beaver triples should be generated in large batches offline, and the communication tends to be the bottleneck. Several works have focused on reducing the communication for maliciously secure triple generation; notable approaches include SPDZ_{2^k} [20], Overdrive2k [41], Mon \mathbb{Z}_{2^k} a [11], MH2k [14], LowGear 2.0 [45], and Multipars [32]. In terms of communication, Multipars is the current state of the art, with a reported amortized communication cost (per party) of 14.9 kbit per triple. In our protocols, this corresponds to approximately 4.5 MB of communication per decryption.

6 Related Work

Threshold FHE is a well studied problem, as it is a cornerstone for constant-round MPC (e.g., [2, 39]) and threshold cryptography more generally (e.g., [5, 13, 31]). And yet, all prior simulation-secure schemes use some flavor of noise flooding in their decryption protocols. The first proposed solution in the literature provided an actively secure protocol (with abort) for a small number of parties (due to the use of replicated secret sharing) [4]. This was later improved to full security and any number of parties [2]. Both works used somewhat expensive zero-knowledge proofs for active security, which was sufficient for low-depth circuits and non-real-time applications such as SPDZ pre-processing [26], but not in general.

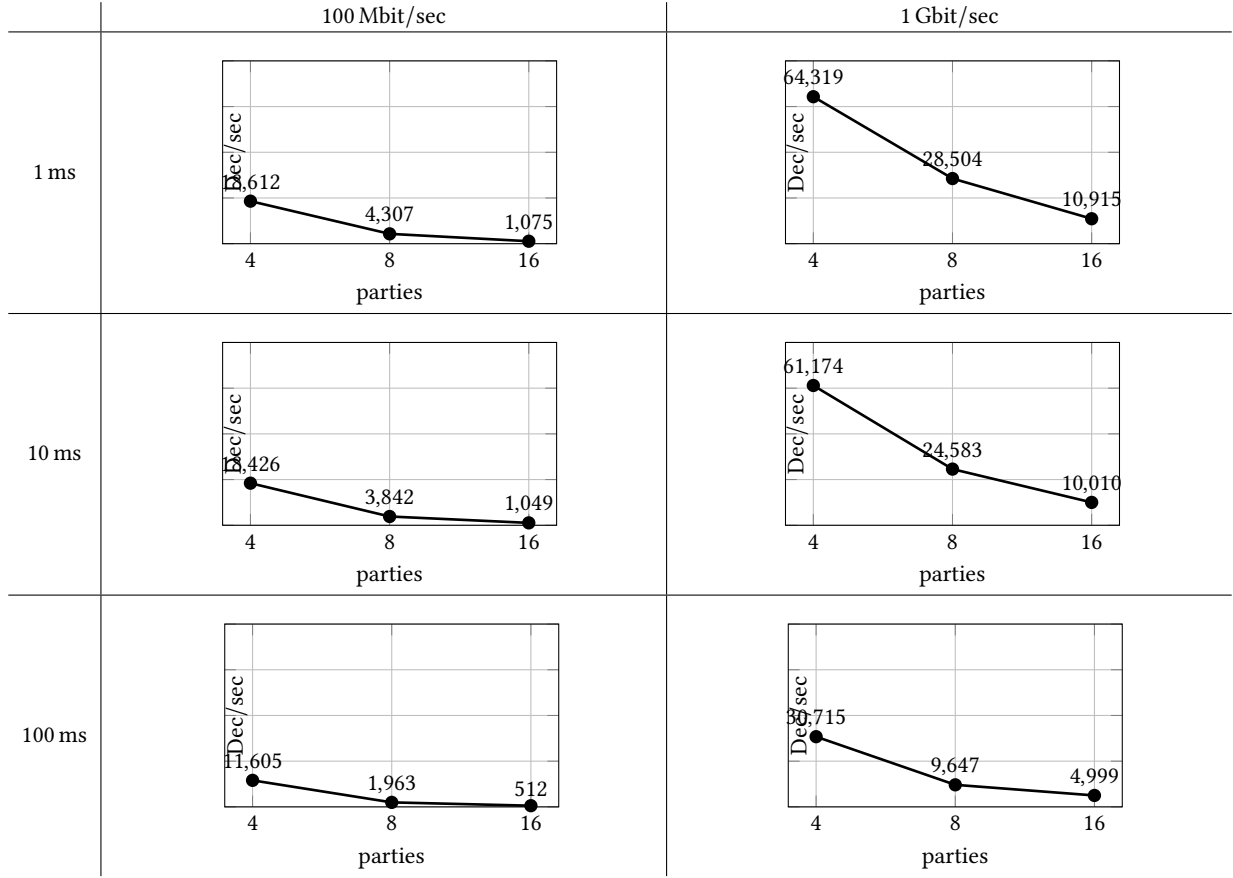


Table 3: Throughput of the online phase under varying network conditions, measured across different number of parties (4, 8, 16), bandwidths (100 Mbit/sec, 1 Gbit/sec), and ping times (1 ms, 10 ms and 100 ms).

Input length	Gates/sec	Communication per gate (kB)
4	94 004	0.35
5	44 770	0.83
6	13 428	1.82
7	10 966	3.84
8	6350	7.90
9	3400	16.06

Table 4: Throughput and communication of our gate-preprocessing protocol (ignoring the cost of triple and random-bit generation), for various bit-input lengths.

Gate	Gates/Dec	Time (sec)	Data (MB)
Sign	8	1.06	63.23
ModLTZ	1	0.29	16.06
Total	9	1.35	79.29

Table 5: Costs (ignoring triple and random-bit generation) of our preprocessing protocols, per thousand decryptions.

The work of [18] obtains a *robust* protocol with much better efficiency by assuming an honest (two-thirds) super-majority and utilizing well-known error-correction techniques.

The works of [5, 13] give a generic “universal thresholdizer” that converts any cryptosystem into a threshold variant. Like all prior schemes, they use noise flooding, which requires a super-polynomial modulus-to-noise ratio, making them too inefficient in practice. In addition, these works propose new variants of linear secret sharing that have large share sizes, leading to additional

computational and communication costs in practice (and a worse dependence on the total number of parties), and also potentially raise security concerns [17]. Recently, the concurrent and independent work of [9] improved this approach by obtaining a smaller ciphertext modulus, and preprocessing the zero-knowledge proofs of validity in an offline phase; however, it still uses noise flooding, and the modulus grows exponentially in the number of parties.

In recent years, probably due to recent advances in practical FHE and its applications (e.g., [1, 38, 48, 50, 53]), we have seen more work on efficient threshold FHE decryption. However, all of these works weaken the model or strengthen the assumptions in at least one of the following ways: adopting game-based security, limiting

the number of decryptions, assuming additional trusted parties, or making new non-standard hardness assumptions. The works of [6, 19, 22] avoid noise flooding and select masking noise from a much narrower range, but as a result they obtain only certain limited forms of game-based security, making it hard to reason about their security in a larger protocol. These schemes also suffer from other limitations and potential security vulnerabilities, as discussed in [42]. Alternatively, the work of [35] provides a very efficient semi-honest and simulation-secure threshold decryption scheme for LWE-based public-key encryption, without resorting to noise flooding. However, the protocol does not perform secure homomorphic operations on ciphertexts; additionally, its practically efficient version relies on a new ad-hoc hardness assumption called *Known-Norm Ring-LWE*. Other recent works [42, 51] obtain better efficiency by adopting specialized, non-standard models that assume incorruptible trusted parties.

We consider the recent work of [21] to be the current state of the art in terms of practical threshold FHE decryption. It uses a technique from [18] to obtain a robust protocol assuming an honest (two-thirds) super-majority. While this work significantly improves LWE parameters needed to support noise flooding, its latency and throughput are orders of magnitude worse than ours (see Section 5). Additionally, this work is limited to an honest majority, whereas ours can support any type of adversary model, as long as it has a realization of \mathcal{F}_{ABB} .

In summary, in contrast to this long line of prior works, ours is the only protocol to construct an *efficient* and *UC-secure* threshold FHE decryption scheme *without noise flooding*. Our protocol supports any adversary model that has a realization of \mathcal{F}_{ABB} realization, including a dishonest-majority adversary via SPDZ_{2k}, and an honest-majority adversary via standard Shamir secret-sharing techniques over Galois rings [27, 49]. For an honest two-thirds super-majority, our protocol is also robust via the same error-correcting techniques used in other works [18, 21].

References

- [1] A. Al Badawi, J. Bates, F. Bergamaschi, D. B. Cousins, S. Erabelli, N. Genise, S. Halevi, H. Hunt, A. Kim, Y. Lee, et al. OpenFHE: Open-source fully homomorphic encryption library. In *Proceedings of the 10th workshop on encrypted computing & applied homomorphic cryptography*, pages 53–63, 2022.
- [2] G. Asharov, A. Jain, A. López-Alt, E. Tromer, V. Vaikuntanathan, and D. Wichs. Multiparty computation with low communication, computation and interaction via threshold FHE. pages 483–501, 2012.
- [3] D. Beaver. Efficient multiparty protocols using circuit randomization. pages 420–432, 1991.
- [4] R. Bendlin and I. Damgård. Threshold decryption and zero-knowledge proofs for lattice-based cryptosystems. In *Theory of Cryptography Conference*, pages 201–218, 2010.
- [5] D. Boneh, R. Gennaro, S. Goldfeder, A. Jain, S. Kim, P. M. R. Rasmussen, and A. Sahai. Threshold cryptosystems from threshold fully homomorphic encryption. pages 565–596, 2018.
- [6] K. Boudgoust and P. Scholl. Simple threshold (fully homomorphic) encryption from LWE with polynomial modulus. pages 371–404, 2023.
- [7] E. Boyle, N. Chandran, N. Gilboa, D. Gupta, Y. Ishai, N. Kumar, and M. Rathee. Function secret sharing for mixed-mode and fixed-point secure computation. pages 871–900, 2021.
- [8] E. Boyle, N. Gilboa, and Y. Ishai. Function secret sharing. pages 337–367, 2015.
- [9] Z. Brakerski, O. Friedman, A. Marmor, D. Mutzari, Y. Spitzer, and N. Trieu. Threshold FHE with efficient asynchronous decryption. *Cryptology ePrint Archive*, Paper 2025/712, 2025.
- [10] R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *FOCS*, pages 136–145, 2001.
- [11] D. Catalano, M. D. Raimondo, D. Fiore, and I. Giacomelli. MonZ_{2k}: Fast maliciously secure two party computation on Z_{2k}. In *Public-Key Cryptography*, pages 357–386, 2020.
- [12] O. Catrina and S. de Hoogh. Improved primitives for secure multiparty integer computation. In *Security and Cryptography for Networks*, pages 182–199, 2010.
- [13] J. H. Cheon, W. Cho, and J. Kim. Improved universal thresholdizer from iterative shamir secret sharing. *Journal of Cryptology*, 38(1):15, 2025.
- [14] J. H. Cheon, D. Kim, and K. Lee. Mhz2k: MPC from HE over Z_{2k} with new packing, simpler reshare, and better ZKP. pages 426–456, 2021.
- [15] K. Chida, K. Hamada, D. Ikarashi, R. Kikuchi, D. Genkin, Y. Lindell, and A. Nof. Fast large-scale honest-majority MPC for malicious adversaries. *Journal of Cryptology*, 36(3):15, 2023.
- [16] I. Chillotti, N. Gama, M. Georgieva, and M. Izabachène. TFHE: Fast fully homomorphic encryption over the torus. *Journal of Cryptology*, 33(1):34–91, 2020.
- [17] W. Cho, J. Kim, and C. Lee. (In)Security of threshold fully homomorphic encryption based on Shamir secret sharing. *Cryptology ePrint Archive*, 2024.
- [18] A. Choudhury, J. Loftus, E. Orsini, A. Patra, and N. P. Smart. Between a rock and a hard place: Interpolating between MPC and FHE. pages 221–240, 2013.
- [19] S. Chowdhury, S. Sinha, A. Singh, S. Mishra, C. Chaudhary, S. Patranabis, P. Mukherjee, A. Chatterjee, and D. Mukhopadhyay. Efficient threshold FHE with application to real-time systems. *Cryptology ePrint Archive*, 2022.
- [20] R. Cramer, I. Damgård, D. Escudero, P. Scholl, and C. Xing. SPDZ_{2k}: Efficient MPC mod 2^k for dishonest majority. pages 769–798, 2018.
- [21] M. Dahl, D. Demmler, S. E. Kazdadi, A. Meyre, J. Orfila, D. Rotaru, N. P. Smart, S. Tap, and M. Walter. Noah’s ark: Efficient threshold-FHE using noise flooding. In *Workshop on Encrypted Computing & Applied Homomorphic Cryptography*, pages 35–46, 2023.
- [22] X. Dai, W. Wu, and Y. Feng. Key lifting: Multi-key fully homomorphic encryption in plain model without noise flooding. *Cryptology ePrint Archive*, 2022.
- [23] I. Damgård, D. Escudero, T. K. Frederiksen, M. Keller, P. Scholl, and N. Volgushev. New primitives for actively-secure MPC over rings with applications to private machine learning. In *IEEE Symposium on Security and Privacy*, pages 1102–1120, 2019.
- [24] I. Damgård, M. Fitzi, E. Kiltz, J. B. Nielsen, and T. Toft. Unconditionally secure constant-rounds multi-party computation for equality, comparison, bits and exponentiation. In *Theory of Cryptography Conference*, pages 285–304, 2006.
- [25] I. Damgård and J. B. Nielsen. Universally composable efficient multiparty computation from threshold homomorphic encryption. pages 247–264, 2003.
- [26] I. Damgård, V. Pastro, N. P. Smart, and S. Zakarias. Multiparty computation from somewhat homomorphic encryption. *Lecture Notes in Computer Science*, pages 643–662, 2012.
- [27] D. Escudero. *Multiparty Computation over Z/2^kZ*. PhD thesis, University of Aarhus, 2021.
- [28] D. Escudero, S. Ghosh, M. Keller, R. Rachuri, and P. Scholl. Improved primitives for MPC over mixed arithmetic-binary circuits. In *CRYPTO*, pages 823–852, 2020.
- [29] D. Escudero, C. Xing, and C. Yuan. More efficient dishonest majority secure computation over Z_{2k} via galois rings. pages 383–412, 2022.
- [30] C. Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, pages 169–178, 2009.
- [31] K. D. Gür, J. Katz, and T. Silde. Two-round threshold lattice-based signatures from threshold homomorphic encryption. In *Post-Quantum Cryptography*, pages 266–300, 2024.
- [32] S. Hasler, P. Reiser, M. Rivinius, and R. Küsters. Multipars: Reduced-communication MPC over Z_{2k}. *Proc. Priv. Enhancing Technol.*, 2024(2):5–28, 2024.
- [33] M. Keller. MP-SPDZ: A versatile framework for multi-party computation. pages 1575–1590, 2020.
- [34] H. Lipmaa and T. Toft. Secure equality and greater-than tests with sublinear online complexity. pages 645–656, 2013.
- [35] D. Micciancio and A. Suhl. Simulation-secure threshold PKE from LWE with polynomial modulus. *IACR Commun. Cryptol.*, 1(4):2, 2024.
- [36] P. Mohassel and Y. Zhang. SecureML: A system for scalable privacy-preserving machine learning. In *IEEE Symposium on Security and Privacy*, pages 19–38, 2017.
- [37] C. Mouchet, J. R. Troncoso-Pastoriza, J. Bossuat, and J. Hubaux. Multiparty homomorphic encryption from ring-learning-with-errors. In *Privacy Enhancing Technologies Symposium*, volume 2021, pages 291–311, 2021.
- [38] C. V. Mouchet, J.-P. Bossuat, J. R. Troncoso-Pastoriza, and J.-P. Hubaux. Lattigo: A multiparty homomorphic encryption library in Go. In *Workshop on Encrypted Computing and Applied Homomorphic Cryptography*, pages 64–70, 2020.
- [39] P. Mukherjee and D. Wichs. Two round multiparty computation via multi-key FHE. pages 735–763, 2016.
- [40] T. Nishide and K. Ohta. Multiparty computation for interval, equality, and comparison without bit-decomposition protocol. In *Public Key Cryptography*, pages 343–360, 2007.
- [41] E. Orsini, N. P. Smart, and F. Vercauteren. Overdrive2k: Efficient secure MPC over Z_{2k} from somewhat homomorphic encryption. In *CT-RSA*, pages 254–283, 2020.
- [42] A. Passelègue and D. Stehlé. Low communication threshold fully homomorphic encryption. In *ASIACRYPT*, pages 297–329, 2024.

- [43] C. Peikert. A decade of lattice cryptography. *Foundations and Trends in Theoretical Computer Science*, 10(4):283–424, 2016.
- [44] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):1–40, 2009. Preliminary version in STOC 2005.
- [45] P. Reisert, M. Rivinius, T. Krips, and R. Küsters. Overdrive LowGear 2.0: Reduced-bandwidth MPC without sacrifice. In *ASIA CCS*, pages 372–386, 2023.
- [46] D. Rotaru and T. Wood. MARBled circuits: Mixing arithmetic and boolean circuits with active security. pages 227–249, 2019.
- [47] T. Ryyffel, P. Tholoniati, D. Pointcheval, and F. R. Bach. AriaNN: Low-interaction privacy-preserving deep learning via function secret sharing. In *Privacy Enhancing Technologies Symposium*, volume 2022, pages 291–316, 2022.
- [48] S. Sav, A. Pyrgelis, J. R. Troncoso-Pastoriza, D. Froelicher, J. Bossuat, J. S. Sousa, and J. Hubaux. POSEIDON: privacy-preserving federated neural network learning. In *Network and Distributed System Security Symposium*, 2021.
- [49] A. Shamir. How to share a secret. *Commun. ACM*, 22(11):612–613, 1979.
- [50] R. Solomon, R. Weber, and G. Almasaqbeh. smartfhe: Privacy-preserving smart contracts from fully homomorphic encryption. In *EuroS&P*, pages 309–331, 2023.
- [51] Y. Sugizaki, H. Tsuchida, T. Hayashi, K. Nuida, A. Nakashima, T. Ishihiki, and K. Mori. Threshold fully homomorphic encryption over the torus. pages 45–65, 2023.
- [52] S. Wagh, D. Gupta, and N. Chandran. SecureNN: 3-party secure computation for neural network training. In *Privacy Enhancing Technologies Symposium*, pages 26–49, 2019.
- [53] G. Zyskind, Y. Erez, T. Langer, I. Grossman, and L. Bondarevsky. FHE-Rollups: Scaling confidential smart contracts on ethereum and beyond. In *ACM International Symposium on Blockchain and Secure Critical Infrastructure*, pages 1–9, 2024.
- [54] G. Zyskind, T. South, and A. Pentland. Don't forget private retrieval: distributed private similarity search for large language models. *CoRR*, abs/2311.12955, 2023.
- [55] G. Zyskind, A. Yanai, and A. S. Pentland. High-throughput three-party DPFs with applications to ORAM and digital currencies. pages 4152–4166, 2024.

A Security of Π_{Decrypt}

We establish the security of our protocol in the universal composability (UC) framework [10]. We first restate Theorem 1 for convenience, then prove it.

Theorem 1. *Protocol Π_{Decrypt} (Figure 3) perfectly realizes $\mathcal{F}_{\text{Decrypt}}$ in the \mathcal{F}_{ABB} -hybrid model, under the same (adaptive corruption) shell.*

PROOF. We define an ideal-process adversary $\mathcal{S}_{\text{Decrypt}}$ (Figure 4)—i.e., a *simulator*—that runs in the ideal execution with $\mathcal{F}_{\text{Decrypt}}$, and simulates all the messages from the dummy adversary \mathcal{A} to the environment \mathcal{Z} in the protocol Π_{Decrypt} . Because Π_{Decrypt} has no direct communication between parties (all interaction is via \mathcal{F}_{ABB}), these messages consist entirely of forwarded messages from \mathcal{F}_{ABB} to \mathcal{A} , which $\mathcal{S}_{\text{Decrypt}}$ must simulate using its access to $\mathcal{F}_{\text{Decrypt}}$. We show that \mathcal{Z} 's view of the execution of Π_{Decrypt} with dummy adversary \mathcal{A} is *identically distributed* to its view of the ideal execution of $\mathcal{F}_{\text{Decrypt}}$ with $\mathcal{S}_{\text{Decrypt}}$. We focus on the case of adaptive corruptions, of which static corruptions are a special case.

The main ideas behind the simulator are as follows. Observe that in the protocol Π_{Decrypt} , the honest parties do not have any secret inputs or outputs; the only secret values are the random bits chosen by \mathcal{F}_{ABB} (and any values derived from them, like the secret key). Also, in the protocol, \mathcal{F}_{ABB} opens only *randomly masked* intermediate values, along with the final decrypted message. So, $\mathcal{S}_{\text{Decrypt}}$ merely runs the code of the honest parties together with the *shell* of \mathcal{F}_{ABB} (which reflects the adversarial model), but not its *core code* (which makes the random choices and performs arithmetic operations). Instead, $\mathcal{S}_{\text{Decrypt}}$ simulates the opened masked values

by choosing them uniformly at random, and it uses the decrypted plaintext provided by $\mathcal{F}_{\text{Decrypt}}$ as the final opened value.

Adaptive corruption of certain subsets of parties can cause the shell to reveal the functionality's internal state to the adversary. Specifically, in the real execution, \mathcal{F}_{ABB} would reveal all its internal random choices to \mathcal{A} ; and in the corresponding ideal execution, $\mathcal{F}_{\text{Decrypt}}$ would reveal just its stored secret key (which is its only internal state) to $\mathcal{S}_{\text{Decrypt}}$. In this event, to simulate \mathcal{F}_{ABB} 's random choices, $\mathcal{S}_{\text{Decrypt}}$ simply computes the corresponding unmasked values using the secret key, then solves for (the bits of) the corresponding masks, and uses these as the simulated internal random choices of \mathcal{F}_{ABB} . The remainder of the simulation follows the Π_{Decrypt} protocol exactly, running the full \mathcal{F}_{ABB} (i.e., shell and core) starting from its simulated internal state.

The shell of \mathcal{F}_{ABB} may (or may not) allow \mathcal{Z} , via \mathcal{A} , to delay messages between \mathcal{F}_{ABB} and the honest parties, in particular modeling security with abort (or guaranteed output delivery, respectively). In the simulated interaction between the honest parties and \mathcal{F}_{ABB} 's shell, $\mathcal{S}_{\text{Decrypt}}$ merely follows \mathcal{Z} 's dictates in this respect. Also, $\mathcal{S}_{\text{Decrypt}}$ handles delivery of each Decrypt message between an honest party and $\mathcal{F}_{\text{Decrypt}}$ according to how \mathcal{Z} controls delivery of the *final* Open command and response between the party and \mathcal{F}_{ABB} (see Line (3) of Figures 3 and 4).

To reflect the modular structure of Π_{Decrypt} , we define $\mathcal{S}_{\text{Decrypt}}$ via a collection of simulator *subroutines* $\mathcal{S}_{\text{Mod}_{I,k}}$, $\mathcal{S}_{\text{LTRand}_{I,k}}$, etc., which respectively simulate the honest parties's interactions with, and internal random choices of, \mathcal{F}_{ABB} in the corresponding protocol subroutines $\text{Mod}_{I,k}$, $\text{LTRand}_{I,k}$, etc. In particular, we assume that there is a correct simulator subroutine $\mathcal{S}_{\text{KeyGen}}$ for the \mathcal{F}_{ABB} -hybrid KeyGen protocol subroutine. Each simulator subroutine takes the same inputs and produces the same outputs as its corresponding protocol subroutine. Recall that none of these are secret—the notation $[z]_k$ is just an alias for the unique *public identifier* id_z of a value z stored in \mathcal{F}_{ABB} .

By the correctness of the $\text{Mod}_{I,k}$ protocol subroutine (Lemma 2) and the fact that the simulation proceeds exactly according to the real protocol execution *except* for the openings of uniformly random masked values, it can be seen by inspection that the simulation is identically distributed to the real protocol execution. \square

Procedure 10: $\mathcal{S}_{\text{Mod}_{I,k}}(\text{sid}, [z]_k)$ (cf. Procedure 1)

- (1) Run $\mathcal{S}_{\text{LTRand}_{I,k}}(\text{sid}, [z]_k)$ to get $z', [r]_k, [u]_k$.
- (2) Send the LinComb command corresponding to the computation of $[e]_k$ as in Line (2) of $\text{Mod}_{I,k}$, and output $[e]_k$.

If the value of z is given, pass it to $\mathcal{S}_{\text{LTRand}_{I,k}}(\text{sid}, [z]_k)$.

The procedures $\mathcal{S}_{\text{SubsetProdsToSigns}_b}$ and $\mathcal{S}_{\text{SubsetProdsToModLTZ}_{d+1}}$ simply send the LinComb commands corresponding to the linear combinations evaluated by the linear functions $\text{SubsetProdsToSigns}_b$ and $\text{SubsetProdsToModLTZ}_{d+1}$, respectively. These functions do not use the Open or RandBit commands, so we do not need to simulate any openings or internal random choices of \mathcal{F}_{ABB} .

Simulator $\mathcal{S}_{\text{Decrypt}}$

Run the honest parties' protocol code together with the shell of \mathcal{F}_{ABB} , simulating Open outputs from \mathcal{F}_{ABB} , as follows.

- When a dummy honest party attempts to send (Init, sid) to $\mathcal{F}_{\text{Decrypt}}$, run that party's protocol code on that command within the subroutine $\mathcal{S}_{\text{KeyGen}}(\text{sid})$.
- When a dummy honest party attempts to send (Decrypt, sid, c) to $\mathcal{F}_{\text{Decrypt}}$:
 - (1) As in Line (1) of Π_{Decrypt} , send (to \mathcal{F}_{ABB} 's shell) the LinComb command corresponding to the computation of $[z]_k$.
 - (2) Call $\mathcal{S}_{\text{ModLTZ}}(\text{sid}, [z]_k)$, which outputs $[e]_k$.
 - (3) As in Line (3) of Π_{Decrypt} :
 - send (to \mathcal{F}_{ABB} 's shell) the LinComb command corresponding to $[\mu']_k = [z]_k - [e]_k$;
 - send (to \mathcal{F}_{ABB} 's shell) the (Open, sid, $[\mu']_k, k$) command, and when \mathcal{Z} delivers it, deliver the above (Decrypt, sid, c) command to $\mathcal{F}_{\text{Decrypt}}$;
 - upon receipt of (Decrypt, sid, c, μ) from $\mathcal{F}_{\text{Decrypt}}$, send (from \mathcal{F}_{ABB} 's shell) the tuple (Open, sid, $[\mu']_k, k, \mu' = L \cdot \mu \in \mathbb{Z}_q$) corresponding to the opening of μ' .

Upon a corruption, if the secret key s is revealed by $\mathcal{F}_{\text{Decrypt}}$, for each ciphertext c from a prior Decrypt command, give $z = \langle c, s \rangle \in \mathbb{Z}_q$ to the corresponding run of $\mathcal{S}_{\text{ModLTZ}}(\text{sid}, [z]_k)$.

Figure 4: Simulator for Π_{Decrypt} (cf. Figure 3)

Procedure 13: $\mathcal{S}_{\text{PrepSign}_{b,d+1}}(\text{sid})$ (cf. Procedure 5)

- (1) Call $\mathcal{S}_{\text{PrepSubsetProds}_{b,d+1}}(\text{sid})$ to simulate the preparation of $[r]_k, [\tilde{p}]_{d+1}$.
- (2) Output $[r]_k, \mathcal{S}_{\text{SubsetProdsToSigns}_b}([\tilde{p}]_{d+1})$.

If the value of r is given, pass it to $\mathcal{S}_{\text{PrepSubsetProds}_{b,d+1}}(\text{sid})$.

Procedure 14: $\mathcal{S}_{\text{PrepModLTZ}_{d+1,k}}(\text{sid})$ (cf. Procedure 7)

- (1) Call $\mathcal{S}_{\text{PrepSubsetProds}_{d+1,m}}(\text{sid})$ to simulate the preparation of $[r]_{d+1}, [\tilde{p}]_m$.
- (2) Output $[r]_{d+1}, \mathcal{S}_{\text{SubsetProdsToModLTZ}_{d+1}}([\tilde{p}]_m)$.

To simulate the application of the gate on $[y]_{d+1}$:

- (1) Send the LinComb command corresponding to the computation of $[y']_{d+1} = [y]_{d+1} + [r]_{d+1}$, and the command (Open, sid, $[y']_{d+1}, d+1$).
- (2) Once \mathcal{F}_{ABB} 's shell would respond, choose $y' \leftarrow \mathbb{Z}_{2D}$ and reply with the tuple (Open, sid, $[y']_{d+1}, d+1, y'$).

If the value of y is given after y' is opened, compute $r = y' - y \bmod 2D$ and give it to $\mathcal{S}_{\text{PrepSubsetProds}_{d+1,m}}(\text{sid})$.

Procedure 11: $\mathcal{S}_{\text{LTrand}_{l,k}}(\text{sid}, [z]_k)$ (cf. Procedure 3)

Preprocessing phase:

- (1) Call $\mathcal{S}_{\text{PrepModLTZ}_{d+1,m}}(\text{sid})$ (Procedure 14) to simulate the preparation of a ModLTZ gate.
- (2) For each $i \in [d-1]$, call $\mathcal{S}_{\text{PrepSign}_{b,d+1}}(\text{sid})$ (Procedure 13) to simulate the preparation of a Sign gate with masking value $[r_i]_k$. For $i = d-1$, do the same but with b', B' in place of b, B (respectively).
- (3) Send the LinComb command corresponding to the computation of $[r]_k$, as in Line (3) of $\text{LTrand}_{l,k}$.

Online phase:

- (1) As in Line (1) of $\text{LTrand}_{l,k}$, send the LinComb command corresponding to $[z']_k = [z]_k + [r]_k$ and the command (Open, sid, $[z']_l, l$). Once \mathcal{F}_{ABB} 's shell would respond, choose $z' \leftarrow [L]$ and reply with the tuple (Open, sid, $[z']_l, l, z'$).
- (2) As in Line (2) of $\text{LTrand}_{l,k}$, express z' in base B .
- (3) For each $i \in [d]$, let $[y_i]_{d+1}$ be (the identifier of) the z'_i th entry of the i th precomputed Sign gate.
- (4) As in Line (4) of $\text{LTrand}_{l,k}$, send the LinComb tuple corresponding to the computation of $[y]_{d+1}$.
- (5) Call $\mathcal{S}_{\text{PrepModLTZ}_{d+1,m}}(\text{sid})$ to simulate the application (as in Section 2.3) of the ModLTZ gate on $[y]_{d+1}$, yielding $[u]_k$.
- (6) Output $z', [r]_k, [u]_k$.

Internal randomness upon corruption:

If the value of z is given after z' is opened in Line (1):

- (1) Let $r = z' - z \bmod L$, and write r in base B , as $r = \sum_{i \in [d]} r_i \cdot B^i$ for $r_i \in [B]$.
- (2) For each $i \in [d]$, give r_i to the corresponding run of $\mathcal{S}_{\text{PrepSign}_{b,d+1}}(\text{sid})$.
- (3) Compute y from the z'_i, r_i (as in Lines (3) and (4) of Procedure 3) and give it to $\mathcal{S}_{\text{PrepModLTZ}_{d+1,m}}$ (if the application of the gate has already been simulated).

Procedure 12: $\mathcal{S}_{\text{PrepSubsetProds}_{a,c}}(\text{sid})$ (cf. Procedure 4)

- (1) If $a = 1$, send the command (RandBit, sid, $[r]_k$), and output $[r]_k, ([1]_c, [r]_c)$.
- (2) As in Line (2) of Procedure 4, write $a = a_0 + a_1$ and call (in parallel) $\mathcal{S}_{\text{PrepSubsetProds}_{a_0,c}}(\text{sid})$ and $\mathcal{S}_{\text{PrepSubsetProds}_{a_1,c}}(\text{sid})$ to simulate the preparation of $[r_0]_k, ([p_{S_0}]_c)_{S_0 \subseteq [a_0]}$ and $[r_1]_k, ([p'_{S_1}]_c)_{S_1 \subseteq [a_1]}$, respectively.
- (3) Send the LinComb command corresponding to $[r]_k = [r_0]_k + 2^{a_0} \cdot [r_1]_k$.
- (4) For each $S_0 \subseteq [a_0]$ and nonempty $S_1 \subseteq [a_1]$, send the command (Mult, sid, $[p_{S_0 \cup (a_0 + S_1)}]_c, [p_{S_0}]_c, [p'_{S_1}]_c$).
- (5) Output $[r]_k, ([p_S]_c)_{S \subseteq [a]}$.

If, upon corruption, the value of $r \in [2^a]$ is given:

- If $a = 1$, reveal to \mathcal{Z} that $r \in \{0, 1\}$ was the random choice made by the above RandBit call (and the value of $[r]_k$).
- Otherwise, express $r = r_0 + 2^{a_0} \cdot r_1$ for $r_i \in [2^{a_i}]$ and $i \in \{0, 1\}$, and give r_i to $\mathcal{S}_{\text{PrepSubsetProds}_{a_i,c}}(\text{sid})$.