

Perfectly-Secure Multiparty Computation with Linear Communication Complexity over Any Modulus

Daniel Escudero¹, Yifan Song^{2,3}, and Wenhao Wang⁴

¹ J.P. Morgan AI Research & J.P. Morgan AlgoCRYPT CoE, NY, USA

² Institute for Theoretical Computer Science, Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing, P. R. CHINA

³ Shanghai Qi Zhi Institute, Beijing, P. R. CHINA

⁴ Yale University, CT, USA

Abstract. Consider the task of secure multiparty computation (MPC) among n parties with perfect security and guaranteed output delivery, supporting $t < n/3$ active corruptions. Suppose the arithmetic circuit C to be computed is defined over a finite ring $\mathbb{Z}/q\mathbb{Z}$, for an arbitrary $q \in \mathbb{Z}$. It is known that this type of MPC over such ring is possible, with communication that scales as $O(n|C|)$, *assuming* that q scales as $\Omega(n)$. However, for *constant-size* rings $\mathbb{Z}/q\mathbb{Z}$ where $q = O(1)$, the communication is actually $O(n \log n|C|)$ due to the need of the so-called ring extensions. In most natural settings, the number of parties is variable but the “datatypes” used for the computation are fixed (*e.g.* 64-bit integers). In this regime, no protocol with linear communication exists.

In this work we provide an MPC protocol in this setting: perfect security, G.O.D. and $t < n/3$ active corruptions, that enjoys linear communication $O(n|C|)$, even for *constant-size* rings $\mathbb{Z}/q\mathbb{Z}$. This includes as important particular cases small fields such as \mathbb{F}_2 , and also the ring $\mathbb{Z}/2^k\mathbb{Z}$. The main difficulty in achieving this result is that widely used techniques such as linear secret-sharing cannot work over constant-size rings, and instead, one must make use of ring extensions that add $\Omega(\log n)$ overhead, while *packing* $\Omega(\log n)$ ring elements in each extension element in order to amortize this cost. We make use of reverse multiplication-friendly embeddings (RMFEs) for this packing, and adapt recent techniques in *network routing* (Goyal *et al.* CRYPTO’22) to ensure this can be efficiently used for non-SIMD circuits. Unfortunately, doing this naively results in a restriction on the minimum width of the circuit, which leads to an extra additive term in communication of $\text{poly}(n) \cdot \text{depth}(C)$. One of our biggest technical contributions lies in designing novel techniques to overcome this limitation by packing elements that are distributed across *different* layers. To the best of our knowledge, all works that have a notion of packing (*e.g.* RMFE or packed secret-sharing) group gates across the same layer, and not doing so, as in our work, leads to a unique set of challenges and complications.

1 Introduction

In secure multiparty computation, or MPC for short, a set of parties P_1, \dots, P_n , each having an input x_i , want to compute a function of their inputs $y = f(x_1, \dots, x_n)$ in such a way that only the output y is produced, and nothing additional about the parties’ inputs is revealed. This can be done assuming a trusted party that receives the inputs, computes the function and returns the output while promising to leak nothing else, but the goal in MPC is to achieve the same guarantees without relying on such trusted party, using only communication among the parties. MPC protocols are resilient against a potential coalition of *corrupted* parties that cooperate, coordinated by an *adversary*, in order to break the remaining parties’ privacy. Several MPC protocols have been proposed through time, leading to practical constructions seen in recent years and even some real-world applications and deployments (*e.g.* see <https://mpc.cs.berkeley.edu/>).

MPC protocols are typically characterized by the amount of corrupted parties t they can tolerate, with three notable settings being $t < n/3$, $t < n/2$ and $t < n$. Protocols with $t \geq n/3$ must allow for some negligible statistical error (if $t < n/2$), or even make use of cryptographic assumptions (if $t \geq n/2$). In the $t < n/3$ case several results are known: not only we can obtain perfectly secure MPC, which safeguards the parties’ inputs regardless of the computational power of the adversary, but we can do so with several appealing features such as guaranteed output delivery

(G.O.D. , which ensures the successful completion of the protocol in spite of any adversarial misbehavior), and a communication complexity that requires each party to send in average an amount of messages that is independent of the total number of parties, and depends only on the circuit size of the function f (here we omit additive terms that are independent of the circuit size but only depend on the number of parties) [GLS19]. More precisely, we can represent the function f as an arithmetic circuit C comprised of input, output, addition and multiplication gates over some finite ring $\mathbb{Z}/q\mathbb{Z}$ of integers modulo an integer q , and if $|C|$ denotes the number of multiplication gates, it is possible to obtain perfectly secure MPC over $\mathbb{Z}/q\mathbb{Z}$ with G.O.D. where the total communication complexity is $O(n|C|)$ ring elements, which we refer to as *linear* communication complexity. A common choice is taking q to be a prime, but this claim holds for any *arbitrary* integer q : the *Chinese remainder theorem* reduces the general case to moduli of the form $q = p^k$, for a prime p and an integer k , and works such as [ACD⁺19] show that perfectly secure MPC with G.O.D. over these moduli is possible, with linear communication complexity. It is relevant to study the general modulus case as it contains several relevant cases such as $q = 2, 2^{32}$ or 2^{64} , which have received considerable attention in recent works due to practical benefits [DEF⁺19].

Unfortunately, the claim of linear communication is not entirely accurate. Shamir secret-sharing is a core technique to enable these results, and for this scheme to work over $\mathbb{Z}/p^k\mathbb{Z}$, p must be strictly larger than n . Once n becomes larger than p , a so-called *Galois ring extension* of degree $\Omega(\log n)$ must be used. It is still true that the communication is $O(n|C|)$ ring elements, but this time each ring element is $\Omega(\log n)$ bits long. Hence, asymptotic communication is truly $O(n \log n |C|)$. In [CCXY18], the authors partially address this issue relying on the technique of reverse multiplication-friendly embeddings (RMFEs), which allow them to remove the $\log n$ overhead when computing many copies of the same circuits, i.e., a SIMD circuit. However, for general circuits, directly applying their techniques does not work as we will discuss in Section 1.3. This leads to the following question:

Can we design perfectly secure MPC protocols for general circuits over a constant-size ring $\mathbb{Z}/q\mathbb{Z}$ whose asymptotic communication complexity scales as $O(n|C|)$?

Again, all current solutions have communication that scales as $O(n|C|)$ ring elements, but assuming that either the ring bit-size scales as $\log n$ or the underlying circuit is a SIMD circuit. For fixed-size rings (which is the natural setting in practice) and general circuits, ring extensions must be used, which leads to a communication of $O(n \log n |C|)$. Recall that this affects the very practically-motivated settings of binary computation (i.e. circuits over \mathbb{Z}_2), and also the `int32` and `int64` cases (i.e. circuits over $\mathbb{Z}/2^k\mathbb{Z}$ for $k = 32$ and $k = 64$).

1.1 Our Contribution

In this work we give an answer in the affirmative to the aforementioned question. We provide an MPC protocol with the following desirable features:

- Perfect security against an active adversary corrupting $t < n/3$ parties
- Securely computes circuits over $\mathbb{Z}/q\mathbb{Z}$ for *any constant* q
- The total number of elements in $\mathbb{Z}/q\mathbb{Z}$ communicated scales as $O(n|C|)$
- Guaranteed output delivery

As we have previously discussed, via CRT this task reduces to the task of computation over a ring $\mathbb{Z}/p^k\mathbb{Z}$ with constant p . For the sake of presentation we focus on $p = 2$, that is, computation over $\mathbb{Z}/2^k\mathbb{Z}$, but the ideas presented directly work for more general p .

Remark 1 (On security with abort). We note that even removing the G.O.D. condition, and settling with security with abort only, obtaining a protocol with the other properties is not simple, and is on its own a relevant and challenging open question. In this work we aim for the stronger notion of G.O.D.

Remark 2 (Cost of addition gates). In our work, (a linear amount of) communication is required for every *addition gate*, which is not the case if one settles for $O(n \log n |C|)$ complexity (where addition gates are for free in terms of communication). Looking ahead, this occurs due to a technique called *network routing*, originated in [GPS22,GPS21], which is used to exploit some notion of

packing while routing values through the circuit correctly. The work of [GPS22] (which is set in a different context than ours) also suffers from communication per addition gates, and avoiding this overhead in our work would likely lead to improvements to [GPS22]. If the circuit does not have substantially many more addition gates than multiplication gates (a reasonable setting in practice), then this condition becomes immaterial for our asymptotic $O(n|C|)$ claim.

1.2 Related Work

Perfectly secure MPC for $t < n/3$ with linear communication and G.O.D. has been studied in multiple works such as the one by Goyal, Liu, and Song [GLS19] and Beerliová-Trubíniová and Hirt [BTH08]. These are set specifically in the context of finite fields \mathbb{F}_{p^d} , where $p^d = \Omega(n)$. The work of Abspoel et al. [ACD⁺19] generalizes this to the ring $\mathbb{Z}/p^k\mathbb{Z}$ by using Galois ring extensions of degree $d = \Omega(n)$. Unfortunately, as we have pointed out, these techniques are not suitable for *constant-sized* rings since they add an $\Omega(\log n)$ overhead.

If it is known that t is *far* from $n/3$, that is, $t < n(\frac{1}{3} - \epsilon)$ for any constant $\epsilon > 0$, then it is possible to obtain perfectly fully secure MPC with a communication of $O(|C|)$ (independent of n) for fields \mathbb{F}_{p^d} with $p^d = \Omega(n)$, and $O(\log n|C|)$ for constant-sized fields [DIK10]. These techniques extend naturally to the case of $\mathbb{Z}/p^k\mathbb{Z}$ by using ring extensions as in [ACD⁺19], and to arbitrary $\mathbb{Z}/q\mathbb{Z}$ via CRT. This is better than our linear communication $O(n|C|)$, but it assumes the aforementioned gap $\epsilon > 0$. Our protocol works for the case $\epsilon = 0$: $t < n/3$, and n can be as small as $n = 3t + 1$.

In the $t < n/2$ case with *statistical* security, a similar situation exists: most protocols with G.O.D. that achieve linear communication complexity require a ring $\mathbb{Z}/p^k\mathbb{Z}$ with $p > n$ [BFO12], [GSZ20,ACD⁺19]. This state of affairs changed for the case of \mathbb{F}_p and *security with abort* for *constant* p in the recent work of [PS21]. The same ideas in [PS21] extend naturally to $\mathbb{Z}/p^k\mathbb{Z}$ for constant p . However, it is not clear how to extend the ideas in [PS21] to G.O.D. since, in the $t < n/2$ regime, the techniques used to achieve full security are considerably much more complex than these for $t < n/3$ (which are already quite intricate). The core difference is that in $t < n/2$ setting the central idea to achieve G.O.D., dispute control, differs from player elimination—in spite of sharing some similarities: unlike player elimination, in $t < n/2$ this pair cannot be removed because one may be inadvertently removing one of the $t + 1$ honest parties, and t honest parties alone cannot have enough joint information to finish the computation.

Finally, we note that if one only wants to achieve statistical security (rather than perfect security) in the G.O.D. setting for $t < n/3$, [IKP⁺16] has achieved $O(\text{poly}(\log(n))|C|)$ elements of communication if $t < n(1/2 - \epsilon)$, for any constant $\epsilon > 0$, which is even better than linear communication.

1.3 Overview of our Techniques

We begin by highlighting the difficulties that existing works face when considering MPC over constant-size rings. Perfectly secure MPC with linear communication (for non-constant fields) was first proposed in [BTH08], and it was later improved in [GLS19] to remove a term dependent on the circuit depth. These ideas can be generalized to rings such as $\mathbb{Z}/p^k\mathbb{Z}$ [ACD⁺19], again with the same complexity if $p = \Omega(n)$. This complexity comes from the use of Shamir secret-sharing, which requires enough interpolation points to operate, and hence it requires a ring of large enough characteristic. Shamir’s is not the only linear secret-sharing scheme one could use, but it is unlikely there exist other schemes that somehow reduce this secret/share size requirement, since this is highly connected to the MDS conjecture (see for example [FR22, Lemma 1]). Since Shamir secret-sharing seems unavoidable, our core idea is to use the ring extensions needed for it, while *packing* multiple entries in a single ring extension secret-shared value, so that the cost per single share is ultimately linear. Ultimately, the challenges lie on efficiently making use of this packing, which is our focus in this overview.

As we have mentioned previously, here and for the rest of our paper we focus on the case $\mathbb{Z}/2^k\mathbb{Z}$. Furthermore, for the sake of this introduction only we focus on security with abort (which, as mentioned previously, is already a challenging task on its own). We highlight towards the end of the section how to tackle the G.O.D. case. Our work makes use of several techniques in the literature such as player elimination [BTH08], reverse multiplication-friendly embeddings (RMFEs) [CCXY18], network routing [GPS22,GPS21], among others. In this section we provide a

high level overview of how these ideas are put together in order to obtain our protocol. First, we introduce some notation, with further details given in Section 2 and beyond. We let $\text{GR}(2^k, m)$ be a Galois ring extension of $\mathbb{Z}/2^k\mathbb{Z}$ of degree m , which can be seen as the ring of polynomials over $\mathbb{Z}/2^k\mathbb{Z}$ modulo an monic polynomial irreducible mod 2 of degree m (for more details see Section 2 or [Wan03]). We can perform Shamir secret-sharing over this ring if $m = \lceil \log n \rceil + 1$ [ACD⁺19], and we denote degree- d sharings of $x \in \text{GR}(2^k, m)$ by $[x]_d$. An RMFE is a pair of $\mathbb{Z}/2^k\mathbb{Z}$ -linear homomorphisms $(\phi : (\mathbb{Z}/2^k\mathbb{Z})^\ell \rightarrow \text{GR}(2^k, m), \psi : \text{GR}(2^k, m) \rightarrow (\mathbb{Z}/2^k\mathbb{Z})^\ell)$ satisfying $\psi(\phi(\mathbf{x}) \cdot \phi(\mathbf{y})) = \mathbf{x} \star \mathbf{y}$ for every $\mathbf{x}, \mathbf{y} \in (\mathbb{Z}/2^k\mathbb{Z})^\ell$. Such objects exist, with $\ell = \Theta(m)$ [CCXY18,ELXY23], and in fact we can also take them such that $\phi(\mathbf{1}) = 1$ [ELXY23]. This ensures that $\psi(\phi(\mathbf{x})) = \psi(\phi(\mathbf{x})\phi(\mathbf{1})) = \mathbf{x} \star \mathbf{1} = \mathbf{x}$.

Embedding using RMFEs. We may take as a starting point the approach in [CCXY18] in order to secret-share elements of $\mathbb{Z}/2^k\mathbb{Z}$ efficiently using Shamir secret-sharing via RMFEs. Recall that one (naive) way of secret-sharing an element $x \in \mathbb{Z}/2^k\mathbb{Z}$ using Shamir SS is to embed it in $\text{GR}(2^k, m)$, and then secret-sharing over $\text{GR}(2^k, m)$, which results in an undesired overhead of $m \approx \log n$. Instead, one may secret-share m elements simultaneously by interpreting them as an element of $\text{GR}(2^k, m)$, removing the overhead of m in an amortized sense (*i.e.* after dividing by the number of secrets m), but unfortunately this approach does not interact well with the MPC setting since one cannot easily multiply elements this way.

RMFEs are introduced in [CCXY18] as a solution to this problem. Instead of secret-sharing m elements by thinking of them as an element of $\text{GR}(2^k, m)$, ℓ elements $\mathbf{x} = (x_1, \dots, x_\ell) \in (\mathbb{Z}/2^k\mathbb{Z})^\ell$ are shared together by first mapping them as $x = \phi(\mathbf{x}) \in \text{GR}(2^k, m)$, and sharing $[x]_t$ instead. This has an overhead of m for ℓ secrets, which is $m/\ell = \Theta(1)$ amortized. Furthermore, the multiplicative property of RMFEs turns out to enable products on secret-shared data, making this embedding particularly suitable for MPC. In fact, this is used in [CCXY18] to remove the $\log n$ overhead of perfect security in the context in which the circuit C is structured as ℓ copies of the same function, run on possibly different inputs. Our goal however is to enable a more general class of circuits that have no specific structure in terms of their wiring, and hence this approach is insufficient.

To illustrate the main challenge, let us discuss how the approach from [CCXY18] works, for ℓ copies of the same circuit. The invariant the parties maintain is that, for every set of ℓ values $\mathbf{x} = (x_1, \dots, x_\ell) \in (\mathbb{Z}/2^k\mathbb{Z})^\ell$ corresponding to the same wire across the ℓ copies, the parties have sharings $[\phi(\mathbf{x})]_t$. Given ℓ copies of a multiplication gate with inputs \mathbf{x} and \mathbf{y} , the invariant is maintained by using a triple $([\phi(\mathbf{a})]_t, [\phi(\mathbf{b})]_t, [\phi(\mathbf{c})]_t)$, where $\mathbf{c} = \mathbf{a} \star \mathbf{b}$. Assuming the parties have $[\phi(\mathbf{x})]_t$ and $[\phi(\mathbf{y})]_t$, they can locally compute $[\phi(\mathbf{x}) + \phi(\mathbf{a})]_t$ and open this as $\phi(\mathbf{u})$, and similarly open $\phi(\mathbf{v})$ where $\mathbf{v} = \mathbf{y} + \mathbf{b}$. Crucially, since this is Shamir reconstruction, this can be done with error correction/detection, ensuring no errors are introduced. Now, the parties compute locally

$$[z]_t = \phi(\mathbf{u}) \cdot \phi(\mathbf{v}) - \phi(\mathbf{v}) \cdot [\phi(\mathbf{a})]_t - \phi(\mathbf{u}) \cdot [\phi(\mathbf{b})]_t + [\phi(\mathbf{c})]_t,$$

which satisfies $\psi(z) = \mathbf{x} \star \mathbf{y}$, thanks to the properties of RMFEs. Finally, the parties can execute a simple re-encoding protocol that applies $\phi \circ \psi : \text{GR}(2^k, m) \rightarrow \text{GR}(2^k, m)$ to $[z]_t$, yielding $[\phi(\mathbf{x} \star \mathbf{y})]_t$, hence preserving the desired invariant.

Now, if the circuit is non-SIMD, it could easily happen that, say, a value encoded in position 1 must be multiplied (or even added) with a value encoded in position 2. The properties of RMFEs only allow for the computation of $\mathbf{x} \star \mathbf{y} = (x_1y_1, \dots, x_\ell y_\ell)$, but they fall short if one somehow needs products that are “not aligned”, such as $(x_1y_2, x_2y_1, x_3y_3, \dots)$. SIMD circuits do not present such misalignments, which is why the techniques in [CCXY18] work in that setting.

Network routing. To alleviate this issue, we may resort to network routing, a general technique introduced in [GPS22,GPS21] to ensure the sharings of output groups can be rearranged in such a way that they are “aligned” when fed as inputs to future groups of gates. In these works, such “wiring” issues appeared in the context of packed secret-sharing where, as in our setting, parties have shares whose underlying secrets are *vectors* that can somehow be added or multiplied component-wise, but cannot be re-routed easily. In our work we show that such techniques can also be used in our context, where the “packing” is done with RMFEs. For the sake of keeping this overview as lightweight as possible, we will not dive into the details on how this is done, and instead we will refer the reader to Sections 3.4 and E where details are provided. For now, it suffices

to say that there is an efficient method for the parties to obtain sharings $[\phi(\mathbf{x})]_t$ for every input group in a given layer, starting from packed sharings of every output group of the previous layers.

For simplicity in the exposition we will assume from now on that the circuit only has multiplication gates, without any addition operation. We discuss at the end of this overview why this is convenient, and how to handle the case of addition gates as well. Naturally, our main protocols include the general case.

On the dependency on circuit width. Using the network routing techniques mentioned above, together with the protocol sketched earlier, we would obtain the desired result: perfectly secure MPC for $t < n/3$ over $\mathbb{Z}/2^k\mathbb{Z}$ with G.O.D. and linear communication $O(n)$. However, this result hides an assumption on the circuit structure: it requires each layer to contain at least $\Omega(n\ell)$ multiplication gates. This is because every layer requires parallel reconstructions for each of its multiplication gates, and robust reconstruction of Shamir sharings with *linear* communication (*i.e.* instead of all parties sending their shares to each other, which would be quadratic) requires $\Omega(n)$ reconstructions to be done in parallel [DN07]. On top of this each such Shamir sharing “packs” $\ell \approx O(\log n)$ values, which results in a requirement of $n\ell$ secrets to be reconstructed (per layer!) to obtain the communication gains.

The above results in a total communication that is not $O(n|C|)$, but rather $O(n|C| + n^2\ell \cdot \text{depth}(C))$. For circuits C such that $\text{depth}(C) \gg |C|/n$ (*e.g.* “skinny” circuits), this extra term dominates communication. In contrast, for the case of MPC over \mathbb{F}_p for large p , it is known that the $n^2\ell \cdot \text{depth}(C)$ term can be eliminated, resulting in true linear communication $O(n|C|)$ [GLS19]. This leaves a gap between what we know over (large) fields (which can be generalized for $\mathbb{Z}/p^k\mathbb{Z}$ for large p), and the case of $\mathbb{Z}/2^k\mathbb{Z}$. What follows is dedicated to discuss how we address this complexity gap.

Computing the circuit optimistically with additive secret-sharing. We first reduce the term $n^2\ell \cdot \text{depth}(C)$ to $n\ell \cdot \text{depth}(C)$ via the following core idea—also used in [GLS19]: derive *additive* sharings from Shamir sharings and use the additive sharings to compute the circuit (which enables cheating but can be done with linear communication without any minimum batch size requirement), and only use the Shamir sharings for a final verification check (which involves reconstructing many robust sharings in parallel and hence can be done with linear communication complexity). For the optimistic computation of the circuit we will make use of additive secret-sharing, which we denote by $\langle x \rangle$ for $x \in \mathbb{Z}/2^k\mathbb{Z}$ (we also extend this notation naturally to vectors over $\mathbb{Z}/2^k\mathbb{Z}$). For input gates, each client having a group of inputs $\mathbf{x} \in (\mathbb{Z}/2^k\mathbb{Z})^\ell$ secret-shares $[\phi(\mathbf{x})]_t$ towards the parties (the parties perform degree checks that ensure the degree is indeed $\leq t$), and then the parties locally derive $\langle \mathbf{x} \rangle$ from $[\phi(\mathbf{x})]_t$. This is done by first locally converting $[\phi(\mathbf{x})]_t$ to $\langle \phi(\mathbf{x}) \rangle$ (a standard procedure involving each party multiplying locally by certain Lagrange coefficients), followed by a *local* application of the mapping ψ , to obtain $\langle \psi(\phi(\mathbf{x})) = \mathbf{x} \rangle$. Now, the parties optimistically compute the circuit. For every pair of values to be multiplied $\langle x \rangle$ and $\langle y \rangle$, letting $i \in [\ell]$ be the index of this gate in its group, the parties can use the triple $([\phi(\mathbf{a})]_t, [\phi(\mathbf{b})]_t, [\phi(\mathbf{c})]_t)$ associated to this group to locally derive the additively shared triple $(\langle a_i \rangle, \langle b_i \rangle, \langle c_i \rangle)$, which they can use to multiply $\langle x \rangle$ and $\langle y \rangle$: open $u_i = x + a_i$ and $v_i = y + b_i$ (with linear communication without any minimum batch size restriction), and compute $\langle xy \rangle = u_i \cdot v_i - v_i \cdot \langle a \rangle - u_i \cdot \langle b \rangle + \langle c \rangle$.

For the check, recall that the parties have Shamir sharings of $[\phi(\mathbf{x})]_t$ for every group \mathbf{x} in the input layer. Consider a multiplication group in the first layer having as inputs \mathbf{x}, \mathbf{y} . Using network routing, the parties can obtain Shamir sharings $[\phi(\mathbf{x})]_t, [\phi(\mathbf{y})]_t$, so they can compute $[\phi(\mathbf{u})]_t = [\phi(\mathbf{x})]_t + [\phi(\mathbf{a})]_t$ and $[\phi(\mathbf{v})]_t = [\phi(\mathbf{y})]_t + [\phi(\mathbf{b})]_t$, hence obtaining robust versions of the reconstructed values \mathbf{u}, \mathbf{v} when these gates were computed optimistically. Furthermore, the parties can compute locally $[z]_t = \phi(\mathbf{u}) \cdot \phi(\mathbf{v}) - \phi(\mathbf{v}) \cdot [\phi(\mathbf{a})]_t - \phi(\mathbf{u}) \cdot [\phi(\mathbf{b})]_t + [\phi(\mathbf{c})]_t$, and apply the re-encoding protocol from [CCXY18] to obtain $[\phi(\mathbf{x} \star \mathbf{y})]_t$. Doing this for all groups in the first layer ensures that the parties obtain robust sharings $[\phi(\mathbf{z})]_t$ of all output groups \mathbf{z} in the first layer. This process can be iterated: the parties use network routing to obtain robust sharings of all input groups in the second layer, obtain robust sharings of all reconstructed values \mathbf{u} and \mathbf{v} during the optimistic computation of this layer, and then compute robust sharings of the outputs groups in this layer. This is repeated until the output layer is reached. Before reconstructing the outputs, however, the parties robustly reconstruct in parallel all Shamir sharings $[\mathbf{u}]_t$ and $[\mathbf{v}]_t$ corresponding

to *all* multiplication groups, cross checking with the values reconstructed optimistically.⁵ The key point is that this robust reconstruction involves enough values as to be able to enjoy amortized linear communication complexity.

Packing across different multiplication layers. Using additive SS optimistically during the circuit computation shaves a factor of n in the term $n^2\ell \cdot \text{depth}(C)$, but it still leaves us with a term $n\ell \cdot \text{depth}(C)$, since we still need to RMFE-pack ℓ gates across the *same* layer. To address this we allow groups to contain gates spanning over potentially *different* layers, with the only restriction being that the set of groups must form a DAG (we say there is an edge from group A to group B if at least one gate in A connects to one gate in B ; in Section 3.1 we show that such grouping can always be done for *any* circuit—without any width restrictions). For notational convenience we assume that the gates in every group are indexed in increasing topological order, meaning that if a gate indexed by i in a group depends on the output of another gate indexed by j in the same group, then $j < i$. As before, we still group each client’s input wires into batches of size ℓ each,⁶ and we also let each client secret-shares a group of inputs \mathbf{x} as $[\phi(\mathbf{x})]_t$.

Careful observation reveals that such relaxation in grouping does not really affect the optimistic computation using additive SS: the parties can still derive additive SS triples $(\langle a_i \rangle, \langle b_i \rangle, \langle c_i \rangle)$ and use these to perform multiplications. However, complications will arise at the verification stage. To illustrate this, let us consider the first multiplication group in topological order, and let us denote its inputs as \mathbf{x} and \mathbf{y} . If it was the case that all of the multiplications in the group belonged to the first layer, then the parties could perform network routing as before to obtain $[\phi(\mathbf{x})]_t$ and $[\phi(\mathbf{y})]_t$. Unfortunately, since we are packing across different layers, it can happen that, say, the left input x_i to the i -th gate is equal to the output $x_j y_j$ of the j -th gate in the same group, with $j < i$. This prevents us from determining $[\phi(\mathbf{x})]_t$ from the input layer alone.

Determining $[\phi(\mathbf{x})]_t$ for the first multiplication group. For illustration, let us keep our focus on the first multiplication group. Our goal is to show how the parties can obtain robust sharings $[\phi(\mathbf{x})]_t$, which, as illustrated above, cannot be derived from the input layer alone since the left input x_i to get i depends on the j -th output $x_j y_j$ of gate $j < i$. First, recall that the parties have opened $\langle \mathbf{u} \rangle = \langle \mathbf{x} \rangle + \langle \mathbf{a} \rangle$ and $\langle \mathbf{v} \rangle = \langle \mathbf{y} \rangle + \langle \mathbf{b} \rangle$, in the process of processing this group optimistically. However, due to the non-robustness of additive SS, these reconstructions may have resulted in $\mathbf{u}' = \mathbf{u} + \boldsymbol{\delta}$ and $\mathbf{v}' = \mathbf{v} + \boldsymbol{\epsilon}$, for some possibly non-zero $\boldsymbol{\delta}$ and $\boldsymbol{\epsilon}$. Using these reconstructions, the parties can compute locally $[\phi(\mathbf{x} + \boldsymbol{\delta})]_t = \phi(\mathbf{u}') - [\phi(\mathbf{a})]_t$ and $[\phi(\mathbf{y} + \boldsymbol{\epsilon})]_t = \phi(\mathbf{v}') - [\phi(\mathbf{b})]_t$, which correspond to robust sharings of the inputs \mathbf{x} and \mathbf{y} , but *potentially incorrect*. Now we let the parties execute *any* correct multiplication protocol (we borrow the protocol from [BTH08,ACD⁺19] for this purpose) to compute the product $[\phi(\mathbf{x} + \boldsymbol{\delta}) \cdot \phi(\mathbf{y} + \boldsymbol{\epsilon})]_t$, followed by re-encoding as in [CCXY18] (*i.e.* applying $\phi \circ \psi$) to obtain $[\phi((\mathbf{x} + \boldsymbol{\delta}) \star (\mathbf{y} + \boldsymbol{\epsilon}))]_t$.

For simplicity let us assume that i is the only index in this group whose corresponding gate depends on other outputs from the same group, with all the other gates receiving inputs directly from the input layer. In particular, both x_j and y_j come from the input layer. Recall that the goal is to obtain $[\phi(\mathbf{x})]_t$ using network routing. Since the parties have robust sharings of all groups in the input layer, they have almost all the pieces needed to obtain $[\phi(\mathbf{x})]_t$, with the only missing part being robust sharings that contain the output of the j -th gate, since this is needed for the i -th left input x_i . Our idea is to use, for the missing j -th output, the sharings $[\phi((\mathbf{x} + \boldsymbol{\delta}) \star (\mathbf{y} + \boldsymbol{\epsilon}))]_t$ computed above, which contain the j -th output $x_j y_j + \gamma$, where $\gamma = x_j \epsilon_j + y_j \delta_j + \delta_j \epsilon_j$, in the j -th entry. In other words, the parties perform network routing on the sharings from the input layer *and* the sharing $[\phi((\mathbf{x} + \boldsymbol{\delta}) \star (\mathbf{y} + \boldsymbol{\epsilon}))]_t$ to obtain $[\phi(\mathbf{x})]_t$. However, since the i -th entry corresponds to $x_i + \gamma$, the actual secret is $[\phi(\mathbf{x} + \gamma \mathbf{e}_i)]_t$. In other words, the parties do not obtain robust sharings of the *correct* \mathbf{x} , but instead, the i -th entry is shifted by γ .

This may raise a red flag at first sight: recall we are using the sharing $[\phi(\mathbf{x} + \gamma \mathbf{e}_i)]_t$ to verify the multiplications in the first group, in particular, verifying that $\langle \mathbf{u} \rangle = \langle \mathbf{x} \rangle + \langle \mathbf{a} \rangle$ was opened correctly. However, we are using an incorrect $[\phi(\mathbf{x} + \gamma \mathbf{e}_i)]_t$, so it may be the case that this somehow helps the

⁵ As shown in [GLS19], certain care is needed to ensure that security is not broken by delaying verification.

We omit these details in this overview.

⁶ We assume each client has $\Omega(\ell) = \Omega(n)$ inputs. Otherwise there is a minor overhead due to packing, but this is only restricted to the input layer.

adversary reconstruct $\langle \mathbf{u} \rangle$ incorrectly. For example, the adversary may be able to reconstruct $\langle \mathbf{u} \rangle$ as $\mathbf{u} + \gamma \mathbf{e}_i$, which is consistent with the robust sharings held by the parties $[\phi(\mathbf{x} + \gamma \mathbf{e}_i) + \phi(\mathbf{a})]_t$, and hence the check will pass, in spite of $\langle \mathbf{u} \rangle$ being reconstructed incorrectly. Yet, observe that this is an attack only if γ is non-zero, for which it must be the case that either δ_j or ϵ_j is not zero; say for simplicity $\delta_j \neq 0$. Fortunately, this will be caught in the check: the parties have the sharings $[\phi(\mathbf{x} + \gamma \mathbf{e}_i) + \phi(\mathbf{a})]_t$, which are incorrect in position i but, crucially, are correct in position j , so the adversary will not be able to conceal the fact that the j -th entry was modified.

Determining $[\phi(\mathbf{x})]_t$ for every input group. The principle above applies more generally. After performing optimistic computation using additive secret-sharing, the parties perform the following for every multiplication group with inputs \mathbf{x}, \mathbf{y} . Let $([\phi(\mathbf{a})]_t, [\phi(\mathbf{b})]_t, [\phi(\mathbf{c})]_t)$ be the Shamir triple associated to the group, and recall that the parties reconstructed (potentially incorrectly) \mathbf{u} and \mathbf{v} as part of the optimistic multiplications. The parties compute locally $[\phi(\mathbf{x})]_t = \phi(\mathbf{u}) - [\phi(\mathbf{a})]_t$ and $[\phi(\mathbf{y})]_t = \phi(\mathbf{v}) - [\phi(\mathbf{b})]_t$, and then they compute the product $[\phi(\mathbf{x} \star \mathbf{y})]_t$ using a secure multiplication protocol followed by re-encoding, as illustrated earlier. At this point, for every set of outputs \mathbf{x} of a given group, the parties hold $[\phi(\mathbf{x})]_t$. Moreover, the following crucial property holds: if no cheating occurred up to (and including) the optimistic evaluation of the gate at index i , then $\mathbf{x}[i]$ holds the correct wire value.

Now, the parties apply network routing to map all the packed sharings of the output groups into packed sharings of input groups $[\phi(\mathbf{x})]_t, [\phi(\mathbf{y})]_t$ for every multiplication group with inputs \mathbf{x}, \mathbf{y} . Importantly, due to the property above, if no cheating has occurred prior to the computation of, say, the i -th gate in a group with inputs \mathbf{x}, \mathbf{y} , then we know that the sharings $[\phi(\mathbf{x})]_t, [\phi(\mathbf{y})]_t$ derived from the network routing satisfy that the i -th entries x_i and y_i are *correct*. This way, if cheating occurs for the first time in this gate, by reconstructing incorrect $u'_i = x_i + a_i + \delta_i$ and $v'_i = y_i + b_i + \epsilon_i$ (which may result in more errors in other entries of the vectors $[\phi(\mathbf{x})]_t, [\phi(\mathbf{y})]_t$, but only with indices $j > i$, not for $j \leq i$), the parties can use $[\phi(\mathbf{x})]_t, [\phi(\mathbf{y})]_t$ (together with the associated triple $([\phi(\mathbf{a})]_t, [\phi(\mathbf{b})]_t, [\phi(\mathbf{c})]_t)$) to check the correctness of the reconstructed u'_i and v'_i .

Dealing with Addition Gates. Recall we assumed for simplicity that the circuit did not have any addition gates. We briefly comment how the general case is handled. First, both addition and multiplication gates are grouped in sets of ℓ gates each (where the gates within each group are of the same type). The optimistic computation phase remains the same: the parties handle addition gates by simply adding their (additive) shares together, locally. For the verification step, however, we need the parties to communicate for every group of addition gates. To see why this is the case, consider a group of addition gates with inputs $\langle \mathbf{x} \rangle, \langle \mathbf{y} \rangle$, and imagine that every output of these gates is later each fed to a multiplication gate. The parties can of course add locally $\langle \mathbf{x} \rangle, \langle \mathbf{y} \rangle$, but recall that for the network routing phase to work, we need the parties to have *packed Shamir* sharings of the outputs, like $[\phi(\mathbf{x} + \mathbf{y})]_t$, but it is not clear how they can obtain these from $\langle \mathbf{x} \rangle, \langle \mathbf{y} \rangle$ alone.

For the case of multiplication gates, this was achieved with the help of the (packed) triple that was used for the product. For the case of addition gates, we will use a similar idea: we make use of an *additive triple* $([\phi(\mathbf{a})]_t, [\phi(\mathbf{b})]_t, [\phi(\mathbf{c})]_t)$, where $\mathbf{c} = \mathbf{a} + \mathbf{b}$, ask the parties to open $\mathbf{u} \leftarrow \langle \mathbf{x} + \mathbf{a} \rangle$ and $\mathbf{v} \leftarrow \langle \mathbf{y} + \mathbf{b} \rangle$, and compute $[\phi(\mathbf{x} + \mathbf{y})]_t = (\mathbf{u} + \mathbf{v}) - [\phi(\mathbf{c})]_t$. This can also be seen as adding an extra step that first converts $\langle \mathbf{x} \rangle$ to $[\phi(\mathbf{x})]_t$ using the “double sharing” $(\langle \mathbf{a} \rangle, [\phi(\mathbf{a})]_t)$ (and similarly for \mathbf{y}), and then add these sharings together.

On Guaranteed Output Delivery. Finally, we comment on how the protocol sketched here is extended to G.O.D., without blowing communication. We use the player elimination framework from [BTH08], in which the parties not only perform a check but, in case of failure, identify a so-called *semi-corrupted pair* in the process, which is a pair of parties that is guaranteed to contain at least one corrupted party. At this point, the pair can be safely removed from the computation (which preserves the $t < n/3$ ratio), and the computation can be restarted. Restarting the computation many times can cause communication to blow up by a large factor. To address this, the circuit is split into *segments* of certain size, and the check described here is performed at the end of each such segment, rather than at the end of the whole circuit. Setting segment sizes appropriately reduces the size of the repeated computations, which keeps communication within $O(n|C|)$. There are

subtle issues, like part of the output of a group of a given segment being fed to the same segment, while some other part is fed to a future segment. This is mostly inconvenient notation-wise, but it does not add heavy technical complications.

To give a more complete picture, it remains to describe more clearly how the parties can identify a semi-corrupted pair during our check. Recall that our verification consists, in essence, of opening sharings of the form $[\phi(\mathbf{z})]_t$, and comparing against a previously opened set of values \mathbf{z}' . How should the parties react in case some mismatch is found? The core idea is to pinpoint to the party who announced an incorrect (additive) sharing in a first place. The main challenge with this is that, even though the parties have a robust version of the underlying *secrets* \mathbf{z} , they do not necessarily have a robust version of the additive sharings that each party should have sent when reconstructing \mathbf{z} , so it is not obvious how to identify which party sent an incorrect additive share. Fortunately, as it turns out, it is indeed possible to derive robust sharings that somehow commit the parties to the additive shares they should send at reconstruction. For this we introduce a notion of *extended* additive sharings, which expands additive SS with the necessary information to check whether a party sent a correct additive share. We provide details in Sections 3.2 and C in the Appendix.

Remark 3 (On sharings of zero). We remark that our overview here is a simplified version of our actual protocol, which must use of several other ingredients not discussed here. One of these is that, in several places, the parties need to re-randomize certain sharings using shares of zero, which is crucial for, among different purposes, preventing leakage of sensitive information when reconstructing optimistically, as in [GLS19].

1.4 Outline of the Document

We begin by presenting several important preliminaries in Section 2. This is followed by Section 3, which contains our protocols for optimistically evaluating a segment (Section 3.2), as well as verifying the computation is performed correctly (Section 3.5). This includes the network routing needed to compute robust sharings of groups in the circuit (Section 3.4), as well as our method to identify semi-corrupted parties once an attack has been detected (Section 3.6). We also discuss in detail how the circuit is partitioned into groups and segments (Section 3.1). Section 4 uses the building blocks from the previous sections to present our main MPC perfectly secure protocol with G.O.D. over constant-size rings, with linear communication complexity.

Our work makes use of several functionalities and protocols. To help the reader navigate, we provide in Section G in the Appendix a list with all of our functionalities and protocols, and their location within the text.

2 Preliminaries

In this work, we focus on functions that can be written as an arithmetic circuit C over the ring $\mathbb{Z}/2^k\mathbb{Z}$ with input, addition, multiplication and output gates. Let $|C|$ denote the size of the circuit C . We will make use of the *client-server* model for secure multiparty computation, in which clients can provide inputs and receive outputs to/from the servers, who are the parties who execute the actual MPC protocol. Note that, if every party plays a single client and a single server, this corresponds to a protocol in the standard MPC model. We assume that every pair of parties, either client and/or server, is connected via a secure (private and authentic) synchronous channel. We measure communication complexity as the total number of bits sent via private channels.⁷

Let c denote the number of clients, n denote the number of servers, and t denote the upper bound of the number of corrupted servers. In this work we focus on the 1/3-corruption setting, *i.e.* $3t + 1 = n$. In this work, we design an MPC protocol where all clients and servers compute the functionality $\mathcal{F}_{\text{Main}}$ with perfect security. Our definition of perfect security is based on the standard simulation-based security which is shown in the work [Can00].

⁷ Since we consider constant-sized rings, this is asymptotically the same as measuring the number of ring elements.

Functionality 1: $\mathcal{F}_{\text{Main}}(C)$

1. Let x denote the input and C denote the circuit. $\mathcal{F}_{\text{Main}}$ receives the input from all clients.
2. $\mathcal{F}_{\text{Main}}$ computes $C(x)$ and distributes the output to all clients.

2.1 Party-Elimination Framework

We make use of the party elimination framework by Hirt, Maurer, and Przydatek [HMP00], which constitutes a general strategy to achieve perfect security with G.O.D. with linear communication complexity. The basic idea is to let the parties perform checks that evaluate the correctness of the computation, identifying a pair of parties (with the help of BA for consensus) that contains at least one corrupted party in case of failure; such pair is referred to as a *semi-corrupted pair*. This pair of parties is then *eliminated* (i.e. removed from the computation), and the protocol is restarted. To avoid the overhead of re-executing as many times as potential eliminated pairs—which is upper bounded by t —the computation is divided into segments, and the check is performed at the end of each segment. This way, the extra cost of re-running is—in the worst case— t times the cost of each segment, so by keeping segments of appropriate size one can obtain efficient protocols with G.O.D.

We use $\mathcal{P}_{\text{active}}$ to denote the set of parties which are active in the current segment, that is, that have not been eliminated. We use $\mathcal{C}_{\text{active}} \subset \mathcal{P}_{\text{active}}$ for the set of active corrupted parties. Let n' be the size of $\mathcal{P}_{\text{active}}$. We use t' for the maximum possible number of the corrupted parties in $\mathcal{P}_{\text{active}}$. Each time a semi-corrupted pair is identified, these two parties are removed from $\mathcal{P}_{\text{active}}$ and hence $\mathcal{C}_{\text{active}}$. It results in $n' := n' - 2$ and $t' := t' - 1$. Initially we have $n = n'$, $t = t'$. Let $T = n' - 2t'$. Therefore, T remains unchanged during the whole protocol.

2.2 Finite Rings

Basic notation. Let \mathbb{Z} denote the ring of integers. For $q \in \mathbb{Z}$, let $q\mathbb{Z}$ denote the ideal $\{q \cdot n : q \in \mathbb{Z}\}$ and let $\mathbb{Z}/q\mathbb{Z}$ denote the quotient ring, which is the ring of integers modulo q . For a ring \mathcal{S} , let $\mathcal{S}[X]$ denote the ring of polynomials in the variable X with coefficients in \mathcal{S} . Also, let \mathcal{S}^* denote the multiplicative subgroup of invertible elements in \mathcal{S} .

Galois Rings. We adopt the notion of *Galois rings* that contains the quotient ring $\mathbb{Z}/2^k\mathbb{Z}$ from [ACD⁺19].

Definition 1 (Galois Ring [ACD⁺19]). A degree- d Galois ring of $\mathbb{Z}/2^k\mathbb{Z}$ is a ring of the form

$$(\mathbb{Z}/2^k\mathbb{Z})[X]/g(X),$$

where h is a positive integer, and $g(X) \in (\mathbb{Z}/2^k\mathbb{Z})[X]$ is a non-constant degree- d polynomial such that its reduction modulo 2 is an irreducible polynomial in the field $\mathbb{F}_2[X]$. We use $\text{GR}(2^k, d)$ to denote degree- d Galois ring of $\mathbb{Z}/2^k\mathbb{Z}$.

In order to interpolate polynomials in a Galois ring, we rely on the following lemma.

Lemma 1 ([ACD⁺19]). Let $\text{GR}(2^k, d)$ be a Galois ring with degree d . There exists a length 2^d sequence of distinct elements in $\text{GR}(2^k, d)$ denoted by $\alpha_1, \dots, \alpha_{2^d}$, such that for any $x_1, \dots, x_{2^d} \in \text{GR}(2^k, d)$, there exists a unique interpolating polynomial of degree at most $(2^d - 1)$ such that $f(\alpha_i) = x_i$ for all $i \in \{1, 2, \dots, d\}$.

Using this lemma, we can define necessary components such as Shamir secret sharings and hyper-invertible matrices over Galois rings. In the following, we will use a Galois ring of $\mathbb{Z}/2^k\mathbb{Z}$ denoted by $\mathcal{R} := \text{GR}(2^k, m)$. Note that the size of \mathcal{R} is $2^{m \cdot k}$. We select m such that $2^m \geq 2n + 1$ so that it is possible to interpolate degree- $2n$ polynomials in \mathcal{R} .

2.3 Secret Sharing Schemes

Shamir Secret Sharing. We will use the standard Shamir secret sharing scheme [Sha79] in this work. For the Galois ring $\mathcal{R} = \text{GR}(2^k, m)$, suppose $(\alpha_i)_{i=1}^n, \beta$ are $n + 1$ distinct points, which can be used to interpolate polynomials according to Lemma 1. A degree- d Shamir sharing of $x \in \mathcal{R}$ among $n' \leq n$ parties is a vector $(s_1, s_2, \dots, s_{n'}) \in \mathcal{R}^{n'}$ that satisfies the property that there exists a polynomial $f(\cdot) \in \mathcal{R}^{\leq d}[x]$ with $f(\beta) = x$ and $f(\alpha_i) = s_i, \forall i \in [n']$. The share held by party P_i is s_i . With any $(d + 1)$ different shares of the same sharing the secret x can be reconstructed.

A degree- d Shamir sharing of $x \in \mathcal{R}$ is denoted as $[x]_d$. The following two properties hold for Shamir sharings:

- For all $x, y \in \mathcal{R}$, $[x + y]_d = [x]_d + [y]_d$.
- For all $x, y \in \mathcal{R}$ and for all d_1, d_2 subject to $d_1 + d_2 < n$, we have $[x \cdot y]_{d_1+d_2} = [x]_{d_1} \cdot [y]_{d_2}$.

Our protocol also rely heavily on the following property of Shamir sharings. Suppose after some party elimination steps we have n' parties where a maximum t' of them can be malicious.

Lemma 2 ([BTH08]). *Suppose n' parties share a degree- d Shamir sharing $[x]_d$, and at most t' of the shares may be incorrect.*

- If $t' < (n' - d)/2$, then $[x]_d$ is correctable after receiving all the shares, e.g. by Berlekamp-Welch Algorithm.
- If $t' < n' - d$, then whether $[x]_d$ is inconsistent is detectable after receiving all the shares. This is due to the fact that two different degree- d polynomials agree on at most d points.

2.4 Reverse Multiplication Friendly Embeddings

Definition 2 (RMFE over Ring [CRX21,ELXY23]). *Let ℓ, m, k be positive integers. Let $\mathcal{R} = \text{GR}(2^k, m)$ denote the degree- m Galois ring of $\mathbb{Z}/2^k\mathbb{Z}$. A pair of mappings $(\phi : (\mathbb{Z}/2^k\mathbb{Z})^\ell \rightarrow \mathcal{R}, \psi : \mathcal{R} \rightarrow (\mathbb{Z}/2^k\mathbb{Z})^\ell)$ is called an $(\ell, m)_{2^k}$ -reverse multiplication friendly embedding (RMFE) if, for all $\mathbf{x}, \mathbf{y} \in (\mathbb{Z}/2^k\mathbb{Z})^\ell$, it holds that*

$$\psi(\phi(\mathbf{x}) \cdot \phi(\mathbf{y})) = \mathbf{x} \star \mathbf{y}.$$

Without loss of generality we can assume that $\psi(\phi(\mathbf{1})) = \mathbf{1}$, which ensures $\psi(\phi(\mathbf{x})) = \mathbf{x}$ for all $\mathbf{x} \in (\mathbb{Z}/2^k\mathbb{Z})^\ell$.

Defining the $\mathbb{Z}/2^k\mathbb{Z}$ -Linear Map $\text{val}(\cdot)$ [PS21]. To compute the summation of all entries of $\psi(y)$ from $y \in \mathcal{R}$, we define an $\mathbb{Z}/2^k\mathbb{Z}$ -linear map $\text{val}(\cdot) : \mathcal{R} \rightarrow \mathbb{Z}/2^k\mathbb{Z}$ as follows:

- For an input y , suppose $\psi(y) = (y_1, y_2, \dots, y_\ell)$.
- $\text{val}(y)$ is defined to be $\sum_{i=1}^{\ell} y_i$.

Let \mathbf{e}_i be a vector in $(\mathbb{Z}/2^k\mathbb{Z})^\ell$ such that all entries are 0 except that the i -th entry is 1, and let \mathbf{x} be a vector in $(\mathbb{Z}/2^k\mathbb{Z})^\ell$ of which the i -th entry is x_i . According to the definition of RMFEs, we have $\mathbf{e}_i \star \mathbf{x} = \psi(\phi(\mathbf{e}_i) \cdot \phi(\mathbf{x}))$. Therefore, we can access x_i by computing $x_i = \text{val}(\phi(\mathbf{e}_i) \cdot \phi(\mathbf{x}))$.

Existence of Constant Rate RMFEs over ring $\mathbb{Z}/2^k\mathbb{Z}$ [ACE⁺21]. In [ACE⁺21] it has been shown that constant rate RMFEs exist, as summarized in Theorem 1.

Theorem 1. *There exists a family of constant rate $(\ell, m)_{2^k}$ -RMFE where $m = \Theta(\ell)$.*

In this work, we will use $(\ell, m)_{2^k}$ -RMFE such that $m = O(\log n)$ and $\ell = O(\log n)$. The Galois ring $\mathcal{R} = \text{GR}(2^k, m)$ satisfies $2^m \geq 2n + 1$.

2.5 Useful Building Blocks

Reconstructing Shamir Sharings. The functionality $\mathcal{F}_{\text{OpenPub}}$ takes N degree- d ($d \leq t$) Shamir secret sharings over \mathcal{R} as input, and it outputs the reconstructed secrets to all parties. We assume that for each input degree- d sharing, the shares of all active honest parties lie on a degree- d polynomial.⁸ The full description of $\mathcal{F}_{\text{OpenPub}}$ appears in Section A in the Appendix. An instantiation of this functionality for our ring case can be easily generalized from the field-case construction in [ACD⁺19], which has communication complexity of $O(N \cdot n + n^2)$ elements in \mathcal{R} (i.e. $O(N \cdot n \cdot m + n^2 \cdot m)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$).

Secure Multiplication. The functionality $\mathcal{F}_{\text{Mult}}$ takes two tuples of N degree- t Shamir sharings over \mathcal{R} as input, which are denoted by $([x_1]_t, \dots, [x_N]_t)$ and $([y_1]_t, \dots, [y_N]_t)$. The output of $\mathcal{F}_{\text{Mult}}$ is the tuple of degree- t Shamir sharings of the results $([x_1 \cdot y_1]_t, \dots, [x_N \cdot y_N]_t)$. We assume that for each input degree- t Shamir sharing, the shares of all active honest parties lie on a degree- t polynomial. The full description of $\mathcal{F}_{\text{Mult}}$ appears in Section A in the Appendix. An instantiation of this functionality for our ring case can be easily generalized from the field-case construction in [BTH08]. Also, another instantiation of this functionality is implied in [ACD⁺19]. The protocol generalized from [BTH08] has communication complexity of $O(N \cdot n + n^3)$ elements in \mathcal{R} (i.e. $O(N \cdot n \cdot m + n^3 \cdot m)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$).

Performing Re-Encode. In our construction, we will need to transform a degree- t Shamir sharing over \mathcal{R} from $[x]_t$ to $[\phi \circ \psi(x)]_t$ in order to evaluate multiplication gates in the circuit. This process is called *re-encode*. The functionality $\mathcal{F}_{\text{ReEncode}}$ takes N degree- t Shamir sharings over \mathcal{R} as input, which are denoted by $[x_1]_t, \dots, [x_N]_t$. The output of $\mathcal{F}_{\text{ReEncode}}$ are N degree- t Shamir sharings of the re-encoded result $[\phi \circ \psi(x_1)]_t, \dots, [\phi \circ \psi(x_N)]_t$. We assume that for each input degree- t Shamir sharing, the shares of all active honest parties lie on a degree- t polynomial. The full description of $\mathcal{F}_{\text{ReEncode}}$ appears in Section A in the Appendix. An instantiation of this functionality for our ring case can be easily generalized from the field-case construction in [CCXY18]. The instantiation has communication complexity of $O(N \cdot n + n^3)$ elements in \mathcal{R} (i.e. $O(N \cdot n \cdot m + n^3 \cdot m)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$).

Verifying Consistency of Unreliable Broadcast Values. In our protocol, we will ask a dealer D to distribute several values that are supposed to be all the same, towards all parties. We do this over point-to-point channel to save the communication. The functionality $\mathcal{F}_{\text{VerifyBC}}$ receives from all parties N such unreliable broadcast values dealt by a dealer D , and verifies whether all parties indeed received the same values. The output of this functionality to each party is either **consistent** or **(inconsistent, E)**, where E is a semi-corrupted pair of parties. The full description of $\mathcal{F}_{\text{VerifyBC}}$ appears in Section A in the Appendix. An instantiation of this functionality for our ring case can be easily generalized from the field-case construction in [BTH08], which has communication complexity of $O(N \cdot n + n^2)$ elements in \mathcal{R} (i.e. $O(N \cdot n \cdot m + n^2 \cdot m)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$).

Input Gates for Shamir Sharings. We introduce the functionality $\mathcal{F}_{\text{InputShamir}}$, where a client C with N inputs in \mathcal{R} shares its inputs to the active parties using Shamir secret sharing. The full description of $\mathcal{F}_{\text{InputShamir}}$ appears in Section A in the Appendix. An instantiation of this functionality for our ring case can be easily generalized from the field-case construction in [BTH08]. The instantiation has a communication complexity of $O(N \cdot n + n^3)$ elements in \mathcal{R} (i.e. $O(N \cdot n \cdot m + n^3 \cdot m)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$).

2.6 Preparing Correlated Randomness

Our protocol relies on different forms of correlated randomness shared by all parties, and these are prepared independently of the inputs of the clients. We give a brief description of the correlations we require below.

⁸ If this is not the case, we ask the functionality to send the active honest parties' inputs to the adversary and allow the adversary to decide the output of active honest parties. Essentially, we give up the security if the shares of active honest parties do not lie on degree- d polynomials.

Random Shamir Sharings. The functionality $\mathcal{F}_{\text{RandShamir}}$ enables all parties to prepare N random degree- t Shamir sharings in the form of $[\phi(\mathbf{r})]_t$, where \mathbf{r} is a random vector in $(\mathbb{Z}/2^k\mathbb{Z})^\ell$. The description and the instantiation of $\mathcal{F}_{\text{RandShamir}}$ can be found in Section B.3 in the Appendix. The total communication complexity for the instantiation of $\mathcal{F}_{\text{RandShamir}}$ to generate N random Shamir sharings is $O(N \cdot n \cdot m + n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

Random Zero Additive Sharings. Once Beaver triples are prepared in the preprocessing phase, parties only need to do reconstructions in the online phase. To protect the shares held by honest parties, for each reconstruction, we will prepare a random additive sharing of 0 among the first $t+1$ parties. The functionality $\mathcal{F}_{\text{RandZeroAdditive}}$ enables all parties to prepare N random zero additive sharings. The description and the instantiation of $\mathcal{F}_{\text{RandZeroAdd}}$ can be found in Section B.3 in the Appendix. The total communication complexity for the instantiation of $\mathcal{F}_{\text{RandZeroAdd}}$ to generate N zero additive sharings is $O(N \cdot n \cdot m + n^3 \cdot m)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

Random Parity Sharings. For an element $p \in \mathcal{R}$, we say that p is *parity element* if $\text{val}(p) = 0$, and a parity sharing is a degree- t Shamir sharing of a parity element. When localizing a fault within a circuit segment, uniformly random parity sharings will be used as masks so that it is possible to check the correctness of the reconstruction. The functionality $\mathcal{F}_{\text{RandParity}}$ enables all parties to prepare N random parity sharings. The description and instantiation of $\mathcal{F}_{\text{RandParity}}$ can be found in Section B.3 in the Appendix. The total communication complexity for the instantiation of $\mathcal{F}_{\text{RandParity}}$ to generate N random parity sharings is $O(N \cdot n \cdot m + n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

Beaver Triples. To evaluate addition gates and multiplication gates, all parties will prepare Beaver triples in the form of $([\phi(\mathbf{a})]_t, [\phi(\mathbf{b})]_t, [\phi(\mathbf{c})]_t)$, where $\mathbf{a} + \mathbf{b} = \mathbf{c}$ when the Beaver triple is additive and $\mathbf{a} \star \mathbf{b} = \mathbf{c}$ when the Beaver triple is multiplicative. We introduce two functionalities $\mathcal{F}_{\text{TripleAdd}}$ and $\mathcal{F}_{\text{TripleMult}}$. $\mathcal{F}_{\text{TripleAdd}}$ enables all parties to prepare N random additive Beaver triples, and $\mathcal{F}_{\text{TripleMult}}$ enables all parties to prepare N random multiplicative Beaver triples. The descriptions and the instantiations of both $\mathcal{F}_{\text{TripleAdd}}$ and $\mathcal{F}_{\text{TripleMult}}$ can be found in Section B.3 in the Appendix. The total communication complexity for the instantiation of either $\mathcal{F}_{\text{TripleAdd}}$ or $\mathcal{F}_{\text{TripleMult}}$ to generate N Beaver Triples is $O(N \cdot n \cdot m + n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

3 Segment Evaluation and Verification

Our protocol first splits the circuit into *segments*, and then assigns gates of each type within a segment into *gate groups*. This is discussed in Section 3.1. Then, in Section 3.2 we show how the parties evaluate the gates of a given circuit *optimistically*, that is, without checking that the computation was carried out correctly. The verification is discussed in Section 3.3

3.1 Groups and Segments

We need to split the circuit into groups in certain way in order to apply packing effectively. We assume that the circuit C satisfies the following conditions:

1. Circuit Segment Conditions:

- Because C is a Directed Acyclic Graph (DAG), there exists a topological ordering among all addition and multiplication gates. We require that each segment consists of addition and multiplication gates whose topological orders are consecutive.
- The size of each segment should be $\Theta(n^2 \cdot m^2)$, except that the size of the last segment can be $O(n^2 \cdot m^2)$.

2. Gate Number Conditions:

- In the input and output layers, the number of input gates belonging to each client and the number of output gates belonging to each client are multiples of ℓ .
- The number of addition and multiplication gates within each circuit segment are multiples of ℓ .

3. Gate Grouping Conditions:

- During the computation, gates that have the same type (i.e., input gates belonging to the same client, output gates belonging to the same client, multiplication gates in the same circuit segment, addition gates in the same circuit segment) are organized into gate groups of size ℓ .
- For the output wires of each gate group, the number of times that those wires are used as input wires in other gates is a multiple of ℓ .

In Section E.1 in the Appendix we show that, if C does not satisfy these properties, then it can be transformed into a circuit C' that does satisfy the properties without affecting our linear communication claim. Based on the conditions above, we can split each segment into *gate groups* consisting of either ℓ multiplication gates (in which case the group is a *multiplication group*), or ℓ addition gates (in which case the group is an *addition group*). A set of ℓ wires corresponding to the left or right inputs of a given gate group is referred to as an *input group*, and *output groups* are defined similarly, but with output wires.

3.2 Segment Evaluation

The focus of this section is to show how the parties can evaluate optimistically a given segment \mathbf{seg} . The overall idea is to use additive secret-sharing with multiplication triples derived from packed triples. However, we consider an “enhanced” version of additive secret-sharing that allows for fault detection in case of cheating. This is described below.

Extended additive sharings. Let (ϕ, ψ) be an $(\ell, m)_2$ -RMFE. Recall that n denotes the number of parties and $\phi : (\mathbb{Z}/2^k\mathbb{Z})^\ell \rightarrow \mathcal{R}$ is an $\mathbb{Z}/2^k\mathbb{Z}$ -linear map. Also, $\mathcal{R} = \text{GR}(2^k, m)$ such that $2^m \geq 2n + 1$. Therefore, Shamir secret sharing is well-defined in \mathcal{R} . In our construction, we will use ϕ to encode a vector of secrets $\mathbf{x} = (x^{(1)}, x^{(2)}, \dots, x^{(\ell)}) \in (\mathbb{Z}/2^k\mathbb{Z})^\ell$. All parties will hold a degree- t Shamir sharing of $\phi(\mathbf{x})$, denoted by $[\phi(\mathbf{x})]_t$. For $x \in \mathbb{Z}/2^k\mathbb{Z}$, we use $\langle x \rangle$ to denote an *additive sharing* of x among the first $t + 1$ parties in $\mathbb{Z}/2^k\mathbb{Z}$. Specifically, the additive sharing of x is $\langle x \rangle = (x_1, \dots, x_{n'})$ where party P_i holds the share $x_i \in \mathbb{Z}/2^k\mathbb{Z}$ such that $\sum_{i=1}^{t+1} x_i = x$ and $x_{t+2}, \dots, x_{n'}$ are all 0. Recall that $\psi : \mathcal{R} \rightarrow (\mathbb{Z}/2^k\mathbb{Z})^\ell$ and $\text{val}(\cdot) : \mathcal{R} \rightarrow \mathbb{Z}/2^k\mathbb{Z}$ are both $\mathbb{Z}/2^k\mathbb{Z}$ -linear. The parties have *extended additive sharings* of $x \in \mathbb{Z}/2^k\mathbb{Z}$, denoted by $\langle\!\langle x \rangle\!\rangle$, if they have a degree- t Shamir sharing $[y]_t$ in \mathcal{R} such that $\text{val}(y) = x$. We write $\langle\!\langle x \rangle\!\rangle := [y]_t$. It is clear that these sharings are additive.

We note that we can derive extended additive sharings of $x \in \mathbb{Z}/2^k\mathbb{Z}$ from Shamir sharings $[\phi(\mathbf{z})]_t$, where the j -th element of \mathbf{z} is x . To see this, observe that, by the property of RMFE, we have that $\psi(\phi(\mathbf{e}_j) \cdot \phi(\mathbf{z})) = \mathbf{e}_j \star \mathbf{z}$. Therefore, $\text{val}(\phi(\mathbf{e}_j) \cdot \phi(\mathbf{z})) = x$. To obtain $\langle\!\langle x \rangle\!\rangle$, all parties locally compute $\langle\!\langle x \rangle\!\rangle = \phi(\mathbf{e}_j) \cdot [\mathbf{z}]_t$. In addition, it is easy to obtain sharings $\langle x \rangle$ from $\langle\!\langle x \rangle\!\rangle$ by using Lagrange coefficients; we give the details in Section C in the Appendix, where we describe at length the notion of extended sharings, together with their properties.

Optimistically evaluating a segment. For a circuit segment \mathbf{seg} , we use Protocol $\Pi_{\text{Eval}}(\mathbf{seg})$ to optimistically evaluate this segment. We evaluate its addition and multiplication gates using extended additive sharings and Beaver triples. A Beaver triple $([\phi(\mathbf{a})]_t, [\phi(\mathbf{b})]_t, [\phi(\mathbf{c})]_t)$ can be used to evaluate ℓ addition gates or ℓ multiplication gates. We assume that all parties have computed Shamir sharings of all gate group outputs of the previous circuit segments. This means that for an i -th element of any $[\phi(\mathbf{z})]_t$ used as a gate input in \mathbf{seg} , all parties can locally compute the extended additive sharing $\langle\!\langle x \rangle\!\rangle := \phi(\mathbf{e}_i) \cdot [\phi(\mathbf{z})]_t$.

For each gate with extended additive input sharings $\langle\!\langle x \rangle\!\rangle, \langle\!\langle y \rangle\!\rangle$, we will use the Beaver triple associated to the gate group containing this gate. Suppose the gate is the j -th gate within the gate group, and suppose $([\phi(\mathbf{a})]_t, [\phi(\mathbf{b})]_t, [\phi(\mathbf{c})]_t)$ is the Beaver triple corresponding to the gate group. All parties compute the extended additive sharings $\langle\!\langle a_j \rangle\!\rangle := \phi(\mathbf{e}_j) \cdot [\phi(\mathbf{a})]_t$, $\langle\!\langle b_j \rangle\!\rangle := \phi(\mathbf{e}_j) \cdot [\phi(\mathbf{b})]_t$ and $\langle\!\langle c_j \rangle\!\rangle := \phi(\mathbf{e}_j) \cdot [\phi(\mathbf{c})]_t$. Then all parties derive the additive sharings $\langle x \rangle, \langle y \rangle, \langle a_j \rangle, \langle b_j \rangle$ from the extended additive sharings using the method described previously.

The next step is reconstructing $\langle x \rangle + \langle a_j \rangle$ and $\langle y \rangle + \langle b_j \rangle$, for which a fixed dealer D will receive all shares, and then sends the reconstructed value to all parties. However, a subtle issue is that all parties must protect the redundancy in their sharings by preparing two random zero additive

sharing $\langle o_1 \rangle, \langle o_2 \rangle$, which can be done with Functionality $\mathcal{F}_{\text{RandZeroAdd}}$. Then, all parties send their shares of $\langle x \rangle + \langle a_j \rangle + \langle o_1 \rangle$ and $\langle y \rangle + \langle b_j \rangle + \langle o_2 \rangle$ to the dealer D , who reconstructs $u := x + a_j$ and $v := y + b_j$, and sends the result to all other parties.

If the gate is an addition gate, all parties locally compute the output extended additive sharing

$$\langle z \rangle := (u + v) \cdot \phi(\mathbf{e}_j) - \langle c_j \rangle.$$

If the gate is a multiplication gate, all parties locally compute the output extended additive sharing

$$\langle z \rangle := (u \cdot v) \cdot \phi(\mathbf{e}_j) - u \cdot \langle b_j \rangle - v \cdot \langle a_j \rangle + \langle c_j \rangle.$$

We describe Π_{Eval} in full detail below, and we show that the communication cost of Π_{Eval} is $O(n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

Protocol 1: $\Pi_{\text{Eval}}(\text{seg})$

1. Suppose seg has $n_{\text{add}} \cdot \ell$ addition gates and $n_{\text{mult}} \cdot \ell$ multiplication gates. All parties select the active party with the smallest index as the dealer of this segment seg . Let D denote the dealer.
2. All parties call $\mathcal{F}_{\text{RandZeroAdd}}$ to prepare $2 \cdot (n_{\text{add}} + n_{\text{mult}}) \cdot \ell$ random additive zero sharings. All parties also receive a set of eliminated parties denoted by S_1 . All parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_1$.
3. All parties call $\mathcal{F}_{\text{TripleAdd}}(n_{\text{add}} \cdot \ell)$ to generate the additive Beaver triples for the segments. All parties receive a set of eliminated parties denoted by S_2 . All parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_2$.
All parties call $\mathcal{F}_{\text{TripleMult}}(n_{\text{mult}} \cdot \ell)$ to generate the multiplicative Beaver triples for the segments. All parties receive a set of eliminated parties denoted by S_3 . All parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_3$.
4. All parties locally get all the extended additive sharings for the gate inputs of seg that are collected from previous layers. For the gate input connected to the i -th wire of the Shamir sharing $[\phi(\mathbf{z})]_t$, all parties locally derive the extended additive sharing by $\phi(\mathbf{e}_i) \cdot [\phi(\mathbf{z})]_t$.
5. All parties evaluate the multiplication gates and addition gates within seg according to topological ordering. For each gate with input extended additive sharings $\langle x \rangle$ and $\langle y \rangle$, we suppose it corresponds to the j -th entry of the Beaver triple $([\phi(\mathbf{a})]_t, [\phi(\mathbf{b})]_t, [\phi(\mathbf{c})]_t)$, denoted by (a_j, b_j, c_j) . All parties consume two unused random additive sharings prepared in Step 2, denoted by $\langle o_1 \rangle$ and $\langle o_2 \rangle$. Then all parties perform the following steps:
 - (a) All parties locally derive the extended additive sharing for a_j, b_j, c_j with $\langle a_j \rangle = \phi(\mathbf{e}_j) \cdot [\phi(\mathbf{a})]_t$, $\langle b_j \rangle = \phi(\mathbf{e}_j) \cdot [\phi(\mathbf{b})]_t$ and $\langle c_j \rangle = \phi(\mathbf{e}_j) \cdot [\phi(\mathbf{c})]_t$.
 - (b) All parties locally computes $\langle x \rangle + \langle a_j \rangle + \langle o_1 \rangle$ and $\langle y \rangle + \langle b_j \rangle + \langle o_2 \rangle$, and send their shares to D .
 - (c) D reconstructs $u := x + a_j$ and $v := y + b_j$. Then D sends u and v to all parties.
 - (d) If the gate is an addition gate, all parties locally compute the output extended additive sharing

$$\langle z \rangle := (u + v) \cdot \phi(\mathbf{e}_j) - \langle c_j \rangle.$$

If the gate is a multiplication gate, all parties locally compute the output extended additive sharing

$$\langle z \rangle := (u \cdot v) \cdot \phi(\mathbf{e}_j) - u \cdot \langle b_j \rangle - v \cdot \langle a_j \rangle + \langle c_j \rangle.$$

6. All parties output the eliminated set of parties $S := S_1 \cup S_2 \cup S_3$.

Cost of Π_{Eval} . Recall that each circuit segment has $O(n^2 \cdot m^2)$ gates, so we have $n_{\text{add}} = O(n^2 \cdot m^2 / \ell)$ and $n_{\text{mult}} = O(n^2 \cdot m^2 / \ell)$. It follows that the communication complexity of Step 2, Step 3, and Step 5 are all $O(n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$. Therefore, the communication complexity of Π_{Eval} is $O(n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

3.3 A First (Inefficient) Verification Protocol

For a given segment seg , once the parties have executed $\Pi_{\text{Eval}}(\text{seg})$, they have obtained extended additive sharings of every intermediate wire value of the circuit. However, a corrupted party may have cheated during this protocol, perhaps by sending incorrect sharings to the dealer D , and/or

D itself sends incorrect reconstructions to the parties. The purpose of this section is discussing how the parties can check whether such cheating indeed took place or not. The parties do this before they proceed to the next segment evaluation. Our strategy, as outlined in Section 1.3, is to get Shamir sharings of all input groups in \mathbf{seg} , which will be used to verify that the openings are done correctly. These sharings are derived from Shamir sharings of the output groups in the previous segments, and possibly of \mathbf{seg} itself. In order to achieve this, we must apply *network routing* techniques, which are developed in the works of [GPS21, GPS22].

We begin by presenting a version of our verification protocol that does *not* yet satisfy the linear communication complexity claim, but is structurally very close to our actual protocol while requiring little preliminaries on network routing for a clear understanding.

Network routing. Consider a segment \mathbf{seg} , and suppose that the parties have sharings $[\phi(\mathbf{z})]_t$ for every output group \mathbf{z} of each segment prior to \mathbf{seg} , and also of \mathbf{seg} itself. Now, let \mathbf{x} be an input group in \mathbf{seg} . Each entry x_i in \mathbf{x} is the output of a previous gate, which appears in an output group of either \mathbf{seg} , or a prior segment. Network routing is a set of techniques that enables the parties to obtain, from the sharings of all previous groups $[\phi(\mathbf{z})]_t$, sharings $[\phi(\mathbf{x})]_t$ of the input group \mathbf{x} in \mathbf{seg} . This is crucial for our verification protocol, and it is achieved by Protocol $\Pi_{\text{NetworkRouting}}(\mathbf{seg})$ (Protocol 2 in p. 17). The construction of $\Pi_{\text{NetworkRouting}}(\mathbf{seg})$ is a natural adaptation to the RMFE setting of the network routing techniques from [GPS21, GPS22], which are originally set in the packed secret-sharing context.

For our first verification protocol, we will use $\Pi_{\text{NetworkRouting}}$ as a “black-box”. Doing so results in a verification protocol with super-linear communication, stemming from the fact that all the calls to $\Pi_{\text{NetworkRouting}}(\mathbf{seg})$ across all segments \mathbf{seg} redo a lot of computation that can be done only once if one “opens the box”. Later in the section we dig into the details of $\Pi_{\text{NetworkRouting}}(\mathbf{seg})$, identifying these steps that are repeated across calls so that they are called only once, avoiding unnecessary repetition and hence achieving the desired linear communication complexity.

Fault detection. We need to introduce one more tool before we present our first (inefficient) verification protocol. Our verification ultimately boils down to ensuring that extended secret-shared values, that have been opened non-robustly using additive shares through a dealer, have actually been opened correctly. If this is not the case, the parties should be able to identify a semi-corrupted pair. This is achieved by means of a protocol $\Pi_{\text{FaultDetection}}$ that takes as input an inconsistent pair of extended additive sharing $\langle s \rangle$ and the masked additive sharing $\langle s \rangle + \langle o \rangle$ corresponding to it. At a high level, in this protocol the dealer will open the shares of the extended additive sharing and the masked additive sharing to find out a dispute between parties. We refer the readers to Section 3.6 for more details.

Inefficient verification. Consider a segment \mathbf{seg} . Let $([\phi(\mathbf{a}^{(i)})]_t, [\phi(\mathbf{b}^{(i)})]_t, [\phi(\mathbf{c}^{(i)})]_t)$ denote the Beaver triple corresponding to the i -th gate group in \mathbf{seg} , and let $\mathbf{x}^{(i)}$ and $\mathbf{y}^{(i)}$ denote the left and right inputs of the gate group. Note that before the verification of \mathbf{seg} , each party locally holds the values $d_j^{(i)} = x_j^{(i)} + a_j^{(i)}$ and $e_j^{(i)} = y_j^{(i)} + b_j^{(i)}$ sent by D for $i \in [N], j \in [\ell]$. All parties first check the consistency of the values sent by D by calling $\mathcal{F}_{\text{VerifyBC}}$. If the values are consistent, all parties can get “temporary” input group sharings $[\phi(\mathbf{x})]_t$ and $[\phi(\mathbf{y})]_t$ from with $\mathbf{d}^{(i)}$, $\mathbf{e}^{(i)}$ and the Beaver triple $([\phi(\mathbf{a}^{(i)})]_t, [\phi(\mathbf{b}^{(i)})]_t, [\phi(\mathbf{c}^{(i)})]_t)$, and they use $\mathcal{F}_{\text{Mult}}$ and $\mathcal{F}_{\text{ReEncode}}$ to obtain $[\phi(\mathbf{x} \star \mathbf{y})]_t$.

After getting all output Shamir sharings in \mathbf{seg} , the parties use $\Pi_{\text{NetworkRouting}}$ to obtain Shamir sharings of all the input groups $[\phi(\mathbf{x})]_t$ and $[\phi(\mathbf{y})]_t$. The parties robustly reconstruct $[\phi(\mathbf{x} + \mathbf{a})]_t$ and $[\phi(\mathbf{y} + \mathbf{b})]_t$ using $\mathcal{F}_{\text{OpenPub}}$, and then they compare these outputs with the values reconstructed in the evaluation phase $\mathbf{d}^{(i)}$ and $\mathbf{e}^{(i)}$. If all entries are consistent, the evaluation of \mathbf{seg} is correct. Otherwise, all parties can locate the first inconsistent $x_j^{(i)} + a_j^{(i)}$ or $y_j^{(i)} + b_j^{(i)}$. Then they eliminate a set of semi-corrupted parties, and re-evaluate the circuit segment afterwards.

The steps of the verification protocol are the following. We reiterate that this does not have linear communication complexity, but only because of the calls to $\Pi_{\text{NetworkRouting}}$. This is addressed in Section 3.5 after we “open the box” of network routing in Section 3.4

1. Suppose \mathbf{seg} has N gate groups. For the i -th gate group within \mathbf{seg} that corresponds to the Beaver triple $([\phi(\mathbf{a}^{(i)})]_t, [\phi(\mathbf{b}^{(i)})]_t, [\phi(\mathbf{c}^{(i)})]_t)$ and has inputs $\mathbf{x}^{(i)}, \mathbf{y}^{(i)} \in (\mathbb{Z}/2^k\mathbb{Z})^\ell$. All parties locally compute $\phi(\mathbf{x}^{(i)} + \mathbf{a}^{(i)})$ and $\phi(\mathbf{y}^{(i)} + \mathbf{b}^{(i)})$ for all $i \in [N]$.

2. All parties call $\mathcal{F}_{\text{VerifyBC}}$ with inputs $\{\phi(\mathbf{x}^{(i)} + \mathbf{a}^{(i)})\}_{i \in [N]} \cup \{\phi(\mathbf{y}^{(i)} + \mathbf{b}^{(i)})\}_{i \in [N]}$ and D . If the result is a semi-honest pair $\{P_{j_1}, P_{j_2}\}$, all parties take it as output and halt. Otherwise, run the following steps.
3. All parties locally compute $[\phi(\mathbf{x}^{(i)})]_t = \phi(\mathbf{x}^{(i)} + \mathbf{a}^{(i)}) - [\phi(\mathbf{a}^{(i)})]_t$ and $[\phi(\mathbf{y}^{(i)})]_t = \phi(\mathbf{y}^{(i)} + \mathbf{b}^{(i)}) - [\phi(\mathbf{b}^{(i)})]_t$ for all $i \in [N]$.
4. For each additive gate group with input $[\phi(\mathbf{x}^{(i)})]_t$ and $[\phi(\mathbf{y}^{(i)})]_t$, all parties locally compute the output sharing

$$[\phi(\mathbf{z}^{(i)})]_t = [\phi(\mathbf{x}^{(i)})]_t + [\phi(\mathbf{y}^{(i)})]_t.$$

For all the multiplicative gate groups, suppose their indices form the set I_{mult} . All parties call $\mathcal{F}_{\text{Mult}}$ with input $([\phi(\mathbf{x}^{(i)})]_t)_{i \in I_{\text{mult}}}$ and $([\phi(\mathbf{y}^{(i)})]_t)_{i \in I_{\text{mult}}}$, and get output $([w^{(i)}]_t)_{i \in I_{\text{mult}}}$ and a set of eliminated parties denoted by S_1 . All parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_1$.

Then all the parties call the functionality $\mathcal{F}_{\text{ReEncode}}$ with input $([w^{(i)}]_t)_{i \in I_{\text{mult}}}$ and get the output sharings $([\phi(\mathbf{z}^{(i)})]_t)_{i \in I_{\text{mult}}}$ and a set of eliminated parties denoted by S_2 . All parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_2$.

5. All parties run the protocol $\Pi_{\text{NetworkRouting}}(\text{seg})$ to get all Shamir sharings of seg 's gate group inputs, denoted by $\{[\phi(\tilde{\mathbf{x}}^{(i)})]_t\}_{i \in [N]}$ and $\{[\phi(\tilde{\mathbf{y}}^{(i)})]_t\}_{i \in [N]}$. All parties also receive a set of eliminated parties denoted by S_3 . All parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_3$.
6. All parties call $\mathcal{F}_{\text{OpenPub}}$ to reconstruct the Shamir sharings $\{[\phi(\tilde{\mathbf{x}}^{(i)})]_t + [\phi(\mathbf{a}^{(i)})]_t\}_{i \in [N]}$ and $\{[\phi(\tilde{\mathbf{y}}^{(i)})]_t + [\phi(\mathbf{b}^{(i)})]_t\}_{i \in [N]}$, and get $\tilde{\mathbf{x}}^{(i)} + \mathbf{a}^{(i)}$ and $\tilde{\mathbf{y}}^{(i)} + \mathbf{b}^{(i)}$ for all $i \in [N]$.
7. Each party locally compares $\tilde{\mathbf{x}}^{(i)} + \mathbf{a}^{(i)}$ with $\mathbf{x}^{(i)} + \mathbf{a}^{(i)}$ and compares $\tilde{\mathbf{y}}^{(i)} + \mathbf{b}^{(i)}$ with $\mathbf{y}^{(i)} + \mathbf{b}^{(i)}$ for all $i \in [N]$. If there are any differences, all parties do the following: let $S := S_0 \cup S_1 \cup S_2 \cup S_3$. If $S \neq \emptyset$, all parties output S and **incorrect** and halts. Otherwise, all parties can select a wire with inconsistent value that has the smallest topological order, denoted by $x_{j_0}^{(i_0)} + a_{j_0}^{(i_0)}$ or $y_{j_0}^{(i_0)} + b_{j_0}^{(i_0)}$. Let $\langle s \rangle + \langle o \rangle$ denote its corresponding additive sharing for reconstruction in the protocol Π_{Eval} , and let $\langle s \rangle$ denote its extended additive sharing. Then all parties run the protocol $\Pi_{\text{FaultDetection}}(D, \langle s \rangle + \langle o \rangle, \langle s \rangle)$, and get a set of eliminated parties S_4 as output. All parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_4$, and output $S \cup S_4$ and **incorrect**.

3.4 Details on Network Routing

In this section we describe in detail how network routing works in order to identify the pieces that can be re-used from one call to $\Pi_{\text{NetworkRouting}}$ to the next, and then we present our actual verification protocol with linear communication in Section 3.5.

Fan-Out Operations. The first ingredient of network routing is a functionality $\mathcal{F}_{\text{FanOut}}$, which we describe in detail as Functionality 18 in Section E.4 in the Appendix. $\mathcal{F}_{\text{FanOut}}$ takes as input a list of sharings $[\phi(\mathbf{z}^{(1)})]_t, \dots, [\phi(\mathbf{z}^{(N)})]_t$, and also $n_j^{(i)} \in \mathbb{Z}^+$ for every $i \in [N]$ and $j \in [\ell]$ where ℓ divides $\sum_{j=1}^{\ell} n_j^{(i)}$. The functionality outputs, for each $i \in [N]$, sharings $[\phi(\mathbf{x}_j^{(i)})]_t$ for $j \in [m^{(i)}]$ where $m^{(i)} := \left(\sum_{j=1}^{\ell} n_j^{(i)}\right) / \ell$, such that each $z_j^{(i)}$ appears exactly $n_j^{(i)}$ times in $[\mathbf{x}_1^{(i)} \parallel \dots \parallel \mathbf{x}_{m^{(i)}}^{(i)}]_t$. Jumping ahead, fan-out will be used to copy each output wire as many times as it is used subsequently in the circuit.

The protocol Π_{FanOut} that implements $\mathcal{F}_{\text{FanOut}}$ can be found in Section E.4 in the Appendix, and its communication complexity of the protocol to generate a total of M fan-out sharings is $O(M \cdot n \cdot m + n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

Secret Collection. The second crucial operation that is needed in network routing is, given a series of Shamir sharings $([\phi(\mathbf{x}_1)]_t, \dots, [\phi(\mathbf{x}_N)]_t)$, obtain another set of sharings $([\phi(\mathbf{y}_1)]_t, \dots, [\phi(\mathbf{y}_N)]_t)$, where $(\mathbf{y}_1 \parallel \dots \parallel \mathbf{y}_N)$ is a permutation of $(\mathbf{x}_1 \parallel \dots \parallel \mathbf{x}_N)$. We refer to this operation as *secret collection*. To gain some intuition on how this helps in network routing, suppose that $\mathcal{F}_{\text{FanOut}}$ has been applied to all wires in the circuit, copying them as many times as they appear in future gates, and let $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ be the all the vectors output by $\mathcal{F}_{\text{FanOut}}$. Let $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_N)$ consist of all input groups to all circuit segments, and also to the output layer. We have then that \mathbf{y} is a *permutation* of \mathbf{x} . Using secret collection, the parties can obtain sharings of each input group $[\phi(\mathbf{y}_i)]_t$, which is precisely what is needed for network routing.

The following theorem from [GPS21] is useful for implementing this secret collection functionality.

Theorem 2 ([GPS21]). Suppose $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ and $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_N)$ satisfy $P \cdot \mathbf{x} = \mathbf{y}$ where P is a permutation matrix. There exists two sets of permutations p_1, \dots, p_N and q_1, \dots, q_N that permute vectors in $(\mathbb{Z}/2^k\mathbb{Z})^\ell$, such that after applying p_i to \mathbf{x}_i and q_j to \mathbf{y}_j for all $i, j \in [N]$, the following property holds for an arbitrary $q_h \cdot \mathbf{y}_h$:

- Suppose that $q_h \cdot \mathbf{y}_h = (y'_1, \dots, y'_\ell)$, and that y'_w corresponds to the $u_{h,w}$ -th entry in $p_{v_{h,w}} \cdot \mathbf{x}_{v_{h,w}}$, for $w \in [\ell]$. Then, $(u_{h,1}, \dots, u_{h,\ell})$ forms a permutation of $(1, 2, \dots, \ell)$

Following this theorem, we can obtain $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_N)$ by first permuting each $[\phi(\mathbf{x}_i)]_t$ as $[\phi(p_i \cdot \mathbf{x}_i)]_t$. Then, for each $h \in [N]$ parties compute locally $\sum_{w=1}^{\ell} \phi(\mathbf{e}_w) \cdot [\phi(p_{v_{h,w}} \cdot \mathbf{x}_{v_{h,w}})]_t$, which creates a vector whose w -th entry is the w -th entry of $p_{v_{h,w}} \cdot \mathbf{x}_{v_{h,w}}$. Thanks to the theorem, this is precisely $q_h \cdot \mathbf{y}_h$, so the parties can obtain $[\phi(q_h \cdot \mathbf{y}_h)]_t$ by applying $\mathcal{F}_{\text{ReEncode}}$. Finally, to obtain the desired $[\phi(\mathbf{y}_h)]_t$ for each $h \in [N]$, the parties can apply the inverse permutation q_h^{-1} of q_h to the sharing $[\phi(q_h \cdot \mathbf{y}_h)]_t$.

The permutations above are done with a functionality $\mathcal{F}_{\text{Permute}}$, which we define and instantiate with Protocol Π_{Permute} in Section E.4 in the Appendix, involving a communication complexity of $O(N \cdot n \cdot m + n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$. The aggregation $\sum_{w=1}^{\ell} \phi(\mathbf{e}_w) \cdot [\phi(p_{v_{h,w}} \cdot \mathbf{x}_{v_{h,w}})]_t$ followed by the re-encoding with $\mathcal{F}_{\text{ReEncode}}$ is abstracted as a functionality $\mathcal{F}_{\text{Collect}}$, which we implement with a protocol Π_{Collect} in Section E.4 in the Appendix. The communication complexity of Π_{Collect} is $O(N \cdot n \cdot m + n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

Network Routing for a Circuit Segment. Consider a segment \mathbf{seg} . We can finally describe Protocol $\Pi_{\text{NetworkRouting}}(\mathbf{seg})$, which computes sharings $[\phi(\mathbf{x})]_t$ for every input group \mathbf{x} of the segment \mathbf{seg} . In a nutshell, the protocol proceeds exactly as sketched above: (1) $\mathcal{F}_{\text{FanOut}}$ is used to copy wires, and (2) Secret collection is performed by calling $\mathcal{F}_{\text{Permute}} \rightarrow \mathcal{F}_{\text{Collect}} \rightarrow \mathcal{F}_{\text{Permute}}$. However, for our “non-black-box” optimization, suppose that $\Pi_{\text{NetworkRouting}}(\mathbf{seg}')$ was called for the segment \mathbf{seg}' that goes right before \mathbf{seg} . In this case, $\mathcal{F}_{\text{FanOut}}$ has been performed on all output groups prior to segment \mathbf{seg} , and therefore we only need to do it for these output groups in \mathbf{seg} . Even more, since $\Pi_{\text{NetworkRouting}}$ is called to verify the inputs of \mathbf{seg} , we only need to use $\mathcal{F}_{\text{FanOut}}$ in these output wires of \mathbf{seg} that are fed as inputs to \mathbf{seg} itself. Later, if the check passes, then we apply $\mathcal{F}_{\text{FanOut}}$ to the remaining wires. A similar optimization happens with the first call to $\mathcal{F}_{\text{Permute}}$, which is done on the output groups.

The description of $\Pi_{\text{NetworkRouting}}$ appears in Protocol 2. If $\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(N)}$ are the output groups of \mathbf{seg} , we denote by $\tilde{n}_j^{(i)}$ for $j \in [\ell], i \in [N]$ the number of times that wire $z_j^{(i)}$ is used *inside* \mathbf{seg} itself. Protocol $\Pi_{\text{NetworkRouting}}$ assumes that the protocol has been called for the previous segment, so these are the only remaining copies needed for getting the input groups of \mathbf{seg} . After this, the $\mathcal{F}_{\text{Permute}} \rightarrow \mathcal{F}_{\text{Collect}} \rightarrow \mathcal{F}_{\text{Permute}}$ sequence is applied.

Protocol 2: $\Pi_{\text{NetworkRouting}}(\mathbf{seg})$

1. All parties call $\mathcal{F}_{\text{FanOut}}$ on all the output sharings of \mathbf{seg} and $\{\tilde{n}_j^{(i)}\}_{j \in [\ell], i \in [N]}$, where the j -th wire of $\mathbf{z}^{(i)}$ is copied $\tilde{n}_j^{(i)}$ times (this is the number of times this wire is used in \mathbf{seg} itself). All parties receive the fan-out sharings which are used for this segment’s gate inputs, and a set of eliminated parties denoted by S_1 . All parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_1$.
2. All parties call $\mathcal{F}_{\text{Permute}}$ with the fan-out sharings and the desired permutations as input. All parties receive the permuted fan-out Shamir secret sharings, and a set of eliminated parties denoted by S_2 . All parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_2$.
3. All parties call $\mathcal{F}_{\text{Collect}}$ to get the collected Shamir sharings of \mathbf{seg} ’s gate group inputs. All parties also receive a set of eliminated parties denoted by S_3 . All parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_3$.
4. All parties call $\mathcal{F}_{\text{Permute}}$ with the collected Shamir sharings as input, and get all Shamir sharings of \mathbf{seg} ’s gate group inputs, denoted by $\{[\phi(\tilde{\mathbf{x}}^{(i)})]_t\}_{i \in [N]}$ and $\{[\phi(\tilde{\mathbf{y}}^{(i)})]_t\}_{i \in [N]}$. All parties also receive a set of eliminated parties denoted by S_4 .
5. All parties output $S := S_1 \cup S_2 \cup S_3 \cup S_4$ and all the input Shamir sharings $\{[\phi(\tilde{\mathbf{x}}^{(i)})]_t\}_{i \in [N]}$ and $\{[\phi(\tilde{\mathbf{y}}^{(i)})]_t\}_{i \in [N]}$.

3.5 Efficient Verification

Now we are ready to “patch” the verification protocol from Section 3.3 to get linear communication. The full protocol, Π_{Verify} , is given in Section D in the Appendix, and here we only discuss the core

differences with respect to the protocol from before. As before, we do call $\Pi_{\text{NetworkRouting}}$ in step 5, except that this time we take into account the fact that $\mathcal{F}_{\text{FanOut}}$ and $\mathcal{F}_{\text{Permute}}$ have been called for all previous segments, and hence only need to be computed for the current segment (as described in $\Pi_{\text{NetworkRouting}}$). Second, recall that the $\mathcal{F}_{\text{FanOut}}$ and $\mathcal{F}_{\text{Permute}}$ calls in $\Pi_{\text{NetworkRouting}}$ are only performed to the output groups needed for the input groups in the current segment seg . If step 7 succeeds, then the parties need to apply $\mathcal{F}_{\text{FanOut}}$ and $\mathcal{F}_{\text{Permute}}$ to the remaining output groups in seg in preparation to the call to $\Pi_{\text{NetworkRouting}}$ for the next segment.

Protocol Π_{Verify} is given in Section D in the Appendix. In that section we analyze its communication, verifying that it indeed grows linearly with n .

3.6 Fault Localization

Finally, we have focused so far in how the parties detect that cheating occurred, but we have not discussed how to react to that, identifying a semi-corrupted pair so that the segment can be re-run. Recall that a parity element $p \in \mathcal{R}$ is an element that satisfies $\text{val}(p) = 0$. To mask the shares of extended additive sharing $\langle s \rangle$ when the dealer D opens all shares, all parties will prepare a random degree- t shamir sharing of a random parity element, denoted by $[p]_t$. Then all parties will send their shares of $\langle s \rangle + [p]_t$ to D .

We note that it is enough for D to localize a semi-corrupted pair just by opening the two sharings $\langle s \rangle + [p]_t$ and $\langle s \rangle + \langle o \rangle$, so we introduce the following steps to let all parties disclose to D the randomness masks that they have distributed and received. We first observe that, due to the way all parties prepare random sharings (see Section B.2 in the Appendix), $\langle o \rangle$ can be written as $\langle o \rangle = \sum_{i=0}^{n'} \langle o_i \rangle$, and $[p]_t$ can be written as $[p]_t = \sum_{i=0}^{n'} [p_i]_t$, where $\langle o_i \rangle$ is the zero additive sharing distributed by P_i and $[p_i]_t$ is the parity sharing distributed by the party P_i . Each parity sharing $[p_i]_t$ corresponds to an additive sharing whose secret is 0, denoted by $\langle o'_i \rangle$. Let $\langle o' \rangle := \sum_{i=1}^{n'} \langle o'_i \rangle$. Note that D can locally compute $\langle o \rangle - \langle o' \rangle$, and $\langle o \rangle - \langle o' \rangle = \sum_{i=1}^{n'} \langle o_i \rangle - \langle o'_i \rangle$.

Following the idea in [BFO12], in order to protect the shares of $\langle o_i \rangle - \langle o'_i \rangle$ (which may leak information), each party P_i distributes another additive zero sharing $\langle o''_i \rangle$ as a mask. Then P_i reveals to D all the shares of $\langle o_i \rangle - \langle o'_i \rangle + \langle o''_i \rangle$, and also the share of $\langle o_j \rangle - \langle o'_j \rangle + \langle o''_j \rangle$ that P_i received from another party P_j . Given this information, D is able to identify disputes between parties. We summarize the full protocol $\Pi_{\text{FaultDetection}}$ as follows. Note that the dealer D has the shares of $\langle s \rangle + \langle o \rangle$ sent by all parties at the beginning of the protocol, and $\langle s \rangle$ is transformed from $\langle s \rangle$.

Protocol 3: $\Pi_{\text{FaultDetection}}(D, \langle s \rangle + \langle o \rangle, \langle s \rangle)$

1. Each party P_i prepares a random additive sharing of zero denoted by $\langle o''_i \rangle$, and distribute the shares of $\langle o''_i \rangle$ to all other parties. Then all parties locally compute $\langle o'' \rangle := \sum_{i=1}^{n'} \langle o''_i \rangle$.
2. All parties send their shares of $\langle o'' \rangle$ to D . Then D broadcasts whether $\langle o'' \rangle$ is a valid additive sharing of 0.
3. If D broadcasts that $\langle o'' \rangle$ is not valid, all parties set $\langle \hat{o} \rangle := \sum_{i=1}^{n'} \langle o''_i \rangle$ and let $\langle \hat{o}_i \rangle := \langle o''_i \rangle$ for all $i \in [n']$.
4. Otherwise, all parties call $\mathcal{F}_{\text{RandParity}}$ to prepare a random parity sharing $[p]_t$. If all parties receive a non-empty set of eliminated parties denoted by S , all parties take S as output and halt. Otherwise, all parties send their shares of $\langle s \rangle + [p]_t$ to D . We note that $[p]_t$ can be written as $[p]_t = \sum_{i=1}^{n'} [p_i]_t$, and $[p_i]_t$ is distributed by the party P_i . Let $\langle o' \rangle$ denote the additive sharing transformed from $[p]_t$, and let $\langle o'_i \rangle$ denote the additive sharing transformed from $[p_i]_t$. All parties set $\langle \hat{o} \rangle := \langle o \rangle - \langle o' \rangle + \langle o'' \rangle$ and let $\langle \hat{o}_i \rangle := \langle o_i \rangle - \langle o'_i \rangle + \langle o''_i \rangle$ for all $i \in [n']$.
5. Let $\hat{o}_i^{(j)}$ denote the share of $\langle \hat{o}_i \rangle$ sent by P_i to P_j . Let $(\hat{o}_j^{(i)})'$ denote the share of $\langle \hat{o}_j \rangle$ received by P_i from P_j . For all $i \in [n']$ and for all $j \in [n']$, P_i sends $\hat{o}_i^{(j)}$ and $(\hat{o}_j^{(i)})'$ to D .
6. Let $\hat{o}^{(j)}$ denote the share of $\langle \hat{o} \rangle$ held by P_j . If for some i , $\langle \hat{o}_i \rangle$ that P_i claims to have sent is not a valid additive sharing of 0, or if $\hat{o}^{(i)} \neq \sum_{j=1}^{n'} \hat{o}_j^{(i)}$, D broadcasts the message $(P_i, \text{incorrect})$. Otherwise, if for some j_1 and j_2 ($j_1 \neq j_2$), the share $\hat{o}_{j_1}^{(j_2)}$ that P_{j_1} claims to have sent to P_{j_2} is inconsistent to $(\hat{o}_{j_1}^{(j_2)})'$ that P_{j_2} claims to have received from P_{j_1} , D broadcasts the message $(P_{j_1}, P_{j_2}, \hat{o}_{j_1}^{(j_2)}, (\hat{o}_{j_1}^{(j_2)})', \text{inconsistent})$.

7. If $(P_i, \text{incorrect})$ is broadcast by D , all parties set the eliminating set $E := \{D, P_i\}$. Otherwise, if $(P_{j_1}, P_{j_2}, \hat{o}_{j_1}^{(j_2)}, (\hat{o}_{j_1}^{(j_2)})', \text{inconsistent})$ is broadcast by D , P_{j_1} and P_{j_2} will broadcast if they agree with this message. If P_{j_1} does not agree, all parties set the eliminating set $E := \{D, P_{j_1}\}$; otherwise if P_{j_2} does not agree, all parties set the eliminating set $E := \{D, P_{j_2}\}$; otherwise all parties set the eliminating set $E := \{P_{j_1}, P_{j_2}\}$. All parties take E as output.

Lemma 3 (Correctness of $\Pi_{\text{FaultDetection}}$). *If the secrets of $\langle s \rangle + \langle o \rangle$ and $\langle s \rangle + [p]_t$ are not consistent, and $\langle s \rangle$ is derived from $\langle s \rangle$, then Protocol 3 can find a semi-corrupted pair E that satisfies $E \cap \mathcal{C}_{\text{active}} \neq \emptyset$ when the protocol does not halt on Step 4.*

The proof of Lemma 3 can be found in Section F.1 in the Appendix.

Cost of $\Pi_{\text{FaultDetection}}$. The cost of Step 1 is $O(n^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$ for each party to share $\langle o_i'' \rangle$. The cost of Step 2 is $O(n)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$ for all parties to send their shares of $\langle o'' \rangle$ to D , and $O(n^2)$ bits for D to broadcast if $\langle o'' \rangle$ is valid. The cost of Step 4 is $O(n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$ for calling $\mathcal{F}_{\text{RandParity}}$. The cost of Step 5 is $O(n^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$ for each party to send $\langle \hat{o}_i \rangle$ that they sent and received to D . The cost of Step 6 is bounded by $O(n^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$ for broadcasting $\hat{o}_{j_1}^{(j_2)}, (\hat{o}_{j_1}^{(j_2)})'$, plus $O(n^2 \log n)$ bits for broadcasting P_i or P_{j_1}, P_{j_2} . The cost of Step 7 is bounded by $O(n^2)$ bits. Recall that $m = O(\log n)$. Therefore, the overall cost of this protocol is $O(n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

4 Main Protocol

In the previous section we saw how to evaluate each segment in the circuit, and how to check if the execution was correct, identifying a semi-corrupt pair if this was not the case. In this section, we show how to put together these protocols in order to evaluate the entire circuit with G.O.D., essentially by making use of the player elimination framework by Beerliová-Trubíniová and Hirt [BTH08], in which the parties that remain after a semi-corrupted pair is removed re-run the failed segment. This is described in Section 4.3. No MPC protocol would be complete without describing how input and output gates are handled. This is explained in Sections 4.1 and 4.2, respectively.

4.1 Input Gates

Since we are in the client-server model, all the inputs belong to the clients. Recall that we assume that the number of inputs for each client is a multiple of ℓ . In this part, we introduce a protocol Π_{Input} , which enables all client to share their inputs to all parties, and then properly performs fan-out and permutations to input sharings to prepare them for later use. We describe the protocol Π_{Input} below.

Protocol 4: Π_{Input}

1. For each client **Client**, suppose its inputs are $\{x_i^{(1)}, \dots, x_i^{(\ell)}\}_{i=1}^N$. All parties and **Client** perform the following steps:
 - (a) Let $\mathbf{x}_i := (x_i^{(1)}, \dots, x_i^{(\ell)}) \in (\mathbb{Z}/2^k\mathbb{Z})^\ell$. **Client** locally computes $\phi(\mathbf{x}_i)$ for all $i \in [N]$.
 - (b) All parties call $\mathcal{F}_{\text{InputShamir}}$ with the inputs $\phi(\mathbf{x}_1), \dots, \phi(\mathbf{x}_N)$ from **Client**. All parties get the output $[y_1]_t, \dots, [y_N]_t$ and a set of eliminated parties denoted by S_1 . All parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_1$.
 - (c) All parties call $\mathcal{F}_{\text{RandShamir}}(N)$ to prepare N random degree- t shamir secret sharings denoted by $[\phi(\mathbf{r}_1)]_t, \dots, [\phi(\mathbf{r}_N)]_t$. All parties also receive a set of eliminated parties denoted by S_2 . All parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_2$.
 - (d) All parties call $\mathcal{F}_{\text{OpenPub}}(N)$ to reconstruct $[y_1]_t + [\phi(\mathbf{r}_1)]_t, \dots, [y_N]_t + [\phi(\mathbf{r}_N)]_t$ and get the secrets $y_1 + \phi(\mathbf{r}_1), \dots, y_N + \phi(\mathbf{r}_N)$.
 - (e) For all $i \in [N]$, All parties locally check if $\phi \circ \psi(y_i + \phi(\mathbf{r}_i)) = y_i + \phi(\mathbf{r}_i)$. If the equation does not hold, all parties set $[y_i]_t$ to $[0]_t$. Note that after this step, each sharing $[y_i]_t$ can be written as $[y_i]_t = [\phi(\mathbf{x}'_i)]_t$, where $\mathbf{x}'_i = \psi(y_i)$.
 - (f) All parties take $[\phi(\mathbf{x}'_1)]_t, \dots, [\phi(\mathbf{x}'_N)]_t$ as the shared inputs of **Client**.

2. All parties call $\mathcal{F}_{\text{FanOut}}$ with the input sharings from all clients, and get the set of eliminated parties denoted by S_3 and the resulting fan-out sharings. All parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_3$.
3. All parties call $\mathcal{F}_{\text{Permute}}$ to permute all the fan-out sharings with the desired permutations. All parties get the set of eliminated parties denoted by S_4 , and the resulting permuted sharings. All parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_4$.
4. All parties output the resulting permuted sharings in the previous step.

Cost of Π_{Input} . To get input from a client with $N \cdot \ell$ inputs in $\mathbb{Z}/2^k\mathbb{Z}$, the communication complexity is $O(N \cdot n \cdot m + n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$. The communication complexity of Step 2 is $O(M \cdot n \cdot m + n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$, and the communication complexity of Step 3 is $O(M \cdot n \cdot m + n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$. Since M is bounded by $O(|C'|/\ell)$, and $m, \ell = O(\log n)$, the total communication complexity of Π_{Input} is $O(|C'| \cdot n + c \cdot n^3 \cdot m^2 + n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

4.2 Output Gates

We introduce the protocol Π_{Output} that reconstructs all outputs towards all clients. In this protocol, all parties first performs network routing to generate the output Shamir sharings, and then they send the shares of the output sharing to each client. The description of Π_{Output} appears below.

Protocol 5: Π_{Output}

1. All parties call $\mathcal{F}_{\text{Collect}}$ to collect secrets for Shamir sharings of the output layer, and all parties receive a set of eliminated parties denoted by S_1 . All parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_1$.
2. All parties call $\mathcal{F}_{\text{Permute}}$ with the collected Shamir sharings and the desired permutations as input. All parties get all output Shamir sharings as output, and get a set of eliminated parties denoted by S_2 . All parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_2$.
3. For each client Client that has output Shamir sharings $[\phi(\mathbf{z}_1)]_t, \dots, [\phi(\mathbf{z}_N)]_t$, all parties send their shares of $[\phi(\mathbf{z}_i)]_t$ to Client for all $i \in [N]$. Then Client reconstructs the secrets $\mathbf{z}_1, \dots, \mathbf{z}_N$.

Cost of Π_{Output} . In Step 1, the communication complexity is bounded by calling $\mathcal{F}_{\text{Collect}}$ that outputs $|C'|/\ell$ sharings, so the communication complexity is bounded by $O((|C'| + \ell \cdot c) \cdot n \cdot m/\ell + n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$. In step 2, the communication complexity is bounded by $O(n \cdot |C'|) = O(|C'| \cdot n + n \cdot c \cdot \ell)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$. Since m, ℓ are both $O(\log n)$, the total communication complexity of Π_{Output} is $O(|C'| \cdot n + c \cdot n \cdot m + n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

4.3 Main Protocol

Given the above protocols, the main protocol that implements the ideal functionality $\mathcal{F}_{\text{Main}}$ follows in a straightforward way. The main protocol is introduced in Π_{Main} .

Protocol 6: $\Pi_{\text{Main}}(C)$

1. Let C denote the circuit. All parties transform C to C' . All parties agree on the gate grouping, and they order of the circuit segments according to topological ordering. All parties set $\mathcal{P}_{\text{active}} := \mathcal{P}$.
2. All parties run the protocol Π_{Input} .
3. All parties evaluate the circuit segments according to their ordering. For each circuit segment denoted by seg :
 - (a) All parties run the protocol $\Pi_{\text{Eval}}(\text{seg})$. All parties get the set of eliminated parties denoted by S .
 - (b) All parties run the protocol $\Pi_{\text{Verify}}(\text{seg}, S)$. If the output is **incorrect**, all parties repeat step 3.(a) and 3.(b) to evaluate seg . Otherwise, all parties continue to evaluate the next circuit segment.
4. All parties run the protocol Π_{Output} .

Cost of Π_{Main} . In Step 3, each time Π_{Eval} or Π_{Verify} outputs a semi-honest pair of parties, at least one corrupted party is removed from $\mathcal{P}_{\text{active}}$, so Π_{Eval} and Π_{Verify} will be repeated at most $t = O(n)$ times. Note that each time Π_{Eval} or Π_{Verify} is repeated, the communication complexity is $O(n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$. Overall, this introduces an overhead of $O(n^4 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$ to the total communication complexity.

In Step 3, when no circuit segment is repeated, the total communication complexity of Π_{Eval} and Π_{Verify} is

$$\sum_{\text{seg}} (O(n^3 \cdot m^2) + O(M_{\text{seg}} \cdot n \cdot m)) = O(|C'| \cdot n + n^3 \cdot m^2 + (\sum_{\text{seg}} M_{\text{seg}}) \cdot n \cdot m).$$

We note that $\sum_{\text{seg}} M_{\text{seg}}$ is bounded by $O(|C'|/\ell)$. So the total communication complexity of Step 3 when no circuit segment is repeated is $O((|C'|/\ell) \cdot n \cdot m + |C'| \cdot n + n^3 \cdot m^2)$.

Recall that $m, \ell = O(\log n)$. Therefore, the overall communication complexity of Π_{Main} is

$$O(|C| \cdot n + c \cdot n^3 \cdot \log^2 n + n^4 \cdot \log^2 n)$$

elements in $\mathbb{Z}/2^k\mathbb{Z}$. Here $|C|$ is the size of the circuit in *both* addition and multiplication gates. Notice that this is *linear*, as desired.

Lemma 4. *Protocol Π_{Main} securely computes $\mathcal{F}_{\text{Main}}$ in the $(\mathcal{F}_{\text{InputShamir}}, \mathcal{F}_{\text{RandShamir}}, \mathcal{F}_{\text{RandZeroAdd}}, \mathcal{F}_{\text{RandParity}}, \mathcal{F}_{\text{OpenPub}}, \mathcal{F}_{\text{ReEncode}}, \mathcal{F}_{\text{Mult}}, \mathcal{F}_{\text{TripleAdd}}, \mathcal{F}_{\text{TripleMult}}, \mathcal{F}_{\text{VerifyBC}}, \mathcal{F}_{\text{FanOut}}, \mathcal{F}_{\text{Permute}}, \mathcal{F}_{\text{Collect}})$ -hybrid model with perfect security against a fully malicious adversary who controls $t < n/3$ parties.*

The proof of Lemma 4 can be found in Section F.3 in the Appendix. This leads to the following Theorem, which is the main result of our work.

Theorem 3. *In the client-server model, let c denote the number of clients, and $n = 3t + 1$ denote the number of parties (servers). Let k be a constant positive integer and let $\mathbb{Z}/2^k\mathbb{Z}$ be a finite ring of constant size. For an arithmetic circuit C over $\mathbb{Z}/2^k\mathbb{Z}$, let $|C|$ denote the size of the circuit. There exists an information-theoretic MPC protocol which securely computes the arithmetic circuit with perfect security in the presence of a fully malicious adversary controlling up to c clients and t parties. The communication complexity of this protocol is $O(|C| \cdot n + \text{poly}(c, n))$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.*

Acknowledgments

This paper was prepared in part for information purposes by the Artificial Intelligence Research group of JPMorgan Chase & Co and its affiliates (“JP Morgan”), and is not a product of the Research Department of JP Morgan. JP Morgan makes no representation and warranty whatsoever and disclaims all liability, for the completeness, accuracy or reliability of the information contained herein. This document is not intended as investment research or investment advice, or a recommendation, offer or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction, and shall not constitute a solicitation under any jurisdiction or to any person, if such solicitation under such jurisdiction or to such person would be unlawful. 2024 JP Morgan Chase & Co. All rights reserved.

Y. Song was supported in part by the National Basic Research Program of China Grant 2011CBA00300, 2011CBA00301, the National Natural Science Foundation of China Grant 61033001, 61361136003.

References

- ACD⁺19. Mark Abspoel, Ronald Cramer, Ivan Damgård, Daniel Escudero, and Chen Yuan. Efficient information-theoretic secure multiparty computation over $\mathbb{Z}/p^k\mathbb{Z}$ via galois rings. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019: 17th Theory of Cryptography Conference, Part I*, volume 11891 of *Lecture Notes in Computer Science*, pages 471–501, Nuremberg, Germany, December 1–5, 2019. Springer, Heidelberg, Germany.

- ACD⁺20. Mark Abspoel, Ronald Cramer, Ivan Damgård, Daniel Escudero, Matthieu Rabaud, Chaoping Xing, and Chen Yuan. Asymptotically good multiplicative LSSS over Galois rings and applications to MPC over $\mathbb{Z}/p^k\mathbb{Z}$. In Shiho Moriai and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2020, Part III*, volume 12493 of *Lecture Notes in Computer Science*, pages 151–180, Daejeon, South Korea, December 7–11, 2020. Springer, Heidelberg, Germany.
- ACE⁺21. Mark Abspoel, Ronald Cramer, Daniel Escudero, Ivan Damgård, and Chaoping Xing. Improved single-round secure multiplication using regenerating codes. In Mehdi Tibouchi and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2021, Part II*, volume 13091 of *Lecture Notes in Computer Science*, pages 222–244, Singapore, December 6–10, 2021. Springer, Heidelberg, Germany.
- BFO12. Eli Ben-Sasson, Serge Fehr, and Rafail Ostrovsky. Near-linear unconditionally-secure multiparty computation with a dishonest minority. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology – CRYPTO 2012*, volume 7417 of *Lecture Notes in Computer Science*, pages 663–680, Santa Barbara, CA, USA, August 19–23, 2012. Springer, Heidelberg, Germany.
- BGP92. Piotr Berman, Juan A Garay, and Kenneth J Perry. Bit optimal distributed consensus. *Computer Science: Research and Applications*, pages 313–321, 1992.
- BTH08. Zuzana Beerliová-Trubíniová and Martin Hirt. Perfectly-secure MPC with linear communication complexity. In Ran Canetti, editor, *TCC 2008: 5th Theory of Cryptography Conference*, volume 4948 of *Lecture Notes in Computer Science*, pages 213–230, San Francisco, CA, USA, March 19–21, 2008. Springer, Heidelberg, Germany.
- Can00. Ran Canetti. Security and composition of multiparty cryptographic protocols. *Journal of Cryptology*, 13(1):143–202, January 2000.
- CCXY18. Ignacio Cascudo, Ronald Cramer, Chaoping Xing, and Chen Yuan. Amortized complexity of information-theoretically secure MPC revisited. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018, Part III*, volume 10993 of *Lecture Notes in Computer Science*, pages 395–426, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Heidelberg, Germany.
- CRX21. Ronald Cramer, Matthieu Rabaud, and Chaoping Xing. Asymptotically-good arithmetic secret sharing over $\mathbb{Z}/p^{\ell}\mathbb{Z}$ with strong multiplication and its applications to efficient MPC. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology – CRYPTO 2021, Part III*, volume 12827 of *Lecture Notes in Computer Science*, pages 656–686, Virtual Event, August 16–20, 2021. Springer, Heidelberg, Germany.
- CW92. Brian A Coan and Jennifer L Welch. Modular construction of a byzantine agreement protocol with optimal message bit complexity. *Information and Computation*, 97(1):61–85, 1992.
- DEF⁺19. Ivan Damgård, Daniel Escudero, Tore Kasper Frederiksen, Marcel Keller, Peter Scholl, and Nikolaj Volgushev. New primitives for actively-secure MPC over rings with applications to private machine learning. In *2019 IEEE Symposium on Security and Privacy*, pages 1102–1120, San Francisco, CA, USA, May 19–23, 2019. IEEE Computer Society Press.
- DIK10. Ivan Damgård, Yuval Ishai, and Mikkel Krøigaard. Perfectly secure multiparty computation and the computational overhead of cryptography. In Henri Gilbert, editor, *Advances in Cryptology – EUROCRYPT 2010*, volume 6110 of *Lecture Notes in Computer Science*, pages 445–465, French Riviera, May 30 – June 3, 2010. Springer, Heidelberg, Germany.
- DN07. Ivan Damgård and Jesper Buus Nielsen. Scalable and unconditionally secure multiparty computation. In Alfred Menezes, editor, *Advances in Cryptology – CRYPTO 2007*, volume 4622 of *Lecture Notes in Computer Science*, pages 572–590, Santa Barbara, CA, USA, August 19–23, 2007. Springer, Heidelberg, Germany.
- ELXY23. Daniel Escudero, Hongqing Liu, Chaoping Xing, and Chen Yuan. Degree- d reverse multiplication-friendly embeddings: Constructions and applications. *Asiacrypt*, 2023.
- FR22. Thibault Feneuil and Matthieu Rivain. Threshold linear secret sharing to the rescue of MPC-in-the-head. Cryptology ePrint Archive, Report 2022/1407, 2022. <https://eprint.iacr.org/2022/1407>.
- FY92. Matthew K. Franklin and Moti Yung. Communication complexity of secure computation (extended abstract). In *24th Annual ACM Symposium on Theory of Computing*, pages 699–710, Victoria, BC, Canada, May 4–6, 1992. ACM Press.
- GLS19. Vipul Goyal, Yanyi Liu, and Yifan Song. Communication-efficient unconditional MPC with guaranteed output delivery. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part II*, volume 11693 of *Lecture Notes in Computer Science*, pages 85–114, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Heidelberg, Germany.
- GPS21. Vipul Goyal, Antigoni Polychroniadou, and Yifan Song. Unconditional communication-efficient MPC via hall’s marriage theorem. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology – CRYPTO 2021, Part II*, volume 12826 of *Lecture Notes in Computer Science*, pages 275–304, Virtual Event, August 16–20, 2021. Springer, Heidelberg, Germany.

- GPS22. Vipul Goyal, Antigoni Polychroniadou, and Yifan Song. Sharing transformation and dishonest majority MPC with packed secret sharing. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022, Part IV*, volume 13510 of *Lecture Notes in Computer Science*, pages 3–32, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Heidelberg, Germany.
- GSZ20. Vipul Goyal, Yifan Song, and Chenzhi Zhu. Guaranteed output delivery comes free in honest majority MPC. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology – CRYPTO 2020, Part II*, volume 12171 of *Lecture Notes in Computer Science*, pages 618–646, Santa Barbara, CA, USA, August 17–21, 2020. Springer, Heidelberg, Germany.
- HMP00. Martin Hirt, Ueli M. Maurer, and Bartosz Przydatek. Efficient secure multi-party computation. In Tatsuaki Okamoto, editor, *Advances in Cryptology – ASIACRYPT 2000*, volume 1976 of *Lecture Notes in Computer Science*, pages 143–161, Kyoto, Japan, December 3–7, 2000. Springer, Heidelberg, Germany.
- IKP⁺16. Yuval Ishai, Eyal Kushilevitz, Manoj Prabhakaran, Amit Sahai, and Ching-Hua Yu. Secure protocol transformations. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology – CRYPTO 2016, Part II*, volume 9815 of *Lecture Notes in Computer Science*, pages 430–458, Santa Barbara, CA, USA, August 14–18, 2016. Springer, Heidelberg, Germany.
- PS21. Antigoni Polychroniadou and Yifan Song. Constant-overhead unconditionally secure multiparty computation over binary fields. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology – EUROCRYPT 2021, Part II*, volume 12697 of *Lecture Notes in Computer Science*, pages 812–841, Zagreb, Croatia, October 17–21, 2021. Springer, Heidelberg, Germany.
- Sha79. Adi Shamir. How to share a secret. *Communications of the Association for Computing Machinery*, 22(11):612–613, November 1979.
- Wan03. Zhe-Xian Wan. *Lectures on finite fields and Galois rings*. World Scientific Publishing Company, 2003.

Appendix

A Full Description of Building Block Functionalities

Full Description of $\mathcal{F}_{\text{OpenPub}}$. The full description of $\mathcal{F}_{\text{OpenPub}}$ appears in Functionality 2.

Functionality 2: $\mathcal{F}_{\text{OpenPub}}(N)$

1. Let $[x_1]_d, [x_2]_d, \dots, [x_N]_d$ denote the input sharings.
2. For all $i \in [N]$, $\mathcal{F}_{\text{OpenPub}}$ receives from the active honest parties their shares of $[x_i]_d$. Then $\mathcal{F}_{\text{OpenPub}}$ reconstructs the secrets x_i . $\mathcal{F}_{\text{OpenPub}}$ further computes the shares of $[x_i]_d$ held by active corrupted parties, and sends the shares to the adversary.
3. For all $i \in [N]$, $\mathcal{F}_{\text{OpenPub}}$ sends the reconstructed x_i to all active parties.

Full Description of $\mathcal{F}_{\text{Mult}}$. The full description of $\mathcal{F}_{\text{Mult}}$ appears in Functionality 3.

Functionality 3: $\mathcal{F}_{\text{Mult}}(N)$

1. Let $[x_1]_t, \dots, [x_N]_t$ and $[y_1]_t, \dots, [y_N]_t$ denote the input sharings. $\mathcal{F}_{\text{Mult}}$ receives from active honest parties their shares of $[x_1]_t, \dots, [x_N]_t$ and $[y_1]_t, \dots, [y_N]_t$. Then $\mathcal{F}_{\text{Mult}}$ reconstructs the secrets x_1, \dots, x_N and y_1, \dots, y_N . $\mathcal{F}_{\text{Mult}}$ further computes the shares of $[x_1]_t, \dots, [x_N]_t$ and $[y_1]_t, \dots, [y_N]_t$ held by active corrupted parties and send the shares to the adversary.
2. $\mathcal{F}_{\text{Mult}}$ receives from the adversary a set of even number of parties $S \subset \mathcal{P}_{\text{active}}$ S such that $|S|/2 \leq |S \cap \mathcal{C}_{\text{active}}|$. $\mathcal{F}_{\text{Mult}}$ updates $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S$. $\mathcal{F}_{\text{Mult}}$ then sends S to all active honest parties.
3. $\mathcal{F}_{\text{Mult}}$ computes $x_1 \cdot y_1, \dots, x_N \cdot y_N$, and then $\mathcal{F}_{\text{Mult}}$ receives from the adversary the shares of active corrupted parties of $[x_1 \cdot y_1]_t, \dots, [x_N \cdot y_N]_t$.
4. $\mathcal{F}_{\text{Mult}}$ samples the whole sharings $[x_1 \cdot y_1]_t, \dots, [x_N \cdot y_N]_t$ so that they are compatible with the active corrupted parties' shares. For each active honest party P_h , $\mathcal{F}_{\text{Mult}}$ sends P_h 's shares of $[x_1 \cdot y_1]_t, \dots, [x_N \cdot y_N]_t$ to P_h .

Full Description of $\mathcal{F}_{\text{ReEncode}}$. The full description of $\mathcal{F}_{\text{ReEncode}}$ appears in Functionality 4.

Functionality 4: $\mathcal{F}_{\text{ReEncode}}(N)$

1. Let $[x_1]_t, \dots, [x_N]_t$ denote the input sharings. $\mathcal{F}_{\text{ReEncode}}$ receives from active honest parties their shares of $[x_1]_t, \dots, [x_N]_t$. Then $\mathcal{F}_{\text{ReEncode}}$ reconstructs the secrets x_1, \dots, x_N . $\mathcal{F}_{\text{ReEncode}}$ further computes the shares of $[x_1]_t, \dots, [x_N]_t$ held by active corrupted parties and send the shares to the adversary.
2. $\mathcal{F}_{\text{ReEncode}}$ receives from the adversary a set of even number of parties $S \subset \mathcal{P}_{\text{active}}$ S such that $|S|/2 \leq |S \cap \mathcal{C}_{\text{active}}|$. $\mathcal{F}_{\text{ReEncode}}$ updates $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S$. $\mathcal{F}_{\text{ReEncode}}$ then sends S to all active honest parties.
3. $\mathcal{F}_{\text{ReEncode}}$ computes $\phi \circ \psi(x_1), \dots, \phi \circ \psi(x_N)$, and then $\mathcal{F}_{\text{ReEncode}}$ receives from the adversary the shares of active corrupted parties of $[\phi \circ \psi(x_1)]_t, \dots, [\phi \circ \psi(x_N)]_t$.
4. $\mathcal{F}_{\text{ReEncode}}$ samples the whole sharings $[\phi \circ \psi(x_1)]_t, \dots, [\phi \circ \psi(x_N)]_t$ so that they are compatible with the active corrupted parties' shares. For each active honest party P_h , $\mathcal{F}_{\text{ReEncode}}$ sends P_h 's shares of $[\phi \circ \psi(x_1)]_t, \dots, [\phi \circ \psi(x_N)]_t$ to P_h .

Full Description of $\mathcal{F}_{\text{VerifyBC}}$. The full description of $\mathcal{F}_{\text{VerifyBC}}$ appears in Functionality 5.

Functionality 5: $\mathcal{F}_{\text{VerifyBC}}(N)$

1. Let D be the suspicious dealer and let N be the total number of ring elements to verify. Let $x^{(1)}, \dots, x^{(N)}$ denote these elements.
2. $\mathcal{F}_{\text{VerifyBC}}$ receives from active honest parties the values of $x^{(1)}, \dots, x^{(N)}$, and send them to the adversary.
3. $\mathcal{F}_{\text{VerifyBC}}$ receives from the adversary one of the following:
 - Default string only if active honest parties shares are all consistent;

- A semi-corrupted pair $\{P_{j_1}, P_{j_2}\}$ where $\{P_{j_1}, P_{j_2}\} \cap \mathcal{C}_{\text{active}} \neq \emptyset$.
- 4. Based on what it has received, $\mathcal{F}_{\text{VerifyBC}}$ does one of the following:
 - On receiving default string, $\mathcal{F}_{\text{VerifyBC}}$ sends **consistent** to active honest parties;
 - On receiving a semi-corrupted pair, $\mathcal{F}_{\text{VerifyBC}}$ sends (**inconsistent**, $\{P_{j_1}, P_{j_2}\}$) to active honest parties.

Full Description of $\mathcal{F}_{\text{InputShamir}}$. The full description of $\mathcal{F}_{\text{InputShamir}}$ appears in Functionality 6.

Functionality 6: $\mathcal{F}_{\text{InputShamir}}(N)$

1. Let $x_1, \dots, x_N \in \mathcal{R}$ denote the inputs of the client denoted by Client.
2. $\mathcal{F}_{\text{InputShamir}}$ receives from the adversary a set of even number of parties $S \subset \mathcal{P}_{\text{active}}$ such that $|S|/2 \leq |S \cap \mathcal{C}_{\text{active}}|$. $\mathcal{F}_{\text{InputShamir}}$ updates $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S$. $\mathcal{F}_{\text{InputShamir}}$ then sends S to all active honest parties.
3. $\mathcal{F}_{\text{InputShamir}}$ receives from the adversary one of the following:
 - The shares of active corrupted parties of $[x_1]_t, \dots, [x_N]_t$.
 - If Client is corrupted, the adversary can choose to change the inputs to x'_1, \dots, x'_N . Then the adversary sends the shares of active corrupted parties of $[x'_1]_t, \dots, [x'_N]_t$.
4. Based on what it has received, $\mathcal{F}_{\text{InputShamir}}$ does one of the following:
 - On receiving the shares of active corrupted parties of $[x_1]_t, \dots, [x_N]_t$, $\mathcal{F}_{\text{InputShamir}}$ samples a random degree- t Shamir secret sharing $[x_1]_t, \dots, [x_N]_t$ so that they are compatible with the active corrupted parties' shares for all $i \in [N]$. For each active honest party P_h , $\mathcal{F}_{\text{InputShamir}}$ sends P_h 's shares of $[x_1]_t, \dots, [x_N]_t$ to P_h .
 - On receiving a new set of inputs x'_1, \dots, x'_N and the shares of active corrupted parties of $[x'_1]_t, \dots, [x'_N]_t$, $\mathcal{F}_{\text{InputShamir}}$ samples a random degree- t Shamir secret sharing $[x'_1]_t, \dots, [x'_N]_t$ so that they are compatible with the active corrupted parties' shares. For each active honest party P_h , $\mathcal{F}_{\text{InputShamir}}$ sends P_h 's shares of $[x'_1]_t, \dots, [x'_N]_t$ to P_h for all $i \in [N]$.

B Preparing Random Sharings

B.1 Preliminaries

We first introduce several important tools to help us instantiate our protocols that prepare different random secret sharings.

Byzantine Agreement. Byzantine agreement (BA) is a protocol that takes a bit from each party as input and enables all honest parties to reach a consensus on a single bit. In the case that all parties hold the same bit as input then the output bit coincides precisely with this input bit. In our protocol, we use BA in a few places, like letting all parties reach a binary consensus on whether an error happened. We also use BA to let one party broadcast one bit to all other parties consistently, meaning all parties receive the same bit, which is the bit the party sent if this party is honest. This is easily instantiable from BA by letting the party send the bit to all other parties, followed by BA to reach a consensus on the bit received.

With $t < n/3$, both bit consensus and broadcast can be achieved by a perfect byzantine agreement protocol communicating $O(n^2)$ bits [BGP92, CW92]. This quadratic complexity is fine in our context since we only use this primitive a small number of times so that it does not affect our overall linear communication.

Packed Shamir Secret Sharings. We will use the packed Shamir secret sharing scheme [FY92] in this work to prepare randomness for network routing. For the Galois ring $\mathcal{R} = \text{GR}(2^k, m)$, suppose $(\alpha_i)_{i=1}^n, (\beta_j)_{j=1}^n$ are $2n$ distinct points to interpolate polynomials according to Lemma 1.

For all $1 \leq r \leq n$, a degree- d ($d \geq r - 1$) packed Shamir sharing of $\mathbf{x} \in \mathcal{R}^r$ is a vector $(s_1, s_2, \dots, s_n) \in \mathcal{R}^n$ that satisfies the property that there exists a polynomial $f(\cdot) \in \mathcal{R}^{\leq d}[x]$ with $f(\beta_j) = x_j, \forall j \in [r]$ and $f(\alpha_i) = s_i, \forall i \in [n]$. The share held by party P_i is s_i . With any $(d + 1)$

different shares of the same sharing, we can reconstruct the secret \mathbf{x} . When $r = 1$, it corresponds to the standard Shamir sharing.

A degree- d packed Shamir sharing for the vector $\mathbf{x} \in \mathcal{R}^r$ is denoted as $[\mathbf{x}]_d$, and a degree- d standard Shamir sharing for the element $y \in \mathcal{R}$ is denoted as $[y]_d$. For packed Shamir sharings, the following two properties hold:

- For all $\mathbf{x}, \mathbf{y} \in \mathcal{R}^r$, $[\mathbf{x} + \mathbf{y}]_d = [\mathbf{x}]_d + [\mathbf{y}]_d$.
- Let \star denote the coordinate-wise multiplication, then for all $\mathbf{x}, \mathbf{y} \in \mathcal{R}^r$ and for all $d_1, d_2 \geq r - 1$ subject to $d_1 + d_2 < n$, we have $[\mathbf{x} \star \mathbf{y}]_{d_1+d_2} = [\mathbf{x}]_{d_1} \cdot [\mathbf{y}]_{d_2}$.

We note that Lemma 2 also holds for packed Shamir secret sharings.

Abstract General Linear Secret Sharing Schemes. We adopt the idea of an abstract definition of a general linear secret sharing scheme (GLSSS) in [ACD⁺20]. We borrow the following notations from [ACD⁺20].

For two non-empty sets U and \mathcal{I} , $U^{\mathcal{I}}$ represents the indexed Cartesian product $\prod_{i \in \mathcal{I}} U$. For some non-empty $A \subset \mathcal{I}$, the natural projection π_A maps one tuple $u = (u_i)_{i \in \mathcal{I}} \in U^{\mathcal{I}}$ to the tuple $(u_i)_{i \in A} \in U^A$. We also suppose that \mathcal{S} is a finite quotient ring.

We use the definition of arithmetic secret sharing scheme from [ACD⁺20].

Definition 3 (Abstract \mathcal{S} -GLSSS [ACD⁺20]). *The syntax of an \mathcal{S} -arithmetic secret sharing scheme Σ consists of the following data:*

- A set of parties $\mathcal{I} = \{1, \dots, n\}$.
- A secret space $Z = \mathcal{S}^r$. r is also denoted as the number of secrets packed within Σ .
- A share space $U = \mathcal{S}^\ell$. ℓ is also denoted as the share size.
- A sharing space $C \subset U^{\mathcal{I}}$, where $U^{\mathcal{I}}$ denotes the intended Cartesian product $\prod_{i \in \mathcal{I}} U$.
- An injective \mathcal{S} -module homomorphism: $\mathbf{share} : Z \times \mathcal{S}^k \rightarrow C$, that maps a secret $\mathbf{x} \in Z$ and a random tape $\boldsymbol{\rho} \in \mathcal{S}^k$ to a sharing $[\mathbf{x}] \in C$. \mathbf{share} is also denoted as the sharing map of Σ .
- A surjective \mathcal{S} -module homomorphism: $\mathbf{rec} : C \rightarrow Z$, which takes as input a sharing $[\mathbf{x}] \in C$ and outputs a secret $\mathbf{x} \in Z$. \mathbf{rec} is also denoted as the reconstruction map of Σ .

The scheme Σ satisfies that for all $\mathbf{x} \in Z$ and $\boldsymbol{\rho} \in \mathcal{S}^k$, $\mathbf{rec}(\mathbf{share}(\mathbf{x}, \boldsymbol{\rho})) = \mathbf{x}$. We may refer to Σ as the 6-tuple $(n, Z, U, C, \mathbf{share}, \mathbf{rec})$.

Definition 4 (Privacy Set and Reconstruction Set [CCXY18]). *Suppose $A \subset \mathcal{I}$ is non-empty. We say A is a privacy set if for all $\mathbf{x}_0, \mathbf{x}_1 \in Z$, and for all vector $\mathbf{v} \in U^A$,*

$$\Pr_{\boldsymbol{\rho}}[\pi_A(\mathbf{share}(\mathbf{x}_0, \boldsymbol{\rho})) = \mathbf{v}] = \Pr_{\boldsymbol{\rho}}[\pi_A(\mathbf{share}(\mathbf{x}_1, \boldsymbol{\rho})) = \mathbf{v}].$$

We say A is a reconstruction set if there is an \mathcal{S} -module homomorphism $\mathbf{rec}_A : \pi_A(C) \rightarrow Z$, such that for all $[\mathbf{x}] \in C$,

$$\mathbf{rec}_A(\pi_A([\mathbf{x}])) = \mathbf{rec}([\mathbf{x}]).$$

Intuitively, for a privacy set A , the shares of all parties in A are independent of the secret; for a reconstruction set A , the shares of all parties in A fully determine the secret.

In [CCXY18], the interleaved GLSSS $\Sigma^{\times m}$ is defined as an n -party secret sharing scheme which corresponds to m Σ -sharings. We introduce the *tensoring-up lemma* from [CCXY18], which can be easily adapted to our ring case:

Proposition 1 (Tensoring-up Lemma [CCXY18]). *Let \mathcal{R} be a degree- m Galois ring containing the quotient ring \mathcal{S} and let Σ be a \mathcal{S} -GLSSS. Then the m -fold interleaved \mathcal{S} -GLSSS $\Sigma^{\times m}$ is naturally viewed as an \mathcal{R} -GLSSS, compatible with its \mathcal{S} -linearity.*

Suppose $[\mathbf{x}]$ is a sharing in Σ , with this proposition we can define the linear operation for every $\lambda \in \mathcal{R}$ such that for all $[\mathbf{x}] = ([x_1], \dots, [x_m]) \in \Sigma^{\times m}$ we have:

- $\lambda \cdot [\mathbf{x}] = (\lambda \cdot [x_1], \dots, \lambda \cdot [x_m])$ for all $\lambda \in \mathcal{S}$,
- $\lambda_1 \cdot [\mathbf{x}] + \lambda_2 \cdot [\mathbf{x}] = (\lambda_1 + \lambda_2) \cdot [\mathbf{x}]$ for all $\lambda_1, \lambda_2 \in \mathcal{R}$ and
- $\lambda_1 \cdot (\lambda_2 \cdot [\mathbf{x}]) = (\lambda_1 \cdot \lambda_2) \cdot [\mathbf{x}]$ for all $\lambda_1, \lambda_2 \in \mathcal{R}$.

An Example of a GLSSS Using The Tensoring-up Lemma. We will use the standard Shamir secret sharing scheme as an example of a GLSSS and show how to use the tensoring-up lemma. For ring \mathcal{S} of size $|\mathcal{S}| \geq n+1$, we may define a secret sharing scheme Σ which takes $x \in \mathcal{S}$ as input and outputs $[x]_t$, a degree- t Shamir sharing. The secret space and the share space of Σ are \mathcal{S} . According to the Lagrange interpolation, the secret x can be written as a \mathcal{S} -linear combination of all the shares. Therefore, the defining map of Σ is \mathcal{S} -linear. Therefore, Σ is a \mathcal{S} -GLSSS.

A sharing $[\mathbf{x}] = ([x_1], \dots, [x_m]) \in \Sigma^{\times m}$ is a vector of m sharings in Σ . Let \mathcal{R} be a degree- m Galois ring of \mathcal{S} . According to the tensoring-up lemma $\Sigma^{\times m}$ is a \mathcal{R} -GLSSS. Therefore \mathcal{R} -linear operations can be performed to the sharings in $\Sigma^{\times m}$.

Hyper-Invertible Matrices. We adopt the definition of hyper-invertible matrices from [BTH08].

Definition 5. An r -by- c matrix \mathcal{M} over the Galois ring \mathcal{R} is hyper-invertible if for any index sets $R \subseteq \{1, 2, \dots, r\}$ and $C \subseteq \{1, 2, \dots, c\}$ with $|R| = |C| > 0$, the matrix \mathcal{M}_R^C is invertible, where \mathcal{M}_R denotes the matrix consisting of the rows $i \in R$ of \mathcal{M} and \mathcal{M}^C denotes the matrix consisting of the columns $j \in C$ of \mathcal{M} , and $\mathcal{M}_R^C = (\mathcal{M}_R)^C$.

We note that an n -by- n hyper-invertible matrix exists if the degree of the Galois ring $\mathcal{R} = \text{GR}(2^k, m)$ satisfies $2^m \geq 2n + 1$ [ACD⁺19]. In our work, we will use n -by- n hyper-invertible matrices over \mathcal{R} .

B.2 Preparing Random Sharings for $\mathbb{Z}/2^k\mathbb{Z}$ -GLSSS

In this section, we present the protocol for preparing random sharings for a given general $\mathbb{Z}/2^k\mathbb{Z}$ -GLSSS scheme, denoted by Σ . Our protocol is adapted from [PS21], while we introduce party elimination framework to ensure perfect security. Let $[x]$ denote a sharing in Σ of secret x . For a set $A \subset \mathcal{I}$, recall that $\pi_A([x])$ refers to the share of $[x]$ held by parties in A . We assume that Σ satisfies the following property:

- Given a set $A \subset \mathcal{I}$ and a set of shares $\{a_i\}_{i \in A}$ for the parties in A , let

$$\Sigma(A, (a_i)_{i \in A}) := \{[x] \mid [x] \in \Sigma \text{ and } \pi_A([x]) = (a_i)_{i \in A}\}.$$

Then, there is an efficient algorithm that outputs either that $\Sigma(A, (a_i)_{i \in A}) = \emptyset$ or a random sharing $[x]$ in $\Sigma(A, (a_i)_{i \in A})$.

The functionality $\mathcal{F}_{\text{Rand}}$ is described in Functionality 7. Essentially, $\mathcal{F}_{\text{Rand}}$ allows the adversary to specify the shares held by active corrupted parties or to generate a valid semi-corrupted pair. In the first case, based on these shares $\mathcal{F}_{\text{Rand}}$ generates a random sharing in Σ and distributes the shares to active honest parties. Note that when the set of active corrupted parties is a privacy set, the secret is independent of the shares chosen by the adversary. In the second case, the semi-corrupted pair will be eliminated from the active parties.

Functionality 7: $\mathcal{F}_{\text{Rand}}(N)$

1. $\mathcal{F}_{\text{Rand}}$ receives from the adversary one of the following:
 - The shares of active corrupted parties $(s_i^{(j)})_{i \in \mathcal{C}_{\text{active}}}$ where $\Sigma(\mathcal{C}_{\text{active}}, (s_i^{(j)})_{i \in \mathcal{C}_{\text{active}}}) \neq \emptyset$ for all $j \in [N]$;
 - A semi-corrupted pair $\{P_{j_1}, P_{j_2}\}$ where $\{P_{j_1}, P_{j_2}\} \cap \mathcal{C}_{\text{active}} \neq \emptyset$.
2. Based on what it has received, $\mathcal{F}_{\text{Rand}}$ does one of the following:
 - On receiving $(s_i^{(j)})_{i \in \mathcal{C}_{\text{active}}}$ for all $j \in [N]$, $\mathcal{F}_{\text{Rand}}$ randomly samples $[r] \in \Sigma(\mathcal{C}_{\text{active}}, (s_i^{(j)})_{i \in \mathcal{C}_{\text{active}}})$ for all $j \in [N]$. For each honest party $P_i \in \mathcal{H}_{\text{active}}$, $\mathcal{F}_{\text{Rand}}$ sends the i -th share of $[r^{(j)}]$ to P_i for all $j \in [N]$.
 - On receiving $\{P_{j_1}, P_{j_2}\}$, $\mathcal{F}_{\text{Rand}}$ sends $\{P_{j_1}, P_{j_2}\}$ to all active honest parties.

We follow the idea in [BTH08] to design a protocol that prepares random sharings in Σ . To sum up, firstly each party generates a random sharing in $\Sigma^{\times m}$. Each sharing $[\mathbf{x}]$ in $\Sigma^{\times m}$ is a vector of m sharings denoted by $([x_1], [x_2], \dots, [x_m])$, and according to Proposition 1 $\Sigma^{\times m}$ is a

\mathcal{R} -GLSSS. Secondly, all parties apply a hyper-invertible matrix over \mathcal{R} to transform the sharings for fault detection. Next, some selected parties will verify that a subset of the transformed sharings have the correct form. Finally, all parties take the other transformed sharings as the output.

Protocol 7: $\Pi_{\text{Rand}}(N)$

1. All parties agree on a hyper-invertible matrix \mathcal{M} over \mathcal{R} of size n' -by- n' .
2. Each party in $\mathcal{P}_{\text{active}}$ uniformly randomly samples $\lceil \frac{N}{m \cdot T} \rceil$ different sharings in $\Sigma^{\times m}$. Let $\llbracket \mathbf{s}_i^{(j)} \rrbracket$ denote the j -th sharing sampled by P_i . Then each party in $\mathcal{P}_{\text{active}}$ distributes the shares to all the other parties in $\mathcal{P}_{\text{active}}$.
3. For all $j \in \{1, 2, \dots, \lceil \frac{N}{m \cdot T} \rceil\}$, all parties in $\mathcal{P}_{\text{active}}$ locally compute

$$(\llbracket \mathbf{r}_1^{(j)} \rrbracket, \llbracket \mathbf{r}_2^{(j)} \rrbracket, \dots, \llbracket \mathbf{r}_{n'}^{(j)} \rrbracket) = \mathcal{M} \cdot (\llbracket \mathbf{s}_1^{(j)} \rrbracket, \llbracket \mathbf{s}_2^{(j)} \rrbracket, \dots, \llbracket \mathbf{s}_{n'}^{(j)} \rrbracket)$$

4. For all $j \in \{1, 2, \dots, \lceil \frac{N}{m \cdot T} \rceil\}$, and for $\ell \in \{T + 1, \dots, n'\}$, all parties in $\mathcal{P}_{\text{active}}$ send their shares of $\llbracket \mathbf{r}_\ell^{(j)} \rrbracket$ to P_h . P_h checks if the sharing $\llbracket \mathbf{r}_\ell^{(j)} \rrbracket$ is a valid sharing in $\Sigma^{\times m}$. If not, P_h sets its happy-bit to **unhappy**.
5. **Fault Detection Phase:**
 - (a) All players broadcast their happy-bits using the byzantine agreement protocol.
 - (b) For each party, if it receives at least one **unhappy** then sets his or her happy-bit to **unhappy**.
 - (c) All players run a consensus protocol on their happy-bits. If the consensus is **happy**, all parties output the output of π and terminate the procedure. Otherwise, they proceed to the following steps.
6. **Fault Localization Phase:**
 - (a) All players agree the party with the smallest index in \mathcal{P}_A as the dealer D . All other players send all their generated values and communication to D .
 - (b) On receiving all the information, D simulates π and the fault detection phase itself. D either prepares the message $(P_i, \text{corrupt})$ if P_i failed to follow the protocol, or the message $(P_{j_1}, P_{j_2}, \text{idx}, \mathbf{b}, \mathbf{b}', \text{disputed})$ if the idx -th bit \mathbf{b} that P_{j_1} sends to P_{j_2} is inconsistent with P_{j_2} 's bit \mathbf{b}' that P_{j_2} claims to have received. Then D broadcasts the prepared message to all players.
 - (c) If $(P_i, \text{corrupt})$ is received, all players set the eliminating set $E = \{D, P_i\}$. Otherwise, if $(P_{j_1}, P_{j_2}, \text{idx}, \mathbf{b}, \mathbf{b}', \text{disputed})$ is received, P_{j_1} and P_{j_2} will broadcast if they agree with this message. If P_{j_1} does not agree, all players set the eliminating set $E = \{D, P_{j_1}\}$; otherwise if P_{j_2} does not agree, all players set the eliminating set $E = \{D, P_{j_2}\}$; otherwise all players set the eliminating set $E = \{P_{j_1}, P_{j_2}\}$.
 - (d) All players update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - E$ and halt.
7. For all $j \in \{1, 2, \dots, \lceil \frac{N}{m \cdot T} \rceil\}$ and for all $i \in [T]$, all active parties separate the sharing $\llbracket \mathbf{r}_i^{(j)} \rrbracket$ into m sharings in Σ , and take them as outputs.

Lemma 5. *In Step 4 of Protocol 7, either at least one active honest party has **unhappy** happy-bit or all active corrupted parties has dealt consistent shares to active honest parties in $\Sigma^{\times m}$ in Step 2 of Π_{Rand} , i.e. $\Sigma(\mathcal{H}_{\text{active}}, (\llbracket \mathbf{s}_\ell^{(j)} \rrbracket)_i) \neq \emptyset$ where $\llbracket \mathbf{s}_\ell^{(j)} \rrbracket_i$ denotes party P_i 's share of the sharing $\llbracket \mathbf{s}_\ell^{(j)} \rrbracket$.*

Proof. If all actives honest parties are happy, then at least t' active parties have verified that $\llbracket \mathbf{r}_\ell^{(j)} \rrbracket$ is valid for all $j \in \lceil \frac{N}{m \cdot T} \rceil$, and we suppose without loss of generality that $\llbracket \mathbf{r}_\ell^{(j)} \rrbracket$ is valid for $\ell \in \{T + 1, \dots, T + t'\}$. And also we suppose without loss of generality that $\llbracket \mathbf{s}_i^{(j)} \rrbracket$ is valid for $i \in \{1, \dots, n' - t'\}$. Then we have that

$$(\llbracket \mathbf{r}_{T+1}^{(j)} \rrbracket, \dots, \llbracket \mathbf{r}_{T+t'}^{(j)} \rrbracket) = \mathcal{M}_{\mathcal{H}_{\text{active}}} \cdot (\llbracket \mathbf{s}_1^{(j)} \rrbracket, \dots, \llbracket \mathbf{s}_{n'-t'}^{(j)} \rrbracket) + \mathcal{M}_{\mathcal{C}_{\text{active}}} \cdot (\llbracket \mathbf{s}_{n'-t'+1}^{(j)} \rrbracket, \dots, \llbracket \mathbf{s}_n^{(j)} \rrbracket),$$

where $\mathcal{M}_{\mathcal{H}_{\text{active}}} = \mathcal{M}_{\{T+1, \dots, T+t'\}^{\{1, \dots, n'-t'\}}}$ and $\mathcal{M}_{\mathcal{C}_{\text{active}}} = \mathcal{M}_{\{T+1, \dots, T+t'\}^{\{n'-t'+1, \dots, n'\}}}$. Because \mathcal{M} is a hyper-invertible matrix, $\mathcal{M}_{\mathcal{C}}$ is an invertible matrix. Therefore, the shares of $\llbracket \mathbf{s}_{n'-t'+1}^{(j)} \rrbracket, \dots, \llbracket \mathbf{s}_n^{(j)} \rrbracket$ held by active honest parties are also consistent.

It follows Lemma 5 that if all active honest parties are happy then the shares of $\llbracket \mathbf{r}_\ell^{(j)} \rrbracket$ held by active honest parties are consistent for all $\ell \in \{1, \dots, n'\}$, and all active honest parties hold $\lceil N/m \rceil$ consistent sharings in $\Sigma^{\times m}$.

Cost of Protocol 7. Suppose the share size of a sharing in Σ is sh ring elements in $\mathbb{Z}/2^k\mathbb{Z}$. Then the share size of a sharing in $\Sigma^{\times m}$ is $m \cdot \text{sh}$ ring elements in $\mathbb{Z}/2^k\mathbb{Z}$. It follows that the communication complexity of $\Pi_{\text{Rand}}(N)$ is $O(N \cdot n \cdot \text{sh} + n^2 \cdot m \cdot \text{sh})$.

Lemma 6. *The protocol Π_{Rand} computes $\mathcal{F}_{\text{Rand}}$ with perfect security when $|\mathcal{C}_{\text{active}}| < |\mathcal{P}_{\text{active}}|/3$.*

Proof. Let \mathcal{A} denote the adversary. We will construct a simulator Sim to simulate the behaviors of honest parties for the adversary. Recall that $\mathcal{C}_{\text{active}}$ denotes the active corrupted parties and that $\mathcal{H}_{\text{active}}$ denotes the active honest parties.

Simulation for Π_{Rand} . Sim simulates the protocol Π_{Rand} as follows:

1. In Step 2, for each honest party, Sim generates random shares of active corrupted parties in $\Sigma^{\times m}$. Then Sim sends these shares of active corrupted parties on behalf of active honest parties to \mathcal{A} . For each corrupted party, Sim receives the active honest parties' shares of the active corrupted parties' sharings from \mathcal{A} . Suppose that the sharings generated by active corrupted parties are $(\llbracket \mathbf{s}_i^{(j)} \rrbracket)_{i \in \mathcal{C}_{\text{active}}}$ for all $j \in \{1, 2, \dots, \lceil \frac{N}{m \cdot T} \rceil\}$. In this step, Sim only receives from \mathcal{A} the shares of active honest parties of these shares.
2. In Step 4, to simulate the happy-bit, Sim computes the shares held by active honest parties for each sharing generated by the corrupted party. Next, Sim can compute the active honest parties' shares according to $\mathcal{M}_{\{T+1, \dots, n'\}^{\mathcal{C}_{\text{active}}}} \cdot (\llbracket \mathbf{s}_{n'-t'+1}^{(j)} \rrbracket, \dots, \llbracket \mathbf{s}_{n'}^{(j)} \rrbracket)$. Suppose the result sharing is denoted by $(\llbracket \mathbf{w}_{T+1}^{(j)} \rrbracket, \dots, \llbracket \mathbf{w}_{n'}^{(j)} \rrbracket)$. Then Sim receives from \mathcal{A} the shares of $(\llbracket \mathbf{w}_{T+1}^{(j)} \rrbracket, \dots, \llbracket \mathbf{w}_{n'}^{(j)} \rrbracket)$ that the active corrupted parties send to the active honest parties. For $\ell \in \{T+1, \dots, n'\}$, if P_h is an honest party, Sim checks if the sharing $\llbracket \mathbf{w}_\ell^{(j)} \rrbracket$ is a valid sharing in $\Sigma^{\times m}$. If not, Sim sets P_h 's happy-bit to be **unhappy** on behalf of P_h . If P_h is corrupted, Sim does nothing.
3. In Fault Detection Phase, Sim faithfully follows the byzantine agreement protocol using active honest parties' happy-bits.
4. If proceeds to Fault Localization Phase, Sim randomly samples the sharings distributed by active honest parties in Step 2 of Π_{Rand} so that the sampled sharings are consistent with the shares of active corrupted parties. Then Sim faithfully follows the protocol.
5. In Step 7, Sim compute the shares of $\llbracket \mathbf{r}_i^{(j)} \rrbracket$ held by active corrupted parties for all $i \in [T]$ and $j \in [\lceil \frac{N}{m \cdot T} \rceil]$, and Sim sends the shares of the first N sharings held by active corrupted parties to $\mathcal{F}_{\text{Rand}}$.

Hybrid Arguments. We show that Sim perfectly simulates the behaviors of honest parties.

Hybrid₀: The execution in the real world.

Hybrid₁: In Step 2, for each $i \in \mathcal{H}_{\text{active}}$, instead of receiving the entire share from P_i , Sim only samples the active corrupted parties' shares. This will not change the distribution of the shares held by active honest parties. The distribution is identical to the distribution of **Hybrid₀**.

Hybrid₂: In Step 4, Sim simulates the happy-bits of active honest parties. Because the happy-bit of some honest parties only depends on the consistency of the active corrupted parties' shares, Sim can perfectly simulate active honest parties' happy-bits. The distribution is identical to the distribution of **Hybrid₁**.

Hybrid₃: In Fault Detection Phase, Sim simulates the active honest parties' behavior using the simulated active honest parties' happy-bits instead. Because the two sets of active honest parties' happy-bits are identical, the distribution of this hybrid is identical to the distribution of **Hybrid₂**.

Hybrid₄: Sim simulates Fault Localization Phase by sampling random sharings that are compatible with active corrupted parties shares for active honest parties in Step 2. Note that the $\Sigma^{\times m}$ sharings sampled by Sim have the same distribution as the actual sharings sampled by active honest parties. Also note that it will not affect the output distribution, since the semi-corrupted pair is only related to the corrupted parties' messages. Therefore, the distribution of this hybrid is identical to the distribution of **Hybrid₃**.

Hybrid₅: According to Lemma 5, if Sim proceeds to Step 7 in Π_{Rand} then all active honest parties have been dealt consistent sharings from active corrupted parties in Step 2. With these sharings, \mathcal{A} locally holds the shares of $(\llbracket \mathbf{s}_i^{(j)} \rrbracket)_{i \in \mathcal{C}_{\text{active}}}$ that are consistent with the shares of active honest

parties. Specifically, active corrupted parties hold consistent $(\llbracket \mathbf{s}_i^{(j)} \rrbracket)_{i \in \mathcal{C}_{\text{active}}}$ for all $j \in \llbracket \frac{N}{m \cdot T} \rrbracket$, where $\llbracket \mathbf{s}_i^{(j)} \rrbracket \in \Sigma(\mathcal{H}_{\text{active}}, (s_i^{(j)})_{i \in \mathcal{H}_{\text{active}}})$ for $i \in \mathcal{C}_{\text{active}}$. We have

$$(\llbracket \mathbf{r}_i^{(j)} \rrbracket)_{i=1}^T = \mathcal{M}_{[T]} \cdot (\llbracket \mathbf{s}_i^{(j)} \rrbracket)_{i=1}^{n'} = \mathcal{M}_{[T]}^{\mathcal{H}_{\text{active}}} \cdot (\llbracket \mathbf{s}_i^{(j)} \rrbracket)_{i \in \mathcal{H}_{\text{active}}} + \mathcal{M}_{[T]}^{\mathcal{C}_{\text{active}}} \cdot (\llbracket \mathbf{s}_i^{(j)} \rrbracket)_{i \in \mathcal{C}_{\text{active}}}.$$

Sim computes the active corrupted parties' shares of $(\llbracket \mathbf{r}_i^{(j)} \rrbracket)_{i=1}^T$ using this equation and gets $(r_i^{(j)})_{r \in \mathcal{C}_{\text{active}}}$ for all $j \in \llbracket \frac{N}{m \cdot T} \rrbracket$. **Sim** sends $(r_i^{(j)})_{r \in \mathcal{C}_{\text{active}}}$ for all $j \in \llbracket \frac{N}{m \cdot T} \rrbracket$ to $\mathcal{F}_{\text{Rand}}$.

According to the property of Σ , the shares of $(\llbracket \mathbf{s}_i^{(j)} \rrbracket)_{i \in \mathcal{H}_{\text{active}}}$ held by active corrupted parties are independent of the sharings $(\llbracket \mathbf{s}_i^{(j)} \rrbracket)_{i \in \mathcal{C}_{\text{active}}}$, and are also independent of the shares of $(\llbracket \mathbf{s}_i^{(j)} \rrbracket)_{i \in \mathcal{H}_{\text{active}}}$ held by active honest parties. Because $|\mathcal{H}_{\text{active}}| \geq T$, $(\llbracket \mathbf{r}_i^{(j)} \rrbracket)_{i=1}^T$ corresponds to $(\llbracket \mathbf{r}_i^{(j)} \rrbracket)_{i \in \mathcal{H}_{\text{active}}}$ and has no correlation. It follows that the active honest parties' shares of $(\llbracket \mathbf{r}_i^{(j)} \rrbracket)_{i=1}^T$ are independent of the view of \mathcal{A} and are uniformly random. Therefore, the distribution of this hybrid is identical to the distribution of **Hybrid₄**.

Note that **Hybrid₅** is the execution in the ideal world.

B.3 Instantiating Functionalities

A Robust Version of $\mathcal{F}_{\text{Rand}}$. To instantiate functionalities, we first define a functionality $\mathcal{F}_{\text{RandRobust}}$ to prepare random sharings for a $\mathbb{Z}/2^k\mathbb{Z}$ -GLSSS Σ with guaranteed output delivery. The description of $\mathcal{F}_{\text{RandRobust}}$ appears in Functionality 8.

Functionality 8: $\mathcal{F}_{\text{RandRobust}}(N)$

1. $\mathcal{F}_{\text{RandRobust}}$ receives from the adversary a set of even number of parties $S \subset \mathcal{P}_{\text{active}}$ such that $|S|/2 \leq |S \cap \mathcal{C}_{\text{active}}|$. $\mathcal{F}_{\text{RandRobust}}$ updates $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S$. $\mathcal{F}_{\text{RandRobust}}$ then sends S to all active honest parties.
2. $\mathcal{F}_{\text{RandRobust}}$ receives from the adversary the shares of active corrupted parties $(s_i^{(j)})_{i \in \mathcal{C}_{\text{active}}}$ where $\Sigma(\mathcal{C}_{\text{active}}, (s_i^{(j)})_{i \in \mathcal{C}_{\text{active}}}) \neq \emptyset$ for all $j \in [N]$.
3. $\mathcal{F}_{\text{RandRobust}}$ randomly samples $\llbracket \mathbf{r} \rrbracket \in \Sigma(\mathcal{C}_{\text{active}}, (s_i^{(j)})_{i \in \mathcal{C}_{\text{active}}})$ for all $j \in [N]$. For each honest party $P_h \in \mathcal{H}_{\text{active}}$, $\mathcal{F}_{\text{RandRobust}}$ sends the h -th share of $\llbracket \mathbf{r}^{(j)} \rrbracket$ to P_h for all $j \in [N]$.

We show the protocol that implements $\mathcal{F}_{\text{RandRobust}}$ as follows.

Protocol 8: $\Pi_{\text{RandRobust}}(N)$

1. All parties set $S := \emptyset$.
2. All parties call $\mathcal{F}_{\text{Rand}}(N)$. If the result is N random Σ sharings $\{\llbracket \mathbf{r}^{(i)} \rrbracket\}_{i=1}^N$, all parties output these sharings and S . Otherwise, if the result is a semi-corrupted pair E , all parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - E$ and $S := S \cup E$, and then repeat this step.
3. All parties output S .

Suppose the share size of a sharing in Σ is sh ring elements in $\mathbb{Z}/2^k\mathbb{Z}$. The cost of this protocol is $O(N \cdot n \cdot \text{sh} + n^3 \cdot m \cdot \text{sh})$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

Lemma 7. *The protocol $\Pi_{\text{RandRobust}}$ computes $\mathcal{F}_{\text{RandRobust}}$ with perfect security in the $\mathcal{F}_{\text{Rand}}$ -hybrid model when $|\mathcal{C}_{\text{active}}| < |\mathcal{P}_{\text{active}}|/3$.*

Proof. Let \mathcal{A} denote the adversary. We will construct a simulator **Sim** to simulate the behaviors of honest parties for the adversary. In Step 2, **Sim** faithfully emulates $\mathcal{F}_{\text{Rand}}$. Because there is no other communication, the distribution of the real world is the same as the distribution of the ideal world.

Random Shamir Secret Sharings. The functionality $\mathcal{F}_{\text{RandShamir}}$ enables all parties to prepare N random degree- t Shamir sharings in the form of $[\phi(\mathbf{r})]_t$, where \mathbf{r} is a random vector in $(\mathbb{Z}/2^k\mathbb{Z})^\ell$. The description of $\mathcal{F}_{\text{RandShamir}}$ appears in Functionality 9.

Functionality 9: $\mathcal{F}_{\text{RandShamir}}(N)$

1. $\mathcal{F}_{\text{RandShamir}}$ receives from the adversary a set of even number of parties $S \subset \mathcal{P}_{\text{active}}$ such that $|S|/2 \leq |S \cap \mathcal{C}_{\text{active}}|$. $\mathcal{F}_{\text{RandShamir}}$ updates $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S$. $\mathcal{F}_{\text{RandShamir}}$ then sends S to all active honest parties.
2. $\mathcal{F}_{\text{RandShamir}}$ receives from the adversary the shares of active corrupted parties $(s_i^{(j)})_{i \in \mathcal{C}_{\text{active}}}$ where $s_i^{(j)} \in \mathcal{R}$ for all $i \in \mathcal{C}_{\text{active}}$ and for all $j \in [N]$.
3. $\mathcal{F}_{\text{RandShamir}}$ randomly samples N vectors $\mathbf{r}^{(j)}$ in $(\mathbb{Z}/2^k\mathbb{Z})^\ell$ for all $j \in [N]$. Then $\mathcal{F}_{\text{RandShamir}}$ generates N degree- t Shamir sharings $([(\phi(\mathbf{r}^{(j)}))]_t)_{j=1}^N$ such that the shares of $P_i \in \mathcal{C}_{\text{active}}$ are $(s_i^{(j)})_{j=1}^N$. For each $P_h \in \mathcal{H}_{\text{active}}$, $\mathcal{F}_{\text{RandShamir}}$ sends the i -th share of $([\phi(\mathbf{r}^{(j)})]_t)$ to P_h for all $j \in [N]$.

The instantiation of $\mathcal{F}_{\text{RandShamir}}$ can be found in the next part where we describe the instantiation of the functionality that prepares random packed Shamir sharings.

Random Packed Shamir secret sharings. The functionality $\mathcal{F}_{\text{RandShamir}}$ enables all parties to prepare N random degree- t Shamir sharings in the form of $[\phi(\mathbf{r})]_t$, where \mathbf{r} is a random vector in $(\mathbb{Z}/2^k\mathbb{Z})^\ell$. The description of $\mathcal{F}_{\text{RandShamir}}$ appears in Functionality 9. We introduce the functionality to let all parties prepare random degree- d packed Shamir sharings in the form of $([\phi(\mathbf{r}_1), \phi(\mathbf{r}_2), \dots, \phi(\mathbf{r}_h)]_d)$, where $\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_h$ are h random vectors in $(\mathbb{Z}/2^k\mathbb{Z})^\ell$. The description of $\mathcal{F}_{\text{RandShPacked}}$ appears in Functionality 10.

Functionality 10: $\mathcal{F}_{\text{RandShPacked}}(d, h, N)$

1. $\mathcal{F}_{\text{RandShPacked}}$ receives from the adversary a set of even number of parties $S \subset \mathcal{P}_{\text{active}}$ such that $|S|/2 \leq |S \cap \mathcal{C}_{\text{active}}|$. $\mathcal{F}_{\text{RandShPacked}}$ updates $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S$. $\mathcal{F}_{\text{RandShPacked}}$ then sends S to all active honest parties.
2. $\mathcal{F}_{\text{RandShPacked}}$ receives from the adversary the shares of active corrupted parties $(s_i^{(j)})_{i \in \mathcal{C}_{\text{active}}}$ where $s_i^{(j)} \in \mathcal{R}$ for all $i \in \mathcal{C}_{\text{active}}$ and for all $j \in [N]$.
3. $\mathcal{F}_{\text{RandShPacked}}$ randomly samples $N \cdot h$ vectors $\mathbf{r}_q^{(j)}$ in $(\mathbb{Z}/2^k\mathbb{Z})^\ell$ for $q \in [h]$ and $j \in [N]$. Then $\mathcal{F}_{\text{RandShPacked}}$ generates N degree- d Shamir sharings $([(\phi(\mathbf{r}_1^{(j)}), \phi(\mathbf{r}_2^{(j)}), \dots, \phi(\mathbf{r}_h^{(j)}))]_d)_{j=1}^N$ such that the shares of $P_i \in \mathcal{C}_{\text{active}}$ are $(s_i^{(j)})_{j=1}^N$. For each $P_i \in \mathcal{H}_{\text{active}}$, $\mathcal{F}_{\text{RandShPacked}}$ sends the i -th share of $([\phi(\mathbf{r}_1^{(j)}), \phi(\mathbf{r}_2^{(j)}), \dots, \phi(\mathbf{r}_h^{(j)}))]_d)$ to P_i for all $j \in [N]$.

We use $\mathcal{F}_{\text{RandRobust}}$ to instantiate $\mathcal{F}_{\text{RandShPacked}}$. Consider a secret sharing scheme Σ which takes h vectors $\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_h \in (\mathbb{Z}/2^k\mathbb{Z})^\ell$ as input and outputs a degree- d packed Shamir sharing of $(\phi(\mathbf{r}_1), \phi(\mathbf{r}_2), \dots, \phi(\mathbf{r}_h))$ in \mathcal{R} , i.e. $([\phi(\mathbf{r}_1), \phi(\mathbf{r}_2), \dots, \phi(\mathbf{r}_h)]_d)$. Because packed Shamir sharings are linear in \mathcal{R} and ϕ is $\mathbb{Z}/2^k\mathbb{Z}$ -linear, Σ is an $\mathbb{Z}/2^k\mathbb{Z}$ -GLSSS. To use $\mathcal{F}_{\text{RandRobust}}$ to prepare random sharings in Σ , we need to show that:

- Given a set $A \subset \mathcal{I}$ and a set of shares $\{a_i\}_{i \in A}$ for parties in A , there exists an efficient algorithm which outputs either $\Sigma(A, (a_i)_{i \in A}) = \emptyset$ or a random sharing $([\phi(\mathbf{r}_1), \phi(\mathbf{r}_2), \dots, \phi(\mathbf{r}_h)]_d) \in \Sigma(A, (a_i)_{i \in A})$.

Depending on the size of A , there are three cases:

- If $|A| \geq d + 1$, by the property of packed Shamir secret sharing scheme, the set of shares $\{a_i\}_{i \in A}$ can fully determine the whole sharing if exists. The algorithm checks whether the shares and the secrets lie on a polynomial $f(\cdot)$ of degree at most d . If not, the algorithm outputs $\Sigma(A, (a_i)_{i \in A}) = \emptyset$. Otherwise, the algorithm checks whether the secret of this Shamir sharing, denoted by w_1, \dots, w_ℓ , is in the image of ϕ . This can be done by checking whether $\phi(\psi(w_i)) = w_i$ holds for all $i = 1, \dots, \ell$. If not, the algorithm outputs $\Sigma(A, (a_i)_{i \in A}) = \emptyset$. Otherwise, the algorithm outputs the Shamir sharing determined by $\{a_i\}_{i \in A}$.
- If $d - h + 2 \leq |A| \leq d$, the set of shares $\{a_i\}_{i \in A}$ may not be independent because of the structure of the secret space. Recall that the secret space is $Z \subset \mathcal{R}^h$, and Z is \mathcal{R} -linear. Suppose Z is defined by the linear mapping $B : (\mathbb{Z}/2^k\mathbb{Z})^{\ell \cdot h} \rightarrow (\mathbb{Z}/2^k\mathbb{Z})^{m \cdot h}$. Since packed Shamir secret sharing scheme is also \mathcal{R} -linear, given the shares $\{a_i\}_{i \in A}$ the secret space can be defined as the image of the mapping $\mathcal{L} : (\mathbb{Z}/2^k\mathbb{Z})^{(d-|A|+2) \cdot h} \rightarrow (\mathbb{Z}/2^k\mathbb{Z})^{m \cdot h}$, and \mathcal{L} can be written in the

form that $\mathcal{L}(\mathbf{a}) = C \cdot \mathbf{a} + \mathbf{b}$, where the matrix $C \in (\mathbb{Z}/2^k\mathbb{Z})^{(m \cdot h) \times ((d-|A|+2) \cdot h)}$ and the vector $\mathbf{b} \in (\mathbb{Z}/2^k\mathbb{Z})^{m \cdot h}$ are determined by the shares $\{a_i\}_{i \in A}$. It is clear that $\Sigma(A, (a_i)_{i \in A}) \neq \emptyset$ if and only if $B \cdot \mathbf{x} = C \cdot \mathbf{y} + \mathbf{b}$ has solution. This can be determined easily by solving if the following linear equation has solutions

$$(B, C) \cdot \begin{pmatrix} -\mathbf{x} \\ \mathbf{y} \end{pmatrix} = \mathbf{b},$$

and this can be done by solving the null space of the matrix (B, C) . If the null space is empty, the algorithm outputs $\Sigma(A, (a_i)_{i \in A}) = \emptyset$. Otherwise, the algorithm picks random $(-\mathbf{x}_0, \mathbf{y}_0)^T$ in the null space. Suppose the secrets corresponding to the sampled vector \mathbf{y}_0 are $\phi(\mathbf{r}_1), \dots, \phi(\mathbf{r}_\ell)$. Then, based on the secrets $\phi(\mathbf{r}_1), \dots, \phi(\mathbf{r}_\ell)$ and the shares $\{a_i\}_{i \in A}$, the algorithm reconstructs the whole packed Shamir sharing $[(\phi(\mathbf{r}_1), \phi(\mathbf{r}_2), \dots, \phi(\mathbf{r}_h))]_d$ and outputs $[(\phi(\mathbf{r}_1), \phi(\mathbf{r}_2), \dots, \phi(\mathbf{r}_h))]_d$.

- If $|A| \leq d - h + 1$, by the property of packed Shamir secret sharing scheme the set of shares $\{a_i\}_{i \in A}$ are independent, so $\Sigma(A, (a_i)_{i \in A}) \neq \emptyset$. The algorithm randomly samples $\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_\ell \in (\mathbb{Z}/2^k\mathbb{Z})^\ell$ and randomly samples $(d + 1 - h - |A|)$ elements in \mathcal{R} as the shares of the first $(d + 1 - h - |A|)$ parties in $\mathcal{T} \setminus A$. Then, based on the secrets $\phi(\mathbf{r}_1), \phi(\mathbf{r}_2), \dots, \phi(\mathbf{r}_h)$, the shares of the first $(d + 1 - h - |A|)$ parties in $\mathcal{T} \setminus A$ and the shares $\{a_i\}_{i \in A}$, the algorithm reconstructs the whole packed Shamir sharing $[(\phi(\mathbf{r}_1), \phi(\mathbf{r}_2), \dots, \phi(\mathbf{r}_h))]_d$ and outputs $[(\phi(\mathbf{r}_1), \phi(\mathbf{r}_2), \dots, \phi(\mathbf{r}_h))]_d$.

Therefore, $\mathcal{F}_{\text{RandShPacked}}$ can be instantiated by $\mathcal{F}_{\text{RandRobust}}$ with the secret sharing scheme Σ defined above. Note that the share size of Σ is $\text{sh} = m$ elements in $\mathbb{Z}/2^k\mathbb{Z}$. It follows that when using $\Pi_{\text{Rand}}(N)$ to instantiate $\mathcal{F}_{\text{RandRobust}}$ the total communication complexity to generate N random sharings in Σ is $O(N \cdot n \cdot m + n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

We note that $\mathcal{F}_{\text{RandShamir}}$ is a special case of this functionality when $d = t$ and $h = 1$, and therefore $\mathcal{F}_{\text{RandShamir}}$ can be instantiated by $\mathcal{F}_{\text{RandRobust}}$. Also note that all parties can run the protocol Π_{Rand} to prepare N random packed Shamir sharings.

Random Zero Additive Sharings. The description of $\mathcal{F}_{\text{RandZeroAdd}}$ appears in Functionality 11.

Functionality 11: $\mathcal{F}_{\text{RandZeroAdd}}(N)$

1. $\mathcal{F}_{\text{RandZeroAdd}}$ receives from the adversary a set of even number of parties $S \subset \mathcal{P}_{\text{active}}$ such that $|S|/2 \leq |S \cap \mathcal{C}_{\text{active}}|$. $\mathcal{F}_{\text{RandZeroAdd}}$ updates $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S$. $\mathcal{F}_{\text{RandZeroAdd}}$ then sends S to all active honest parties.
2. $\mathcal{F}_{\text{RandZeroAdd}}$ receives from the adversary the shares of active corrupted parties $(s_i^{(j)})_{i \in \mathcal{C}_{\text{active}} \cap \{1, 2, \dots, t+1\}}$ where $s_i^{(j)} \in \mathbb{Z}/2^k\mathbb{Z}$ for all $i \in \mathcal{C}_{\text{active}} \cap \{1, 2, \dots, t+1\}$ and for all $j \in [N]$.
3. $\mathcal{F}_{\text{RandZeroAdd}}$ randomly samples N additive sharings $(\langle o_j \rangle)_{j=1}^N$ such that $o_j = 0$ for all $j \in [N]$. For each $P_h \in \mathcal{H}_{\text{active}} \cap \{1, 2, \dots, t+1\}$, $\mathcal{F}_{\text{RandZeroAdd}}$ sends the h -th share of $\langle o_j \rangle$ to P_h for all $j \in [N]$.

We will use $\mathcal{F}_{\text{RandRobust}}$ to instantiate $\mathcal{F}_{\text{RandZeroAdd}}$. Consider a secret sharing scheme Σ which outputs an additive sharing of 0 in $\mathbb{Z}/2^k\mathbb{Z}$ among the first $t + 1$ parties, and the shares of the remaining parties are fixed to be 0. It is shown in [PS21] that the following properties holds for Σ :

- Σ is an $\mathbb{Z}/2^k\mathbb{Z}$ -GLSSS.
- Given a set $A \subset \mathcal{T}$ and a set of shares $\{a_i\}_{i \in A}$ for parties in A , there exists an efficient algorithm which outputs either $\Sigma(A, (a_i)_{i \in A}) = \emptyset$ or a random sharing $\langle o \rangle \in \Sigma(A, (a_i)_{i \in A})$.

Therefore, $\mathcal{F}_{\text{RandZeroAdd}}$ can be instantiated by $\mathcal{F}_{\text{RandRobust}}$ with the secret sharing scheme Σ defined above. Note that the share size of Σ is $\text{sh} = 1$ element in $\mathbb{Z}/2^k\mathbb{Z}$. It follows that when using $\Pi_{\text{Rand}}(N)$ to instantiate $\mathcal{F}_{\text{RandRobust}}$ the total communication complexity to generate N random sharings in Σ is $O(N \cdot n + n^3 \cdot m)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

Random Zero Packed Shamir Sharings. When we prepare randomness needed for sharing transformation, all parties need to prepare random degree- d packed Shamir sharings whose secrets are 0. More specifically, all parties need to prepare sharings in the form of $([(o_1^{(j)}, o_2^{(j)}, \dots, o_h^{(j)})]_d)_{j=1}^N$,

where o_i is an element in \mathcal{R} that satisfies $\psi(o_i) = \mathbf{0} \in (\mathbb{Z}/2^k\mathbb{Z})^\ell$ for all $i \in [h]$ and for all $j \in [N]$. We add constraints to d and h such that $d - h + 1 \geq t'$ and $d \leq n' - 1$. The description of $\mathcal{F}_{\text{RandZeroSh}}$ appears in Functionality 12.

Functionality 12: $\mathcal{F}_{\text{RandZeroSh}}(d, h, N)$

1. $\mathcal{F}_{\text{RandZeroSh}}$ receives from the adversary one of the following:
 - The shares of active corrupted parties $(s_i^{(j)})_{i \in \mathcal{C}_{\text{active}}}$ where $s_i^{(j)} \in \mathcal{R}$ for all $i \in \mathcal{C}_{\text{active}}$ and for all $j \in [N]$;
 - A semi-corrupted pair $\{P_{j_1}, P_{j_2}\}$ where $\{P_{j_1}, P_{j_2}\} \cap \mathcal{C}_{\text{active}} \neq \emptyset$.
2. Based on what it has received, $\mathcal{F}_{\text{RandZeroSh}}$ does one of the following:
 - On receiving $(s_i^{(j)})_{i \in \mathcal{C}_{\text{active}}}$ for all $j \in [N]$, $\mathcal{F}_{\text{RandZeroSh}}$ generates N degree- d Shamir sharings $[(o_1^{(j)}, o_2^{(j)}, \dots, o_h^{(j)})]_d$ such that the shares of $P_i \in \mathcal{C}_{\text{active}}$ are $(s_i^{(j)})_{j=1}^N$. For each $P_i \in \mathcal{H}_{\text{active}}$, $\mathcal{F}_{\text{RandZeroSh}}$ sends the i -th share of $[(o_1^{(j)}, o_2^{(j)}, \dots, o_h^{(j)})]_d$ to P_i for all $j \in [N]$.
 - On receiving $\{P_{j_1}, P_{j_2}\}$, $\mathcal{F}_{\text{RandZeroSh}}$ sends $\{P_{j_1}, P_{j_2}\}$ to all active honest parties.

We use $\mathcal{F}_{\text{RandRobust}}$ to instantiate $\mathcal{F}_{\text{RandZeroSh}}$. Consider a secret sharing scheme Σ which outputs a degree- d packed Shamir sharing of (o_1, o_2, \dots, o_h) in \mathcal{R} , i.e. $[(o_1, o_2, \dots, o_h)]_d$. Let V denote the set of all elements o in \mathcal{R} such that $\psi(o) = \mathbf{0} \in (\mathbb{Z}/2^k\mathbb{Z})^\ell$. Recall that $\psi(\cdot)$ is $\mathbb{Z}/2^k\mathbb{Z}$ -linear. Therefore, for all $o_1, o_2 \in V$ and for all $\alpha, \beta \in \mathbb{Z}/2^k\mathbb{Z}$, we have

$$\psi(\alpha \cdot o_1 + \beta \cdot o_2) = \alpha \cdot \psi(o_1) + \beta \cdot \psi(o_2) = \mathbf{0}$$

which implies that $\alpha \cdot o_1 + \beta \cdot o_2 \in V$. Therefore, V is a $\mathbb{Z}/2^k\mathbb{Z}$ -subspace. Since packed Shamir sharings are linear in \mathcal{R} , Σ is an $\mathbb{Z}/2^k\mathbb{Z}$ -GLSSS.

Also, there exists an efficient algorithm to sample a uniformly random element in V . Let $\mathbf{1}$ be a vector in $(\mathbb{Z}/2^k\mathbb{Z})^\ell$ such that all entries are 1 in $\mathbb{Z}/2^k\mathbb{Z}$. The algorithm random samples $x \in \mathcal{R}$ and computes $\mathbf{y} := \psi(x)$. Then, the algorithm outputs $x - \phi(\mathbf{1}) \cdot \phi(\mathbf{y})$. Note that $\psi(x - \phi(\mathbf{1}) \cdot \phi(\mathbf{y})) = \psi(x) - \mathbf{1} \star \mathbf{y} = \mathbf{0}$. Therefore, $x - \phi(\mathbf{1}) \cdot \phi(\mathbf{y}) \in V$. Also note that for all $o \in V$, o will be output by the algorithm only when x satisfies that $o + \phi(\mathbf{1}) \cdot \phi(\mathbf{y}) = x$ for some $\mathbf{y} \in (\mathbb{Z}/2^k\mathbb{Z})^\ell$. Because x is uniformly random in \mathcal{R} , the probability that the algorithm outputs o is $|(\mathbb{Z}/2^k\mathbb{Z})^\ell|/|\mathcal{R}|$. Therefore, the algorithm outputs a uniformly random element in V .

To use $\mathcal{F}_{\text{RandRobust}}$ to prepare random sharings in Σ , we need to show that:

- Given a set $A \subset \mathcal{I}$ and a set of shares $\{a_i\}_{i \in A}$ for parties in A , there exists an efficient algorithm which outputs either $\Sigma(A, (a_i)_{i \in A}) = \emptyset$ or a random sharing $[(o_1, o_2, \dots, o_h)]_d \in \Sigma(A, (a_i)_{i \in A})$.

To this property, depending on the size of A , there are two cases:

- If $|A| \geq d + 1$, by the property of packed Shamir secret sharing scheme, the set of shares $\{a_i\}_{i \in A}$ can fully determine the whole sharing if exists. The algorithm checks whether the shares and the secrets lie on a polynomial $f(\cdot)$ of degree at most d . If not, the algorithm outputs $\Sigma(A, (a_i)_{i \in A}) = \emptyset$. Otherwise, the algorithm outputs the Shamir sharing determined by $\{a_i\}_{i \in A}$.
- If $d - h + 2 \leq |A| \leq d$, the set of shares $\{a_i\}_{i \in A}$ may not be independent because of the structure of the secret space. Recall that the secret space is $V \subset \mathcal{R}^h$, and V is \mathcal{R} -linear. Suppose V is defined by the linear mapping $B : (\mathbb{Z}/2^k\mathbb{Z})^{m \cdot h} \rightarrow (\mathbb{Z}/2^k\mathbb{Z})^{m \cdot h}$. Since packed Shamir secret sharing scheme is also \mathcal{R} -linear, given the shares $\{a_i\}_{i \in A}$ the secret space can be defined as the image of the mapping $\mathcal{L} : (\mathbb{Z}/2^k\mathbb{Z})^{(d - |A| + 2) \cdot h} \rightarrow (\mathbb{Z}/2^k\mathbb{Z})^{m \cdot h}$, and \mathcal{L} can be written in the form that $\mathcal{L}(\mathbf{a}) = C \cdot \mathbf{a} + \mathbf{b}$, where the matrix $C \in (\mathbb{Z}/2^k\mathbb{Z})^{(m \cdot h) \times ((d - |A| + 2) \cdot h)}$ and the vector $\mathbf{b} \in (\mathbb{Z}/2^k\mathbb{Z})^{m \cdot h}$ are determined by the shares $\{a_i\}_{i \in A}$. It is clear that $\Sigma(A, (a_i)_{i \in A}) \neq \emptyset$ if and only if $B \cdot \mathbf{x} = C \cdot \mathbf{y} + \mathbf{b}$ has solution. This can be determined easily by solving if the following linear equation has solutions

$$(B, C) \cdot \begin{pmatrix} -\mathbf{x} \\ \mathbf{y} \end{pmatrix} = \mathbf{b},$$

and this can be done by solving the null space of the matrix (B, C) . If the null space is empty, the algorithm outputs $\Sigma(A, (a_i)_{i \in A}) = \emptyset$. Otherwise, the algorithm picks random $(-\mathbf{x}_0, \mathbf{y}_0)^T$ in the null space. Suppose the secrets corresponding to the sampled vector \mathbf{y}_0 are o_1, o_2, \dots, o_h . Then, based on the secrets o_1, o_2, \dots, o_h and the shares $\{a_i\}_{i \in A}$, the algorithm reconstructs the whole packed Shamir sharing $[(o_1, o_2, \dots, o_h)]_d$ and outputs $[(o_1, o_2, \dots, o_h)]_d$.

- If $|A| \leq d - h + 1$, by the property of packed Shamir secret sharing scheme, the set of shares $\{a_i\}_{i \in A}$ is independent, so $\Sigma(A, (a_i)_{i \in A}) \neq \emptyset$. The algorithm randomly samples $(d + 1 - h - |A|)$ elements in \mathcal{R} as the shares of the first $(d + 1 - h - |A|)$ parties in $\mathcal{I} \setminus A$. Then, based on the secrets o_1, \dots, o_h , the shares of the first $(d + 1 - h - |A|)$ parties in $\mathcal{I} \setminus A$ and the shares $\{a_i\}_{i \in A}$, the algorithm reconstructs the whole packed Shamir sharing $[(o_1, o_2, \dots, o_h)]_d$ and outputs $[(o_1, o_2, \dots, o_h)]_d$.

Therefore, $\mathcal{F}_{\text{RandZeroSh}}$ can be instantiated by $\mathcal{F}_{\text{RandRobust}}$ with the secret sharing scheme Σ defined above. Note that the share size of Σ is $\text{sh} = m$ elements in $\mathbb{Z}/2^k\mathbb{Z}$. It follows that when using $\Pi_{\text{RandRobust}}(N)$ to instantiate $\mathcal{F}_{\text{RandRobust}}$ the total communication complexity to generate N random sharings in Σ is $O(N \cdot n \cdot m + n^2 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

Random Parity Sharings. The functionality $\mathcal{F}_{\text{RandParity}}$ that prepares random parity sharings is shown in Functionality 13.

Functionality 13: $\mathcal{F}_{\text{RandParity}}(N)$

1. $\mathcal{F}_{\text{RandParity}}$ receives from the adversary a set of even number of parties $S \subset \mathcal{P}_{\text{active}}$ such that $|S|/2 \leq |S \cap \mathcal{C}_{\text{active}}|$. $\mathcal{F}_{\text{RandParity}}$ updates $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S$. $\mathcal{F}_{\text{RandParity}}$ then sends S to all active honest parties.
2. $\mathcal{F}_{\text{RandParity}}$ receives from the adversary the shares of active corrupted parties $(u_i^{(j)})_{i \in \mathcal{C}_{\text{active}}}$ where $s_i^{(j)} \in \mathcal{R}$ for all $i \in \mathcal{C}_{\text{active}}$ and for all $j \in [N]$.
3. $\mathcal{F}_{\text{RandParity}}$ randomly samples N parity elements $p_1, \dots, p_N \in \mathcal{R}$. Then $\mathcal{F}_{\text{RandParity}}$ generates N degree- t Shamir sharing $([p_j]_t)_{j=1}^N$ such that the shares of $P_i \in \mathcal{C}_{\text{active}}$ are $(u_i^{(j)})_{j=1}^N$. For each $P_h \in \mathcal{H}_{\text{active}}$, $\mathcal{F}_{\text{RandParity}}$ sends the i -th share of $[p_j]_t$ to P_h for all $j \in [N]$.

Next we show that we can use Π_{Rand} to instantiate $\mathcal{F}_{\text{RandRobust}}$. Consider a secret sharing scheme Σ which takes a parity element p and outputs a Shamir sharing of p in \mathcal{R} .

It is shown in [PS21] that the following properties hold for Σ :

- Σ is an $\mathbb{Z}/2^k\mathbb{Z}$ -GLSSS.
- Given a set $A \subset \mathcal{I}$ and a set of shares $\{a_i\}_{i \in A}$ for parties in A , there exists an efficient algorithm which outputs either $\Sigma(A, (a_i)_{i \in A}) = \emptyset$ or a random sharing $[p]_t \in \Sigma(A, (a_i)_{i \in A})$.

Therefore, $\mathcal{F}_{\text{RandParity}}$ can be instantiated by $\mathcal{F}_{\text{RandRobust}}$ with the secret sharing scheme Σ defined above. Note that the share size of Σ is $\text{sh} = m$ elements in $\mathbb{Z}/2^k\mathbb{Z}$. It follows that when using $\Pi_{\text{Rand}}(N)$ to instantiate $\mathcal{F}_{\text{RandRobust}}$ the total communication complexity to generate N random sharings in Σ is $O(N \cdot n \cdot m + n^2 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

Random Additive Beaver Triples. The functionality $\mathcal{F}_{\text{TripleAdd}}$ that prepares random additive Beaver triples is shown in Functionality 14.

Functionality 14: $\mathcal{F}_{\text{TripleAdd}}(N)$

1. $\mathcal{F}_{\text{TripleAdd}}$ receives from the adversary a set of even number of parties $S \subset \mathcal{P}_{\text{active}}$ such that $|S|/2 \leq |S \cap \mathcal{C}_{\text{active}}|$. $\mathcal{F}_{\text{TripleAdd}}$ updates $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S$. $\mathcal{F}_{\text{TripleAdd}}$ then sends S to all active honest parties.
2. $\mathcal{F}_{\text{TripleAdd}}$ receives from the adversary the shares of active corrupted parties $(u_i^{(j)}, v_i^{(j)}, w_i^{(j)})_{i \in \mathcal{C}_{\text{active}}}$ where $u_i^{(j)}, v_i^{(j)}, w_i^{(j)} \in \mathcal{R}$ and $w_i^{(j)} = u_i^{(j)} + v_i^{(j)}$ for all $i \in \mathcal{C}_{\text{active}}$ and for all $j \in [N]$.
3. $\mathcal{F}_{\text{TripleAdd}}$ randomly samples N random triples $(\mathbf{a}_1, \mathbf{b}_1, \mathbf{c}_1), \dots, (\mathbf{a}_N, \mathbf{b}_N, \mathbf{c}_N)$, where each triple entry is in $(\mathbb{Z}/2^k\mathbb{Z})^\ell$ and each triple $(\mathbf{a}_j, \mathbf{b}_j, \mathbf{c}_j)$ satisfies $\mathbf{a}_j + \mathbf{b}_j = \mathbf{c}_j$. Then $\mathcal{F}_{\text{TripleAdd}}$ generates $3N$ degree- t Shamir sharing $([\phi(\mathbf{a}_j)]_t, [\phi(\mathbf{b}_j)]_t, [\phi(\mathbf{c}_j)]_t)_{j=1}^N$ such that the shares of $P_i \in \mathcal{C}_{\text{active}}$ are $(u_i^{(j)}, v_i^{(j)}, w_i^{(j)})_{j=1}^N$. For each $P_i \in \mathcal{H}_{\text{active}}$, $\mathcal{F}_{\text{TripleAdd}}$ sends the i -th share of the sharings in the triple $([\phi(\mathbf{a}_j)]_t, [\phi(\mathbf{b}_j)]_t, [\phi(\mathbf{c}_j)]_t)$ to P_i for all $j \in [N]$.

The protocol $\Pi_{\text{TripleAdd}}$ realizes $\mathcal{F}_{\text{TripleAdd}}$. The description of $\Pi_{\text{TripleAdd}}$ appears in Protocol 9.

Protocol 9: $\Pi_{\text{TripleAdd}}(N)$

1. All parties invoke $\mathcal{F}_{\text{RandShamir}}(2N)$ to prepare $2N$ regular degree- t Shamir secret sharings in the form of $([\phi(\mathbf{a}_i)]_t, [\phi(\mathbf{b}_i)]_t)_{i=1}^N$. All parties also receive a set of eliminated parties denoted by S . All parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S$.
2. All parties locally compute $[\phi(\mathbf{c}_i)]_t = [\phi(\mathbf{a}_i)]_t + [\phi(\mathbf{b}_i)]_t$ for all $i \in [N]$. All parties output S and $([\phi(\mathbf{a}_i)]_t, [\phi(\mathbf{b}_i)]_t, [\phi(\mathbf{c}_i)]_t)_{i=1}^N$.

The communication complexity of calling $\mathcal{F}_{\text{RandShamir}}(2N)$ is $O(N \cdot n \cdot m + n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$. Therefore, the communication complexity of $\Pi_{\text{TripleAdd}}(N)$ is $O(N \cdot n \cdot m + n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

Lemma 8. *The protocol $\Pi_{\text{TripleAdd}}$ computes $\mathcal{F}_{\text{TripleAdd}}$ with perfect security in the $\mathcal{F}_{\text{RandShamir}}$ -hybrid model when $|\mathcal{C}_{\text{active}}| < |\mathcal{P}_{\text{active}}|/3$.*

Proof. Let \mathcal{A} denote the adversary. We will construct a simulator Sim to simulate the behaviors of honest parties for the adversary. In Step 1, Sim faithfully emulates $\mathcal{F}_{\text{RandShamir}}$. Because there is no other communication, the distribution of the real world is the same as the distribution of the ideal world.

Random Multiplicative Beaver Triples. The description of $\mathcal{F}_{\text{TripleMult}}$ appears in Functionality 15.

Functionality 15: $\mathcal{F}_{\text{TripleMult}}(N)$

1. $\mathcal{F}_{\text{TripleMult}}$ receives from the adversary a set of even number of parties $S \subset \mathcal{P}_{\text{active}}$ such that $|S|/2 \leq |S \cap \mathcal{C}_{\text{active}}|$. $\mathcal{F}_{\text{TripleMult}}$ updates $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S$. $\mathcal{F}_{\text{TripleMult}}$ then sends S to all active honest parties.
2. $\mathcal{F}_{\text{TripleMult}}$ receives from the adversary the shares of active corrupted parties $(u_i^{(j)}, v_i^{(j)}, w_i^{(j)})_{i \in \mathcal{C}_{\text{active}}}$ where $u_i^{(j)}, v_i^{(j)}, w_i^{(j)} \in \mathcal{R}$ for all $i \in \mathcal{C}_{\text{active}}$ and for all $j \in [N]$.
3. $\mathcal{F}_{\text{TripleMult}}$ randomly samples N random triples $(\mathbf{a}_1, \mathbf{b}_1, \mathbf{c}_1), \dots, (\mathbf{a}_N, \mathbf{b}_N, \mathbf{c}_N)$, where each triple entry is in $(\mathbb{Z}/2^k\mathbb{Z})^\ell$ and each triple $(\mathbf{a}_j, \mathbf{b}_j, \mathbf{c}_j)$ satisfies $\mathbf{a}_j \star \mathbf{b}_j = \mathbf{c}_j$. Then $\mathcal{F}_{\text{TripleMult}}$ generates $3N$ degree- t Shamir sharing $([\phi(\mathbf{a}_j)]_t, [\phi(\mathbf{b}_j)]_t, [\phi(\mathbf{c}_j)]_t)_{j=1}^N$ such that the shares of $P_i \in \mathcal{C}_{\text{active}}$ are $(u_i^{(j)}, v_i^{(j)}, w_i^{(j)})_{j=1}^N$. For each $P_h \in \mathcal{H}_{\text{active}}$, $\mathcal{F}_{\text{TripleMult}}$ sends the ℓ -th share of the sharings in the triple $([\phi(\mathbf{a}_j)]_t, [\phi(\mathbf{b}_j)]_t, [\phi(\mathbf{c}_j)]_t)$ to P_h for all $j \in [N]$.

The protocol $\Pi_{\text{TripleMult}}$ realizes $\mathcal{F}_{\text{TripleMult}}$. The description of $\Pi_{\text{TripleMult}}$ appears in Protocol 10.

Protocol 10: $\Pi_{\text{TripleMult}}(N)$

1. All parties invoke $\mathcal{F}_{\text{RandShamir}}(2N)$ to prepare $2N$ regular degree- t Shamir secret sharings in the form of $([\phi(\mathbf{a}_i)]_t, [\phi(\mathbf{b}_i)]_t)_{i=1}^N$. All parties also receive a set of eliminated parties denoted by S_1 . All parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_1$.
2. All parties invoke $\mathcal{F}_{\text{Mult}}(N)$ with input $([\phi(\mathbf{a}_i)]_t)_{i=1}^N$ and $([\phi(\mathbf{b}_i)]_t)_{i=1}^N$. They get the output $([z_i]_t)_{i=1}^N$ and a set of eliminated parties denoted by S_2 . All parties set $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_2$. We have for all $i \in [N]$, $z_i = \phi(\mathbf{a}_i) \cdot \phi(\mathbf{b}_i)$.
3. All parties invoke $\mathcal{F}_{\text{ReEncode}}(N)$ with inputs $([z_i]_t)_{i=1}^N$. They get the output $([\phi \circ \psi(z_i)]_t)_{i=1}^N$ and a set of eliminated parties denoted by S_3 . All parties set $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_3$. Let \mathbf{c}_i denote $\psi(z_i)$ for all $i \in [N]$. All parties take $S := S_1 \cup S_2 \cup S_3$ and $([\phi(\mathbf{a}_i)]_t, [\phi(\mathbf{b}_i)]_t, [\phi(\mathbf{c}_i)]_t)_{i=1}^N$ as output.

The communication complexity of calling $\mathcal{F}_{\text{RandShamir}}(2N)$ is $O(N \cdot n \cdot m + n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$. The communication complexity of the instantiation of $\mathcal{F}_{\text{Mult}}(N)$ is $O(N \cdot n \cdot m + n^3 \cdot m)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$, and the communication complexity of the instantiation of $\mathcal{F}_{\text{ReEncode}}(N)$ is $O(N \cdot n \cdot m + n^3 \cdot m)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$. Therefore, the communication complexity of $\Pi_{\text{TripleMult}}(N)$ is $O(N \cdot n \cdot m + n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

Lemma 9. *The protocol $\Pi_{\text{TripleMult}}$ computes $\mathcal{F}_{\text{TripleMult}}$ with perfect security in the $(\mathcal{F}_{\text{RandShamir}}, \mathcal{F}_{\text{Mult}}, \mathcal{F}_{\text{ReEncode}})$ -hybrid model when $|\mathcal{C}_{\text{active}}| < |\mathcal{P}_{\text{active}}|/3$.*

Proof. Let \mathcal{A} denote the adversary. We will construct a simulator Sim to simulate the behaviors of honest parties for the adversary. In Step 1, Sim faithfully emulates $\mathcal{F}_{\text{RandShamir}}$. Because there is no other communication, the distribution of the real world is the same as the distribution of the ideal world.

C Extended Additive Secret Sharing

Let (ϕ, ψ) be an $(\ell, m)_2$ -RMFE. Recall that n denotes the number of parties and $\phi : (\mathbb{Z}/2^k\mathbb{Z})^\ell \rightarrow \mathcal{R}$ is an $\mathbb{Z}/2^k\mathbb{Z}$ -linear map. Also, $|\mathcal{R}| = 2^m \geq 2n + 1$. Therefore, the Shamir secret sharing scheme is well-defined in \mathcal{R} . In our construction, we will use ϕ to encode a vector of secrets $\mathbf{x} = (x^{(1)}, x^{(2)}, \dots, x^{(\ell)}) \in (\mathbb{Z}/2^k\mathbb{Z})^\ell$. All parties will hold a degree- t Shamir sharing of $\phi(\mathbf{x})$, denoted by $[\phi(\mathbf{x})]_t$.

Defining Additive Sharings and Extended Additive Sharings. For $x \in \mathbb{Z}/2^k\mathbb{Z}$, we use $\langle x \rangle$ to denote an *additive sharing* of x among the first $t + 1$ parties in $\mathbb{Z}/2^k\mathbb{Z}$. Specifically, the additive sharing of x is $\langle x \rangle = (x_1, \dots, x_{n'})$ where party P_i holds the share $x_i \in \mathbb{Z}/2^k\mathbb{Z}$ such that $\sum_{i=1}^{t+1} x_i = x$ and $x_{t+2}, \dots, x_{n'}$ are all 0.

Recall that $\psi : \mathcal{R} \rightarrow (\mathbb{Z}/2^k\mathbb{Z})^\ell$ and $\text{val}(\cdot) : \mathcal{R} \rightarrow \mathbb{Z}/2^k\mathbb{Z}$ are both $\mathbb{Z}/2^k\mathbb{Z}$ -linear. We define that a sharing $\langle x \rangle$ is an *extended additive sharing* of $x \in \mathbb{Z}/2^k\mathbb{Z}$ if $\langle x \rangle$'s shares form a degree- t Shamir sharing $[y]_t$ in \mathcal{R} . The secret y satisfies $\text{val}(y) = x$. We write $\langle x \rangle := [y]_t$. It is clear that extended additive sharings are additive.

Generating Extended Additive Sharings. We generate generate the extended additive sharing of $x \in \mathbb{Z}/2^k\mathbb{Z}$ from the Shamir sharing $[\phi(\mathbf{z})]_t$. Suppose the j -th element of \mathbf{z} is x . By the property of RMFE, we have $\psi(\phi(\mathbf{e}_j) \cdot \phi(\mathbf{z})) = \mathbf{e}_j \cdot \mathbf{z}$. Therefore, $\text{val}(\phi(\mathbf{e}_j) \cdot \phi(\mathbf{z})) = x$. To obtain $\langle x \rangle$, all parties locally compute $\langle x \rangle = \phi(\mathbf{e}_j) \cdot [z]_t$.

Converting Extended Additive Sharings to Additive Sharings. Now we show how to obtain $\langle x \rangle$ from $\langle x \rangle$. We suppose that $\langle x \rangle$ is equivalent to the Shamir sharing $[y]_t$.

Recall that the degree- t Shamir sharing $[y]_t$ corresponds to a degree- t polynomial $f(\cdot) \in \mathcal{R}[X]$ such that $f(\alpha_i)$ is the share of the i -th party P_i and $f(0) = y$, where $\alpha_1, \dots, \alpha_n$ are distinct non-zero elements in \mathcal{R} . In particular, relying on Lagrange interpolation, $f(0)$ can be written as a linear combination of the first $t + 1$ shares. For $i \in [t + 1]$, let $c_i = \prod_{j \neq i, j \in [t+1]} \frac{\alpha_j}{\alpha_j - \alpha_i}$. We have $f(0) = \sum_{i=1}^{t+1} c_i f(\alpha_i)$. Therefore, the Shamir sharing $[y]_t$ can be locally converted to an additive sharing of x among the first $t + 1$ parties by letting $P_i (i \in [t + 1])$ taking $\text{val}(c_i \cdot f(\alpha_i))$ as its share, while for $i \notin [t + 1]$ the party P_i takes 0 as its share. Note that

$$\sum_{i=1}^{t+1} \text{val}(c_i \cdot f(\alpha_i)) = \text{val}\left(\sum_{i=1}^{t+1} c_i \cdot f(\alpha_i)\right) = \text{val}(y) = x,$$

this is indeed an additive sharing of x .

D Segment Verification Protocol

Protocol 11: $\Pi_{\text{Verify}}(\text{seg}, S_0)$

1. Suppose seg has N gate groups. For the i -th gate group within seg that corresponds to the Beaver triple $([\phi(\mathbf{a}^{(i)})]_t, [\phi(\mathbf{b}^{(i)})]_t, [\phi(\mathbf{c}^{(i)})]_t)$ and has inputs $\mathbf{x}^{(i)}, \mathbf{z}^{(i)} \in (\mathbb{Z}/2^k\mathbb{Z})^\ell$. All parties locally compute $\phi(\mathbf{x}^{(i)} + \mathbf{a}^{(i)})$ and $\phi(\mathbf{z}^{(i)} + \mathbf{b}^{(i)})$ for all $i \in [N]$.
2. All parties call $\mathcal{F}_{\text{VerifyBC}}$ with inputs $\{\phi(\mathbf{x}^{(i)} + \mathbf{a}^{(i)})\}_{i \in [N]} \cup \{\phi(\mathbf{z}^{(i)} + \mathbf{b}^{(i)})\}_{i \in [N]}$ and D . If the result is a semi-honest pair $\{P_{j_1}, P_{j_2}\}$, all parties take it as output and halt. Otherwise, run the following steps.
3. All parties locally compute $[\phi(\mathbf{x}^{(i)})]_t = \phi(\mathbf{x}^{(i)} + \mathbf{a}^{(i)}) - [\phi(\mathbf{a}^{(i)})]_t$ and $[\phi(\mathbf{z}^{(i)})]_t = \phi(\mathbf{z}^{(i)} + \mathbf{b}^{(i)}) - [\phi(\mathbf{b}^{(i)})]_t$ for all $i \in [N]$.

4. For each additive gate group with input $[\phi(\mathbf{x}^{(i)})]_t$ and $[\phi(\mathbf{y}^{(i)})]_t$, all parties locally compute the output sharing

$$[\phi(\mathbf{z}^{(i)})]_t = [\phi(\mathbf{x}^{(i)})]_t + [\phi(\mathbf{y}^{(i)})]_t.$$

For all the multiplicative gate groups, suppose their indices form the set I_{mult} . All parties call $\mathcal{F}_{\text{Mult}}$ with input $([\phi(\mathbf{x}^{(i)})]_t)_{i \in I_{\text{mult}}}$ and $([\phi(\mathbf{y}^{(i)})]_t)_{i \in I_{\text{mult}}}$, and get output $([w^{(i)}]_t)_{i \in I_{\text{mult}}}$ and a set of eliminated parties denoted by S_1 . All parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_1$.

Then all the parties call the functionality $\mathcal{F}_{\text{ReEncode}}$ with input $([w^{(i)}]_t)_{i \in I_{\text{mult}}}$ and get the output sharings $([\phi(\mathbf{z}^{(i)})]_t)_{i \in I_{\text{mult}}}$ and a set of eliminated parties denoted by S_2 . All parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_2$.

5. All parties run the protocol $\Pi_{\text{NetworkRouting}}(\text{seg})$. Note that during fan-out the j -th wire of $[\phi(\mathbf{z}^{(i)})]_t$ is copied $\tilde{n}_j^{(i)}$ times. All parties get all Shamir sharings of seg 's gate group inputs, denoted by $\{[\phi(\tilde{\mathbf{x}}^{(i)})]_t\}_{i \in [N]}$ and $\{[\phi(\tilde{\mathbf{y}}^{(i)})]_t\}_{i \in [N]}$. All parties also receive a set of eliminated parties denoted by S_3 . All parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_3$.
6. All parties call $\mathcal{F}_{\text{OpenPub}}$ to reconstruct the Shamir sharings $\{[\phi(\tilde{\mathbf{x}}^{(i)})]_t + [\phi(\mathbf{a}^{(i)})]_t\}_{i \in [N]}$ and $\{[\phi(\tilde{\mathbf{y}}^{(i)})]_t + [\phi(\mathbf{b}^{(i)})]_t\}_{i \in [N]}$, and get $\tilde{\mathbf{x}}^{(i)} + \mathbf{a}^{(i)}$ and $\tilde{\mathbf{y}}^{(i)} + \mathbf{b}^{(i)}$ for all $i \in [N]$.
7. Each party locally compares $\tilde{\mathbf{x}}^{(i)} + \mathbf{a}^{(i)}$ with $\mathbf{x}^{(i)} + \mathbf{a}^{(i)}$ and compares $\tilde{\mathbf{y}}^{(i)} + \mathbf{b}^{(i)}$ with $\mathbf{y}^{(i)} + \mathbf{b}^{(i)}$ for all $i \in [N]$. Based on whether there are any differences, all parties do the following:
- (a) If for all $i \in [N]$, $\tilde{\mathbf{x}}^{(i)} + \mathbf{a}^{(i)} = \mathbf{x}^{(i)} + \mathbf{a}^{(i)}$ and $\tilde{\mathbf{y}}^{(i)} + \mathbf{b}^{(i)} = \mathbf{y}^{(i)} + \mathbf{b}^{(i)}$, all parties call $\mathcal{F}_{\text{FanOut}}$ with the output sharings $\{[\phi(\mathbf{z}^{(i)})]_t\}_{i \in [N]}$ and $\{(n_j^{(i)} - \tilde{n}_j^{(i)})\}_{i \in [N], j \in [\ell]}$, where the j -th wire of $[\phi(\mathbf{z}^{(i)})]_t$ is copied $n_j^{(i)} - \tilde{n}_j^{(i)}$ times. All parties receive the fan-out sharings, and a set of eliminated parties denoted by S_4 . All parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_4$. Then all parties run $\mathcal{F}_{\text{Permute}}$ with the fan-out sharings and the desired permutations as input. All parties receive the permuted fan-out sharings, and a set of eliminated parties denoted by S_5 . All parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_5$. All parties output **correct**.
- (b) Otherwise, let $S := S_0 \cup S_1 \cup S_2 \cup S_3$. If $S \neq \emptyset$, all parties output S and **incorrect** and halts. Otherwise, all parties can select a wire with an inconsistent value that has the smallest topological order, denoted by $x_{j_0}^{(i_0)} + a_{j_0}^{(i_0)}$ or $y_{j_0}^{(i_0)} + b_{j_0}^{(i_0)}$. Let $\langle s \rangle + \langle o \rangle$ denote its corresponding additive sharing for reconstruction in the protocol Π_{Eval} , and let $\langle \mathfrak{s} \rangle$ denote its extended additive sharing. Then all parties run the protocol $\Pi_{\text{FaultDetection}}(D, \langle s \rangle + \langle o \rangle, \langle \mathfrak{s} \rangle)$, and get a set of eliminated parties S_4 as output. All parties update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S_4$, and output $S \cup S_4$ and **incorrect**.

Lemma 10 (Correctness of Π_{Verify}). *In Π_{Verify} , either all active honest parties have got the correct output Shamir sharings, or all parties output a set of eliminated parties where at least half of the parties in the set are corrupted.*

The proof of Lemma 10 can be found in Section F.2.

Cost of Π_{Verify} . The communication complexity of Step 2 for calling $\mathcal{F}_{\text{VerifyBC}}$ is $O(N \cdot n \cdot m + n^2 \cdot m)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$. The communication complexity of Step 4 for calling $\mathcal{F}_{\text{Mult}}$ and $\mathcal{F}_{\text{ReEncode}}$ is bounded by $O(N \cdot n \cdot m + n^3 \cdot m)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$. The communication complexity of Step 5 for running $\Pi_{\text{NetworkRouting}}$ is $O(N \cdot n \cdot m + n^2 \cdot m)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$, since all parties generate $\sum_{i \in [N], j \in [\ell]} \tilde{n}_j^{(i)} / \ell \leq 3N$ fan-out sharings in this step. The communication complexity of Step 6 for calling $\mathcal{F}_{\text{OpenPub}}$ is $O(N \cdot n \cdot m + n^2 \cdot m)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$. Note that $N = O(n^2 \cdot m)$ according to Theorem 4. Therefore, the total communication complexity before Step 7 is bounded by $O(n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

For Step 7, will discuss the cost of Π_{Verify} in two different scenarios based on whether the verification passes.

- If the verification passes, all parties will run $\mathcal{F}_{\text{FanOut}}$ and $\mathcal{F}_{\text{Permute}}$ that output

$$M_{\text{seg}} := \frac{\sum_{i \in [N], j \in [\ell]} n_j^{(i)} - \tilde{n}_j^{(i)}}{\ell}$$

Shamir sharings. Then the communication complexity of Step 7.(a) is $O(M_{\text{seg}} \cdot n \cdot m + n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$. Therefore, the total communication complexity in this case is $O(M_{\text{seg}} \cdot n \cdot m + n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

- If the verification does not pass, all parties will run $\Pi_{\text{FaultDetection}}$. And the communication complexity of running $\Pi_{\text{FaultDetection}}$ is $O(n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$. Therefore, the total communication complexity in this case is $O(n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

E Details on Network Routing

E.1 Circuit Transformation

The following theorem shows that we can safely assume that the circuit at hand satisfies the grouping properties for network routing.

Theorem 4 (Circuit Transformation). *Given an arithmetic circuit C over $\mathbb{Z}/2^k\mathbb{Z}$ with input coming from c clients, there exists an efficient algorithm which takes C as input and outputs an arithmetic circuit C' over $\mathbb{Z}/2^k\mathbb{Z}$ with the following properties:*

- For all input x , $C(x) = C'(x)$.
- C' satisfies the Circuit Segment Conditions, the Gate Number Conditions, and the Gate Grouping Conditions which are stated previously.
- Circuit Size: $|C'| = O(|C| + \ell \cdot c)$, where c is the number of clients that provide inputs.

Proof. Let $\text{Client}_1, \text{Client}_2, \dots, \text{Client}_c$ denote the clients that provide inputs. We show the algorithm to transform the circuit C' from C as follows:

- **Step 1:** We start by creating a virtual client Client_0 . We will insert two new types of gates: input gates belonging to Client_0 and output gates belonging to Client_0 . The input gates belonging to Client_0 are used to provide constant values for the computation. The output gates belonging to Client_0 are used to collect the wires that will not be output to any clients. Initially, for each input wire carrying a constant value in C , we insert an input gate belonging to Client_0 for this wire. Without loss of generality, we assume that each input of Client_0 is only used once in the circuit. To see this, an input of Client_0 is known to all parties, and if this input is used more than once, we can simply insert multiple input gates that take the same value. Note that the same argument does not work for other clients since it allows a client to use different values for its input wires in C' which should have the same value in C .

This step does not affect the circuit size.

- **Step 2:** Order all the addition gates and multiplication gates in C according to their topological ordering. We divide $n \cdot \ell$ consecutive gates into one circuit segment. For each circuit segment, let M_i and A_i denote the number of multiplication gates and the number of addition gates in it. We insert $\lceil \frac{M_i}{\ell} \rceil \cdot \ell - M_i$ dummy multiplication gates and $\lceil \frac{A_i}{\ell} \rceil \cdot \ell - A_i$ dummy addition gates in this circuit segment. Each of these new dummy gates takes two inputs from Client_0 and outputs to Client_0 as well. The inputs are set to be 0. In this way, we insert 2 input gates belonging to Client_0 and one output gate belonging to Client_0 for each of the new dummy multiplication gates and dummy addition gates. Then we organize ℓ consecutive addition gates into one addition gate group, and organize ℓ consecutive multiplication gates into one multiplication gate group. After this step, the circuit is divided into circuit segments where each circuit segment has size $\Theta(n^2 \cdot m^2)$ except that the last circuit segment has size $O(n^2 \cdot m^2)$. Also, the multiplication gates and the addition gates within each circuit are organized into gate groups of size ℓ . Note that this step increases the circuit size by $O(|C|/n)$.
- **Step 3:** For each client Client_i ($i = 0, 1, \dots, c$), let I_i denote the number of input gates belonging to Client_i . Then we insert $\lceil \frac{I_i}{\ell} \rceil \cdot \ell - I_i$ dummy input gates belonging to Client_i , which will take 0 as input. We also insert the same number of dummy output gates belonging to Client_0 , where each dummy output gate directly takes the input of one dummy input gate as output. Then we divide the input gates belonging to Client_i into groups of ℓ , and Client_i has $\lceil \frac{I_i}{\ell} \rceil$ input gate groups.

After this step, the number of input gates belonging to each Client_i is a multiple of ℓ . Note that this step increases the circuit size by $O(\ell \cdot c)$.

- **Step 4:** For each gate group (i.e. input gate group, multiplication gate group, or addition gate group), let $(w_1, w_2, \dots, w_\ell)$ denote the output wires, and $(n_1, n_2, \dots, n_\ell)$ denote the number of times that each output wire is used. Note that an output wire in a multiplication or addition

gate group can be the input wire of some gate in the same gate group because we do not require that the gates in a multiplication or addition gate group should belong to the same circuit layer. Let $n_0 = \lceil \frac{\sum_{i=1}^{\ell} n_i}{\ell} \rceil \cdot \ell - \sum_{i=1}^{\ell} n_i$. We insert n_0 dummy output gates belonging to Client_0 , and these output gates take w_1 as the input wire.

After this step, the number of times that w_1, w_2, \dots, w_ℓ are used as input wires in other gates is a multiple of ℓ . Note that each wire w_j is used at least once, so this step increases the circuit size by at most a factor of 2.

- **Step 5:** For each client Client_i ($i = 0, 1, \dots, c$), let O_i denote the number of output gates belonging to Client_i . We insert $\lceil \frac{O_i}{\ell} \rceil \cdot \ell - O_i$ input gates belonging to Client_0 , which take 0 as input. We also insert the same number of output gates belonging to Client_i which take these inputs as output. Then we divide the output gates belonging to each client Client_i into gate groups of size ℓ , and we also divide the added input gates for Client_0 into gate groups of size ℓ . After this step, the number of output gates belonging to each Client_i is a multiple of ℓ . Also, after this step, the number of input gates belonging to Client_0 remains a multiple of ℓ . To see this, recall that after Step 2 and Step 3, the total number of output wires of all clients is a multiple of ℓ , and the number of output wires of Client_0 is a multiple of ℓ . In addition, since all the added dummy input gates for Client_0 in this step will be used only once, the number of times that the output wires of each gate group formed by these input gates are used is a multiple of ℓ . Note that this step increases the circuit size by $O(\ell \cdot c)$.

Next, we show that the circuit C' after the transformation has the desired properties:

- First we show $C(x) = C'(x)$ for all input x . Note that we do not change the original gates and wires in C . Therefore, it is sufficient to show that the gates and wires we add in the transformation do not change the functionality. For the input layer, we create several new input gates belonging to each Client_i ($i \geq 1$). These new gates directly connect to the output gates of Client_0 , which means that these values are never used in the computation. For the output layer, we create several new output gates belonging to each Client_i ($i \geq 1$). These new gates directly take the value 0 from the input gates of Client_0 , which do not affect the final result. Therefore, $C(x) = C'(x)$ for all input x .
- Next we show C' satisfies the Circuit Segment Conditions, the Gate Number Conditions, and the Gate Grouping Conditions.
 - After Step 2, the size of each segment except the last segment is $\Theta(n^2 \cdot m^2)$, and the size of the last segment is $O(n^2 \cdot m^2)$. Also, we divide the segment so that each segment contains gates of consecutive topological order. Therefore, C' satisfies the Circuit Segment Conditions.
 - After Step 2, the number of addition gates and the number of multiplication gates within each circuit segment are multiples of ℓ . After Step 1, Step 3, and Step 5, the number of input gates and the number of output gates belonging to each client are multiples of ℓ . Therefore, C' satisfies the Gate Number Conditions.
 - After Step 2, the addition gates and the multiplication gates within each segment are organized into gate groups of size ℓ . After Step 3 and Step 5, the input gates and the output gates belonging to each client are organized into gate groups of size ℓ . Also, after Step 4, for the output wires of each multiplication gate group and of each addition gate group, the number of times that those wires are used as input wires in other gates is a multiple of ℓ . After Step 4 and Step 5, for the output wires of each input gate group, the number of times that those wires are used as input wires in other gates is a multiple of ℓ . Therefore, C' satisfies the Gate Grouping Conditions.
- Step 1 will not affect the size of the circuit. In Step 2 and Step 4, the circuit size increases by at most a factor of 2. In Step 3 and Step 5, the circuit size increases by $O(\ell \cdot c)$. Therefore, the size of C' after the transformation is bounded by $O(|C| + \ell \cdot c)$.

E.2 Preparing Random Sharings for Sharing Transformations

Definition of Ideal Functionality. Let $r = \lfloor \frac{n'-2\ell+1}{2} \rfloor$. Let $\Pi_1, \Pi_2, \dots, \Pi_{\ell \cdot r}$ be $\ell \cdot r$ arbitrary $\mathbb{Z}/2^k\mathbb{Z}$ -arithmetic secret sharing schemes, where each Π_i is defined as follows:

- The secret space $Z_i = (\mathbb{Z}/2^k\mathbb{Z})^\ell$.

- The share space $U_i = (\mathcal{R})^2$, which can be viewed as $(\mathbb{Z}/2^k\mathbb{Z})^{2m}$.
- The sharing space $C_i = (\mathcal{R})^{2n'}$
- For a secret $\mathbf{x} \in Z_i$, the sharing of \mathbf{x} is $([\phi(\mathbf{x})]_t, [\phi(L_i(\mathbf{x}))]_t)$, where L_i is a linear mapping from the vector space $(\mathbb{Z}/2^k\mathbb{Z})^\ell$ to the vector space $(\mathbb{Z}/2^k\mathbb{Z})^\ell$.
- For a sharing $([\phi(\mathbf{x})]_t, [\phi(L_i(\mathbf{x}))]_t)$, we reconstruct the secret \mathbf{x} using $[\phi(\mathbf{x})]_t$.
- The share_{Π_i} can be viewed as $\text{share}_{\Pi_i} : Z \times (\mathbb{Z}/2^k\mathbb{Z})^{2t \cdot m} \rightarrow C$.

The functionality $\mathcal{F}_{\text{RandSharTrans}}$ presented in Functionality 16 generates a random sharing in an $\mathbb{Z}/2^k\mathbb{Z}$ -arithmetic secret sharing scheme $\Pi_i = (n', Z_i, U_i, C_i, \text{share}_{\Pi_i}, \text{rec}_{\Pi_i})$ for all $i \in [\ell \cdot r]$.

Functionality 16: $\mathcal{F}_{\text{RandSharTrans}}(N)$

1. $\mathcal{F}_{\text{RandSharTrans}}$ receives from the adversary one of the following:
 - The shares of active corrupted parties $(\mathbf{u}_j^{(i)})_{j \in \mathcal{C}_{\text{active}}}$ where $\Pi_i(\mathcal{C}_{\text{active}}, (\mathbf{u}_j^{(i)})_{j \in \mathcal{C}_{\text{active}}}) \neq \emptyset$ for all $i \in [\ell \cdot r]$;
 - A semi-corrupted pair $\{P_{j_1}, P_{j_2}\}$ where $\{P_{j_1}, P_{j_2}\} \cap \mathcal{C}_{\text{active}} \neq \emptyset$.
2. Based on what it has received, $\mathcal{F}_{\text{RandSharTrans}}$ does one of the following:
 - On receiving $(\mathbf{u}_j^{(i)})_{j \in \mathcal{C}_{\text{active}}}$ for all $i \in [\ell \cdot r]$, $\mathcal{F}_{\text{RandSharTrans}}$ uniformly randomly samples a random Π_i -sharing $[\mathbf{x}] \in \Pi(\mathcal{C}_{\text{active}}, (\mathbf{u}_j^{(i)})_{j \in \mathcal{C}_{\text{active}}})$ for all $i \in [\ell \cdot r]$. For each honest party $P_h \in \mathcal{P}_{\text{active}}$, $\mathcal{F}_{\text{RandSharTrans}}$ sends the ℓ -th share of $[\mathbf{x}]$ to P_h for all $i \in [\ell \cdot r]$.
 - On receiving $\{P_{j_1}, P_{j_2}\}$, $\mathcal{F}_{\text{RandSharTrans}}$ sends $\{P_{j_1}, P_{j_2}\}$ to all active honest parties.

Protocol Description. For $\mathbf{x} \in (\mathbb{Z}/2^k\mathbb{Z})^{r \cdot \ell} = (\mathbf{x}_1, \dots, \mathbf{x}_r)$ where each \mathbf{x}_i is a vector in $(\mathbb{Z}/2^k\mathbb{Z})^\ell$, let $\Phi(\mathbf{x})$ denote the vector $(\phi(\mathbf{x}_1), \dots, \phi(\mathbf{x}_r))$. For $\mathbf{y} \in \mathcal{R}^r = (y_1, \dots, y_r)$ where each y_i is an element in \mathcal{R} , let $\Psi(\mathbf{y})$ denote the vector $(\psi(y_1), \dots, \psi(y_r))$.

We have each secret in the secret space needs ℓ elements of randomness in $\mathbb{Z}/2^k\mathbb{Z}$, and each random tape needs $2t \cdot m$ elements of randomness in $\mathbb{Z}/2^k\mathbb{Z}$. Following the idea in [GPS22], we can use random packed Shamir sharings over \mathcal{R} to provide such randomness for the Π_i -sharings. Note that $2t \cdot m + \ell < 2n' \cdot m$, so for a packed Shamir secret sharing scheme whose secret space is in \mathcal{R}^r , we need to prepare $2n' \cdot m$ random sharings in this scheme to provide enough randomness for $\ell \cdot r$ random Π_i sharings.

In our protocol, we will use the following packed secret sharing scheme Σ over $\mathbb{Z}/2^k\mathbb{Z}$ [GPS22]:

- The secret space is $Z = (\mathbb{Z}/2^k\mathbb{Z})^{r \cdot \ell}$.
- The share space is $U = \mathcal{R}$, which can be viewed as $(\mathbb{Z}/2^k\mathbb{Z})^m$.
- For $\mathbf{x} \in Z$, the sharing of \mathbf{x} is computed as follows: First we divide \mathbf{x} into r vectors of dimension ℓ . Then we map each vector to a ring element of \mathcal{R} using $\phi(\cdot)$, i.e. we map \mathbf{x} into $\Phi(\mathbf{x})$. Next, we use a degree- d' packed Shamir secret sharing scheme over \mathcal{R} to store the r elements in \mathcal{R} .
- For a sharing $[\mathbf{x}] \in \Sigma$, we first view it as a degree- d' packed Shamir secret sharing over \mathcal{R} and reconstruct its secret $\mathbf{s} \in \mathcal{R}^r$. To recover the secret, we apply ψ , the inverse map of ϕ , to each element in \mathbf{s} .

Also, we define another secret sharing scheme Σ' over $\mathbb{Z}/2^k\mathbb{Z}$ as follows [GPS22]:

- The secret space is $Z' = (\mathbb{Z}/2^k\mathbb{Z})^{r \cdot \ell}$.
- The share space is $U' = \mathcal{R}$, which can be viewed as $(\mathbb{Z}/2^k\mathbb{Z})^m$.
- For $\mathbf{x} \in Z$, the sharing of \mathbf{x} is computed as follows: First we divide \mathbf{x} into r vectors of dimension ℓ . Then we map \mathbf{x} into $\Phi(\mathbf{x})$. We sample a vector of elements $\mathbf{o} = (o_1, \dots, o_r) \in \mathcal{R}^r$ where $\Psi(\mathbf{o}) = \mathbf{0} \in (\mathbb{Z}/2^k\mathbb{Z})^{r \cdot \ell}$. It has been shown in section 2.6 that an efficient sampling algorithm for $o_i \in \mathcal{R}$ such that $\psi(o_i) = \mathbf{0} \in (\mathbb{Z}/2^k\mathbb{Z})^\ell$ exists. Next, we use a degree- d' packed Shamir secret sharing scheme over \mathcal{R} to store the r elements of $\Phi(\mathbf{x}) + \mathbf{o}$ in \mathcal{R} .
- For a sharing $[\mathbf{x}] \in \Sigma'$, we first view it as a degree- $(d' + r - 1)$ packed Shamir secret sharing over \mathcal{R} and reconstruct its secret $\mathbf{s} \in \mathcal{R}^r$. To recover the secret, we apply $\Psi(\cdot)$ to \mathbf{s} and get $\Psi(\mathbf{s})$.

We pick $d' = \lceil \frac{n'-1}{2} \rceil$ and $r = \lfloor \frac{n'-2t'+1}{2} \rfloor$, to guarantee that the degree- d' packed Shamir sharings can hold r secrets in $(\mathbb{Z}/2^k\mathbb{Z})^\ell$ ($d' - r + 1 \geq t$) and that the degree- $(d' + r - 1)$ packed Shamir sharings can be reconstructed towards an honest party while detecting errors ($d' + r - 1 < n' - t'$ by Lemma 2). Therefore, Σ is a degree- $\lceil \frac{n'-1}{2} \rceil$ packed Shamir secret sharing scheme that holds

$r = \lfloor \frac{n'-2t'+1}{2} \rfloor$ secrets in $(\mathbb{Z}/2^k\mathbb{Z})^\ell$, and Σ' is a degree- $(n' - t' - 1)$ packed Shamir secret sharing scheme that holds $r = \lfloor \frac{n'-2t'+1}{2} \rfloor$ secrets in $(\mathbb{Z}/2^k\mathbb{Z})^\ell$.

In the following, we will use $[\Phi(\mathbf{x})]_{d'}$ to denote a Σ -sharing of the secret $\mathbf{x} \in (\mathbb{Z}/2^k\mathbb{Z})^{r \cdot \ell}$, and we will use $[\mathbf{y}]_{n'-t'-1}$ to denote a Σ' -sharing of $\mathbf{z} \in (\mathbb{Z}/2^k\mathbb{Z})^{r \cdot \ell}$ such that $\Psi(\mathbf{y}) = \mathbf{z}$. It has been shown in [GPS22] that Σ and Σ' satisfy that for all $\mathbf{c} \in (\mathbb{Z}/2^k\mathbb{Z})^{r \cdot \ell}$

$$[\Phi(\mathbf{c}) \star \Phi(\mathbf{x})]_{n'-t'-1} = \Phi(\mathbf{c}) \star [\Phi(\mathbf{x})]_{d'}.$$

To reconstruct the secret $\mathbf{c} \star \mathbf{x}$ from $[\Phi(\mathbf{c}) \star \Phi(\mathbf{x})]_{n'-t'-1}$, we apply $\Psi(\cdot)$ to $\Phi(\mathbf{c}) \star \Phi(\mathbf{x})$ and get $\Psi(\Phi(\mathbf{c}) \star \Phi(\mathbf{x})) = \mathbf{c} \star \mathbf{x}$.

Our construction will use the ideal functionality $\mathcal{F}_{\text{RandShPacked}}$ (Functionality 10) that prepare random Σ -sharings, and the ideal functionality $\mathcal{F}_{\text{RandZeroSh}}$ (Functionality 12) that prepares random Σ' -sharings of $\mathbf{0} \in \mathcal{R}^r$.

The goal of the protocol is to prepare $\ell \cdot r$ random sharings $[\mathbf{x}_1], \dots, [\mathbf{x}_{\ell \cdot r}]$ such that $[\mathbf{x}_i]$ is a random Π_i -sharing, i.e. realizing $\mathcal{F}_{\text{RandSharTrans}}(\Pi_1, \dots, \Pi_{\ell \cdot r})$. The protocol that implements $\mathcal{F}_{\text{RandSharTrans}}$ is shown in Protocol 12.

Protocol 12: $\Pi_{\text{RandSharTrans}}(\Pi_1, \dots, \Pi_{\ell \cdot r})$

1. All parties run the protocol Π_{Rand} to obtain $2n' \cdot m$ random Σ -sharings denoted by

$$[\Phi(\mathbf{u}_1)]_{d'}, \dots, [\Phi(\mathbf{u}_{2n' \cdot m})]_{d'},$$

or get a semi-corrupted pair denoted by E_1 . On receiving E_1 , all parties take E_1 as output and halt. Otherwise, suppose $\mathbf{u}_i = (\mathbf{u}_i^{(1)}, \dots, \mathbf{u}_i^{(r)})$ where $\mathbf{u}_i^{(j)} \in (\mathbb{Z}/2^k\mathbb{Z})^\ell$ and denote the ℓ -th element of $\mathbf{u}_i^{(j)}$ by $u_i^{(j,\ell)}$. For all $i \in [r]$ and for all $q \in [\ell]$, let

$$\boldsymbol{\tau}_{(q-1) \cdot r + i} = (u_1^{(i,q)}, u_2^{(i,q)}, \dots, u_{\ell}^{(i,q)}) \in (\mathbb{Z}/2^k\mathbb{Z})^{\ell \cdot \ell},$$

and

$$\boldsymbol{\rho}_{(q-1) \cdot r + i} = (u_{\ell+1}^{(i,q)}, u_{\ell+2}^{(i,q)}, \dots, u_{\ell+\tilde{q}_i}^{(i,q)}) \in (\mathbb{Z}/2^k\mathbb{Z})^{2t \cdot m \cdot \ell}.$$

The goal of this protocol is to compute the Π_i -sharing $[\mathbf{x}_i] = \text{share}_i(\boldsymbol{\tau}_i, \boldsymbol{\rho}_i)$.

2. All parties run the protocol Π_{Rand} and obtain $2n' \cdot m$ random Σ' -sharings of $\mathbf{0} \in \mathcal{R}^r$, denoted by

$$\{[\mathbf{o}_j^{(1)}]_{n'-t'-1}, [\mathbf{o}_j^{(2)}]_{n'-t'-1}, \dots, [\mathbf{o}_j^{(2m)}]_{n'-t'-1}\}_{j=1}^{n'},$$

or get a semi-corrupted pair denoted by E_2 . On receiving E_2 , all parties take E_2 as output and halt.

3. For all $i \in [r \cdot \ell]$, for all $j \in [n']$, and for all $p \in [2m]$, let $\mathcal{L}_j^{i,p} : Z_i \times (\mathbb{Z}/2^k\mathbb{Z})^{2t \cdot m} \rightarrow \mathbb{Z}/2^k\mathbb{Z}$ denote the $\mathbb{Z}/2^k\mathbb{Z}$ linear transformation such that for all $\boldsymbol{\tau} \in \tilde{Z}_i$ and $\boldsymbol{\rho} \in (\mathbb{Z}/2^k\mathbb{Z})^{2t \cdot m}$, $\mathcal{L}_j^{i,p}(\boldsymbol{\tau}, \boldsymbol{\rho})$ outputs the p -th element of the j -th share of the Π_i -sharing $\text{share}_i(\boldsymbol{\tau}, \boldsymbol{\rho})$. Then there exist $c_{j,1}^{(i,p)}, \dots, c_{j,\ell+2t \cdot m}^{(i,p)} \in \mathbb{Z}/2^k\mathbb{Z}$ such that

$$\mathcal{L}_j^{i,p}(\boldsymbol{\tau}, \boldsymbol{\rho}) = \sum_{w=1}^{\ell} c_{j,w}^{(i,p)} \cdot \tau_w + \sum_{w=1}^{2t \cdot m} c_{j,\ell+w}^{(i,p)} \cdot \rho_w.$$

For all $j \in [n']$, for all $p \in [2m]$, and for all $w \in [2n' \cdot m]$, let

$$\mathbf{c}_{j,w}^{(\star,p)} = (c_{j,w}^{(1,p)}, c_{j,w}^{(2,p)}, \dots, c_{j,w}^{(r \cdot \ell, p)}) \in (\mathbb{Z}/2^k\mathbb{Z})^{r \cdot \ell},$$

where $c_{j,w}^{(i,p)} = 0$ for all $w > \ell + 2t \cdot m$.

4. For all $i \in [r \cdot \ell]$, for all $j \in [n']$, and for all $p \in [2m]$, let $v_j^{(i,p)} = \mathcal{L}_j^{i,p}(\boldsymbol{\tau}_i, \boldsymbol{\rho}_i)$. Let

$$\mathbf{v}_j^{(\star,p)} = (v_j^{(1,p)}, v_j^{(2,p)}, \dots, v_j^{(r \cdot \ell, p)}) \in (\mathbb{Z}/2^k\mathbb{Z})^{r \cdot \ell}.$$

Then for all $j \in [n']$ and for all $p \in [2m]$, all parties locally compute the Σ' -sharing

$$[\mathbf{s}_j^{(\star,p)}]_{n'-t'-1} = [\mathbf{o}_j^{(p)}]_{n'-t'-1} + \sum_{w=1}^{2n' \cdot m} \Phi(\mathbf{c}_{j,w}^{(\star,p)}) \star [\Phi(\mathbf{u}_w)]_{d'}.$$

Note that $\mathbf{s}_j^{(\star,p)} \in \mathcal{R}^r$, and $\Psi(\mathbf{s}_j^{(\star,p)}) = \mathbf{v}_j^{(\star,p)}$. Then all parties send their shares of $[\mathbf{s}_j^{(\star,p)}]_{n'-t'-1}$ to P_j .

5. For all $j \in [n']$ and for all $p \in [2m]$, P_j reconstructs the Σ' sharing $[\mathbf{s}_j^{(*,p)}]_{n'-t'-1}$ and gets $\mathbf{v}_j^{(*,p)} = (\Psi(\mathbf{s}_j^{(1,p)}), \Psi(\mathbf{v}_j^{(2,p)}), \dots, \Psi(\mathbf{v}_j^{(r,p)}))$, or P_j detects error in the sharing $[\mathbf{s}_j^{(*,p)}]_{n'-t'-1}$ and gets **unhappy**.
6. **Fault Detection Phase:**
 - (a) All players broadcast their happy-bits using the byzantine agreement protocol.
 - (b) For each party, if it receives at least one **unhappy** then sets his or her happy-bit to **unhappy**.
 - (c) All players run a consensus protocol on their happy-bits. If the consensus is **happy**, all parties output the output of π and terminate the procedure. Otherwise, they proceed to the following steps.
7. **Fault Localization Phase:**
 - (a) All players agree the party with the smallest index in \mathcal{P}_A as the dealer D . All other players send all their generated values and communication to D .
 - (b) On receiving all the information, D simulates π and the fault detection phase itself. D either prepares the message $(P_i, \text{corrupt})$ if P_i failed to follow the protocol, or the message $(P_{j_1}, P_{j_2}, \text{idx}, \mathbf{b}, \mathbf{b}', \text{disputed})$ if the idx -th bit \mathbf{b} that P_{j_1} sends to P_{j_2} is inconsistent with P_{j_2} 's bit \mathbf{b}' that P_{j_2} claims to have received. Then D broadcasts the prepared message to all players.
 - (c) If $(P_i, \text{corrupt})$ is received, all players set the eliminating set $E = \{D, P_i\}$. Otherwise, if $(P_{j_1}, P_{j_2}, \text{idx}, \mathbf{b}, \mathbf{b}', \text{disputed})$ is received, P_{j_1} and P_{j_2} will broadcast if they agree with this message. If P_{j_1} does not agree, all players set the eliminating set $E = \{D, P_{j_1}\}$; otherwise if P_{j_2} does not agree, all players set the eliminating set $E = \{D, P_{j_2}\}$; otherwise all players set the eliminating set $E = \{P_{j_1}, P_{j_2}\}$.
 - (d) All players update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - E$ and halt.
8. For all $i \in [r \cdot \ell]$, P_j sets its share of the Π_i sharing $[\mathbf{x}_i]$ to be $\mathbf{v}_j^{(i)} = (v_j^{(i,1)}, v_j^{(i,2)}, \dots, v_j^{(i,2m)})$. All parties take $[\mathbf{x}_1], [\mathbf{x}_2], \dots, [\mathbf{x}_{r \cdot \ell}]$ as output.

Cost of Protocol 12. The communication complexity of running Π_{Rand} in Step 2 and Step 3 are both $O(n^2 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$. Step 5 has communication complexity of $O(n^2 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$. Therefore, the total communication complexity of Protocol 12 is $O(n^2 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

Lemma 11. *In Step 5 of Protocol 12, either at least one active honest party has **unhappy** happy-bit or the shares of the output sharings $\{[\mathbf{x}_i]\}_{i=1}^r$ held by active honest parties are valid and consistent.*

Proof. If no honest party get **unhappy**, by Lemma 2 all the Shamir sharings $[\mathbf{s}_j^{(*,p)}]_{n'-t'-1}$ are consistent for $j \in \mathcal{H}_{\text{active}}$. This indicates that no corrupted parties have deviated from the protocol in Step 4 by sending wrong sharings, so all the parties have acted faithfully. Therefore, the output sharings $\{[\mathbf{x}_i]\}_{i=1}^r$ held by active honest parties are valid and consistent.

Lemma 12. *The protocol $\Pi_{\text{RandSharTrans}}$ computes $\mathcal{F}_{\text{RandSharTrans}}$ with perfect security when $|\mathcal{C}_{\text{active}}| < |\mathcal{P}_{\text{active}}|/3$.*

Proof. Let \mathcal{A} denote the adversary. We will construct a simulator **Sim** to simulate the behaviors of honest parties for the adversary. Recall that $\mathcal{C}_{\text{active}}$ denotes the active corrupted parties and that $\mathcal{H}_{\text{active}}$ denotes the active honest parties.

Simulation for $\Pi_{\text{RandSharTrans}}$. **Sim** simulates the protocol $\Pi_{\text{RandSharTrans}}$ as follows:

1. In Step 1, **Sim** calls the simulator for Π_{Rand} to simulate the protocol Π_{Rand} . **Sim** receives either a semi-corrupted pair or the active corrupted parties' shares of the sharings $[\Phi(\mathbf{u}_1)]_{d'}, \dots, [\Phi(\mathbf{u}_{2n'-m})]_{d'}$. On receiving a semi-corrupted pair, **Sim** sends it to $\mathcal{F}_{\text{RandSharTrans}}$ and halts.
2. In Step 2, **Sim** calls the simulator for Π_{Rand} to simulate the protocol Π_{Rand} . **Sim** receives either a semi-corrupted pair or the active corrupted parties' share of the sharings

$$\{[\mathbf{o}_j^{(1)}]_{n'-t'-1}, [\mathbf{o}_j^{(2)}]_{n'-t'-1}, \dots, [\mathbf{o}_j^{(2m)}]_{n'-t'-1}\}_{j=1}^{n'}.$$

On receiving a semi-corrupted pair, **Sim** sends it to $\mathcal{F}_{\text{RandSharTrans}}$ and halts.

3. In step 5, **Sim** simulates the happy-bits of active honest parties as follows. For each honest party P_i , and for each corrupted party P_j , **Sim** compares the real shares held by P_j and the shares that P_j sent to P_i . If they are different, **Sim** sets P_i 's happy-bit to **unhappy** on behalf of P_i .
4. In Fault Detection Phase, **Sim** simulates the byzantine agreement protocol faithfully using active honest parties' happy-bits.
5. If the simulation proceeds to Fault Localization Phase, **Sim** samples random Σ sharings that are compatible with active corrupted parties' shares on behalf of active honest parties as inputs of Π_{Rand} in Step 1. Also, **Sim** samples random Σ' sharings that are compatible with active corrupted parties' shares on behalf of active honest parties as inputs of Π_{Rand} in Step 2. Then **Sim** faithfully follows the protocol.
6. If the simulation does not proceed to Fault Localization Phase, in Step 4, for each honest party P_j , **Sim** receives from \mathcal{A} the active corrupted parties' authentic shares of the sharing $[\mathbf{s}_j^{(*,p)}]_{n'-t'-1}$ for all $p \in [2m]$. Then **Sim** gets from \mathcal{A} the shares sent to P_j of the sharing $[\mathbf{s}_j^{(*,p)}]_{n'-t'-1}$ for all $p \in [2m]$. **Sim** simulates the active honest parties' shares of $([\mathbf{s}_j^{(*,p)}]_{n'-t'-1})_{j \in \mathcal{C}_{\text{active}}}$ for all $p \in [2m]$ as follows:
 - (a) For all $i \in [r \cdot \ell]$, **Sim** samples a random Π_i -sharing $[\tilde{\mathbf{x}}_i]$ that is compatible with active corrupted parties' shares sent to $\mathcal{F}_{\text{RandSharTrans}}$ from \mathcal{A} . Because Π_i has threshold t , **Sim** can always sample such sharings. Suppose party P_j 's share is denoted by $\mathbf{v}_j^{(i)} = (v_j^{(i,1)}, v_j^{(i,2)}, \dots, v_j^{(i,2m)})$.
 - (b) For all $j \in [n']$ and $q \in [2m]$, **Sim** computes $\mathbf{v}_j^{(*,q)}$.
 - (c) Because Σ' has threshold $n' - t' - r \geq t'$, the shares of $[\mathbf{s}_j^{(*,p)}]_{n'-t'-1}$ held by active honest parties are independent of the secret $\mathbf{v}_j^{(*,p)}$ and the shares held by active corrupted parties. For all $j \in \mathcal{C}_{\text{active}}$ and for all $q \in [2m]$, **Sim** samples the random Σ' -sharings $[\mathbf{s}_j^{(*,p)}]_{n'-t'-1}$ based on the secrets $\mathbf{v}_j^{(*,p)}$ and the shares of active corrupted parties simulated by **Sim**.
 - (d) **Sim** sends the shares of $[\mathbf{s}_j^{(*,p)}]_{n'-t'-1}$ held by active honest parties to \mathcal{A} .

Hybrid Arguments. We show that **Sim** perfectly simulates the behaviors of honest parties.

Hybrid₀: The execution in the real world.

Hybrid₁: In this hybrid, **Sim** calls the simulator for Π_{Rand} to simulate the protocol Π_{Rand} in Step 1 of $\Pi_{\text{RandSharTrans}}$, and calls the simulator for Π_{Rand} to simulate the protocol Π_{Rand} in Step 3 of $\Pi_{\text{RandSharTrans}}$. Then the distribution of this hybrid is identical to the distribution of **Hybrid₀** by the correctness of these two simulators.

Hybrid₂: In this hybrid, **Sim** simulates the happy-bits of active honest parties and uses them to simulate the byzantine agreement protocol in the Fault Detection Phase. Because the happy-bits of active honest parties only depends on the communication of \mathcal{A} , **Sim** can perfectly simulate this process.

Therefore, the distribution of this hybrid is identical to the distribution of **Hybrid₁**.

Hybrid₃: In this hybrid, **Sim** simulates the Fault Localization Phase by sampling random sharings that are compatible with active corrupted parties' shares for active honest parties in Step 1 and Step 2. Note that the Σ sharings and Σ' sharings sampled by **Sim** have the same distribution as the actual sharings sampled by active honest parties. Also note that it will not affect the output distribution, since the semi-corrupted pair is only related to the corrupted parties' messages. Therefore, the distribution of this hybrid is identical to the distribution of **Hybrid₂**.

Hybrid₄: In this hybrid, **Sim** simulates the shares of $([\mathbf{s}_j^{(*,p)}]_{n'-t'-1})_{j \in \mathcal{C}_{\text{active}}}$ held by active honest parties for all $p \in [2m]$ instead of getting them from active honest parties.

- We first show that the correct shares of the sharings $([\mathbf{s}_i^{(*,p)}]_{n'-t'-1})_{p=1}^{2m}$ held by active corrupted parties (i.e. $\{\tilde{\mathbf{v}}_j^{(i)}\}_{j \in \mathcal{C}_{\text{active}}}$) are the same as the distribution of that in **Hybrid₃**. To see this, in **Hybrid₃** the shares held by active corrupted parties are computed by following $\text{share}_i(\boldsymbol{\tau}_i, \boldsymbol{\rho}_i)$ for all $i \in [r]$. Because Σ is a degree- d' packed Shamir secret sharing and $d' - r + 1 \geq t'$, $\boldsymbol{\tau}_i$ and $\boldsymbol{\rho}_i$ are uniformly random. While in this hybrid, **Sim** randomly samples the shares of active corrupted parties by first sampling random sharings of Π_1, \dots, Π_r and then obtaining

the shares of active corrupted parties. Therefore, in this hybrid $\{\tilde{\mathbf{v}}_j^{(i)}\}_{j \in \mathcal{C}_{\text{active}}}$ has the same distribution as that in **Hybrid₃**.

- Next we argue that the active honest parties' shares of $[\mathbf{s}_j^{(*,p)}]_{n'-t'-1}$ has the same distribution as that in **Hybrid₃**. Recall that $[\mathbf{o}_j^{(p)}]_{n'-t'-1}$ are random Σ' sharings of $\mathbf{0}$. Because Σ' has a threshold no less than t' , in **Hybrid₃** $[\mathbf{s}_j^{(*,p)}]_{n'-t'-1}$ is a random Σ' sharing of the secret $\mathbf{v}_j^{(*,p)}$ given the shares of the corrupted parties and the secret. Since $\mathbf{v}_j^{(*,p)}$ is determined by $\{\boldsymbol{\tau}_i, \boldsymbol{\rho}_i\}_{i=1}^r$, which are independent of the shares of $\{[\Phi(\mathbf{u}_w)]_{d'}\}_{w=1}^{2n' \cdot m}$ and $[\mathbf{o}_j^{(p)}]_{n'-t'-1}$ held by active corrupted parties, $\mathbf{v}_j^{(*,p)}$ is independent of the shares of $[\mathbf{s}_j^{(*,p)}]_{n'-t'-1}$ held by active corrupted parties.

In this hybrid, because **Sim** computes the shares of $[\mathbf{s}_j^{(*,p)}]_{n'-t'-1}$ held by active corrupted parties given the shares of $[\mathbf{o}_j^{(p)}]_{n'-t'-1}$ held by active corrupted parties, and because the secret $\mathbf{v}_j^{(*,p)}$, a part of $\{\mathbf{v}_j^{(i)}\}_{j \in \mathcal{C}_{\text{active}}}$, has the same distribution as the one in **Hybrid₃**, we have that the secret $\mathbf{v}_j^{(*,p)}$ and the shares of $[\mathbf{s}_j^{(*,p)}]_{n'-t'-1}$ held by active corrupted parties has the same distribution as in **Hybrid₃**. Since the shares of $[\mathbf{s}_j^{(*,p)}]_{n'-t'-1}$ held by active honest parties in this hybrid are sampled according to the secret $\mathbf{v}_j^{(*,p)}$ and the shares of $[\mathbf{s}_j^{(*,p)}]_{n'-t'-1}$ held by active corrupted parties computed by **Sim**, they have the same distribution as that in **Hybrid₃**.

- Finally we show that the distribution of active honest parties' output in $\Pi_{\text{RandSharTrans}}$ given the view of corrupted parties is the same as that in **Hybrid₃**. We have (1) for all $w \in [2n' \cdot m]$, the shares of $[\Phi(\mathbf{u}_w)]_{d'}$ held by active corrupted parties are independent of $\{\boldsymbol{\tau}_i, \boldsymbol{\rho}_i\}_{i=1}^r$; (2) for all $j \in [n']$ and for all $p \in [2m]$, the shares of $[\mathbf{o}_j^{(p)}]_{n'-t'-1}$ held by active corrupted parties are independent of $\{\boldsymbol{\tau}_i, \boldsymbol{\rho}_i\}_{i=1}^r$; (3) for all $j \in \mathcal{C}_{\text{active}}$, the shares of $[\mathbf{s}_j^{(*,p)}]_{n'-t'-1}$ held by active honest parties are only dependent on $\{\mathbf{v}_j^{(i)}\}_{j \in \mathcal{C}_{\text{active}}}$. Recall that $[\mathbf{x}_i] = \text{share}_i(\boldsymbol{\tau}_i, \boldsymbol{\rho}_i)$. Therefore, in **Hybrid₃** the view of corrupted parties is independent of the joint of shares of $[\mathbf{x}_i] = \text{share}_i(\boldsymbol{\tau}_i, \boldsymbol{\rho}_i)$ held by active honest parties and the secret \mathbf{x}_i .

In this hybrid, **Sim** perfectly simulates the shares of $[\mathbf{x}_i] = \text{share}_i(\boldsymbol{\tau}_i, \boldsymbol{\rho}_i)$ held by active corrupted parties and send them to $\mathcal{F}_{\text{RandSharTrans}}$, and then $\mathcal{F}_{\text{RandSharTrans}}$ samples a random Π_i -sharing based on the shares of active corrupted parties. The output of active honest parties is their shares of $[\mathbf{x}_i]$ for all $i \in [r \cdot \ell]$.

Therefore, the output of active honest parties given the joint view of corrupted parties in this hybrid has the same distribution as that in **Hybrid₃**.

Therefore, the distribution of this hybrid is identical to the distribution of **Hybrid₃**.

Note that this hybrid is the execution in the ideal world.

E.3 Performing Linear Transformations

Let L_1, \dots, L_N be N linear mappings in the linear mapping set $\mathcal{L}((\mathbb{Z}/2^k\mathbb{Z})^\ell, (\mathbb{Z}/2^k\mathbb{Z})^\ell)$, and let $[\phi(\mathbf{x}_1)]_t, \dots, [\phi(\mathbf{x}_N)]_t$ be N Shamir sharings. We introduce the functionality that applies the linear transformation L_i on the Shamir sharing $[\phi(\mathbf{x}_i)]_t$ and gets $[\phi(L_i(\mathbf{x}_i))]_t$ for all $i \in [N]$, or outputs a semi-corrupted pair $\{P_i, P_j\}$ such that $\{P_i, P_j\} \cap \mathcal{C}_{\text{active}} \neq \emptyset$. The functionality $\mathcal{F}_{\text{LinearTrans}}$ is described in Functionality 17. We assume that for each degree- t Shamir secret sharing, the shares of all active honest parties lie on a degree- t polynomial.

Functionality 17: $\mathcal{F}_{\text{LinearTrans}}(N)$

1. Let $[\phi(\mathbf{x}_1)]_t, \dots, [\phi(\mathbf{x}_N)]_t$ denote the input sharings. Let L_1, \dots, L_N be N linear mappings in the linear mapping set $\mathcal{L}((\mathbb{Z}/2^k\mathbb{Z})^\ell, (\mathbb{Z}/2^k\mathbb{Z})^\ell)$.
2. $\mathcal{F}_{\text{LinearTrans}}$ receives from active honest parties their shares of $[\phi(\mathbf{x}_1)]_t, \dots, [\phi(\mathbf{x}_N)]_t$, and reconstructs $\mathbf{x}_1, \dots, \mathbf{x}_N$. $\mathcal{F}_{\text{LinearTrans}}$ further computes the shares of $[\phi(\mathbf{x}_1)]_t, \dots, [\phi(\mathbf{x}_N)]_t$ held by active corrupted parties and send them to the adversary.
3. $\mathcal{F}_{\text{LinearTrans}}$ receives from the adversary a set of even number of parties $S \subset \mathcal{P}_{\text{active}}$ such that $|S|/2 \leq |S \cap \mathcal{C}_{\text{active}}|$. $\mathcal{F}_{\text{LinearTrans}}$ updates $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S$. $\mathcal{F}_{\text{LinearTrans}}$ then sends S to all active honest parties.

4. $\mathcal{F}_{\text{LinearTrans}}$ computes $L_1(\mathbf{x}_1), \dots, L_N(\mathbf{x}_N)$, and then $\mathcal{F}_{\text{LinearTrans}}$ receives from the adversary the shares of active corrupted parties of $[\phi(L_1(\mathbf{x}_1))]_t, \dots, [\phi(L_N(\mathbf{x}_N))]_t$;
5. $\mathcal{F}_{\text{LinearTrans}}$ samples the whole sharing $[\phi(L_1(\mathbf{x}_1))]_t, \dots, [\phi(L_N(\mathbf{x}_N))]_t$ so that they are compatible with the active corrupted parties' shares. For each active honest party P_h , $\mathcal{F}_{\text{LinearTrans}}$ sends P_h 's shares of $[\phi(L_1(\mathbf{x}_1))]_t, \dots, [\phi(L_N(\mathbf{x}_N))]_t$ to P_h .

Protocol Description. Recall that in Section E.2, for the linear mapping $L_i : (\mathbb{Z}/2^k\mathbb{Z})^\ell \rightarrow (\mathbb{Z}/2^k\mathbb{Z})^\ell$ we construct a corresponding $\mathbb{Z}/2^k\mathbb{Z}$ -arithmetic secret sharing scheme Π_i as follows:

- The secret space $Z_i = (\mathbb{Z}/2^k\mathbb{Z})^\ell$.
- The share space $U_i = (\mathcal{R})^2$.
- For a secret $\mathbf{x} \in Z_i$, the sharing of \mathbf{x} is $([\phi(\mathbf{x})]_t, [\phi(L_i(\mathbf{x}))]_t)$.
- For a sharing $([\phi(\mathbf{x})]_t, [\phi(L_i(\mathbf{x}))]_t)$, we reconstruct the secret \mathbf{x} using $[\phi(\mathbf{x})]_t$.

Therefore, we can use $\mathcal{F}_{\text{RandSharTrans}}$ to prepare random Π_i sharings, and perform linear transformation to a sharing $[\phi(\mathbf{x})]_t$ following the idea in [GPS22]. Our protocol that implements $\mathcal{F}_{\text{LinearTrans}}$ is described in Protocol 13.

Protocol 13: $\Pi_{\text{LinearTrans}}(N, \{[\phi(\mathbf{x}_i)]_t\}_{i=1}^N, \{L_i\}_{i=1}^N)$

1. Recall that $[\phi(\mathbf{x}_1)]_t, \dots, [\phi(\mathbf{x}_N)]_t$ are the input sharings, and that L_1, \dots, L_N are N linear mappings in the linear mapping set $\mathcal{L}((\mathbb{Z}/2^k\mathbb{Z})^\ell, (\mathbb{Z}/2^k\mathbb{Z})^\ell)$. Let Π_i denote the $\mathbb{Z}/2^k\mathbb{Z}$ -arithmetic secret sharing scheme corresponding to L_i as defined previously. All parties set the eliminated party set $S = \emptyset$.
2. All parties invoke $\mathcal{F}_{\text{RandSharTrans}}$ $\lceil \frac{N}{r \cdot \ell} \rceil$ times. All parties receive the random Π_1, \dots, Π_N sharings $([\phi(\mathbf{r}_1)]_t, [\phi(L_1(\mathbf{r}_1))]_t), \dots, ([\phi(\mathbf{r}_N)]_t, [\phi(L_N(\mathbf{r}_N))]_t)$. If any call of $\mathcal{F}_{\text{RandSharTrans}}$ outputs a semi-corrupted pair E , all parties eliminate the semi-corrupted pair, set $S := S \cup E$ and call $\mathcal{F}_{\text{RandSharTrans}}$ again.
3. All parties invoke $\mathcal{F}_{\text{OpenPub}}(N)$ on the sharings $[\phi(\mathbf{x}_1)]_t + [\phi(\mathbf{r}_1)]_t, \dots, [\phi(\mathbf{x}_N)]_t + [\phi(\mathbf{r}_N)]_t$ and get the output $\phi(\mathbf{x}_1) + \phi(\mathbf{r}_1), \dots, \phi(\mathbf{x}_N) + \phi(\mathbf{r}_N)$. Then all parties locally compute $\mathbf{x}_1 + \mathbf{r}_1, \dots, \mathbf{x}_N + \mathbf{r}_N$ by applying $\psi(\cdot)$ to $\phi(\mathbf{x}_1) + \phi(\mathbf{r}_1), \dots, \phi(\mathbf{x}_N) + \phi(\mathbf{r}_N)$.
4. All parties output the eliminated party set S . For all $i \in [N]$, all parties locally compute $\phi(L_i(\mathbf{x}_i + \mathbf{r}_i))$, and then take $L_i(\mathbf{x}_i + \mathbf{r}_i) - [\phi(L_i(\mathbf{r}_i))]_t$ as output.

Cost of Protocol 13. In Step 2, all parties invoke $\mathcal{F}_{\text{RandSharTrans}}$ $O(\frac{N}{r \cdot \ell} + t)$ times, so the cost of Step 2 is $O(\frac{N}{r \cdot \ell} \cdot n^2 \cdot m^2 + n^2 \cdot m^2 \cdot t)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$. In Step 3, the communication complexity of calling $\mathcal{F}_{\text{OpenPub}}(N)$ is $O(n \cdot m \cdot N + n^2 \cdot m)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$. Note that $m/\ell = O(1)$, $r = O(n)$ and $t = O(n)$, so the total communication complexity of Protocol 13 is $O(N \cdot n \cdot m + n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

Lemma 13. *The protocol $\Pi_{\text{LinearTrans}}$ computes $\mathcal{F}_{\text{linear-transformation}}$ with perfect security in the $(\mathcal{F}_{\text{RandSharTrans}}, \mathcal{F}_{\text{OpenPub}})$ -hybrid model when $|\mathcal{C}_{\text{active}}| < |\mathcal{P}_{\text{active}}|/3$.*

Proof. Let \mathcal{A} denote the adversary. We will construct a simulator Sim to simulate the behaviors of honest parties for the adversary. In Step 2, Sim faithfully emulates $\mathcal{F}_{\text{RandSharTrans}}$. In Step 3, Sim faithfully emulates $\mathcal{F}_{\text{OpenPub}}$. Because there is no other communication, the distribution of the real world is the same as the distribution of the ideal world.

E.4 Fan-Out, Permute, and Collect

Instantiating $\mathcal{F}_{\text{FanOut}}$. The functionality $\mathcal{F}_{\text{FanOut}}$ is described in Functionality 18.

Functionality 18: $\mathcal{F}_{\text{FanOut}}(N)$

1. Let $[\phi(\mathbf{z}^{(1)})]_t, \dots, [\phi(\mathbf{z}^{(N)})]_t$ denote the input sharings. Let $n_j^{(i)}$ denote the number of times the j -th wire of $\mathbf{z}^{(i)}$ needs to be copied and let $m^{(i)} = \frac{\sum_{j=1}^{\ell} n_j^{(i)}}{\ell}$. For all $i \in [N]$, let $f_j^{(i)}$ denote the desired fan-out linear transformations for $j = 1, \dots, m^{(i)}$.
2. $\mathcal{F}_{\text{FanOut}}$ receives from active honest parties their shares of $[\phi(\mathbf{z}^{(1)})]_t, \dots, [\phi(\mathbf{z}^{(N)})]_t$. Then $\mathcal{F}_{\text{FanOut}}$ reconstructs the secrets $\phi(\mathbf{z}^{(1)}), \dots, \phi(\mathbf{z}^{(N)})$. $\mathcal{F}_{\text{FanOut}}$ further computes the shares of $[\phi(\mathbf{z}^{(1)})]_t, \dots, [\phi(\mathbf{z}^{(N)})]_t$ held by active corrupted parties and send the shares to the adversary.
3. $\mathcal{F}_{\text{FanOut}}$ locally computes $\mathbf{z}^{(i)} = \psi(\phi(\mathbf{z}^{(i)}))$ for all $i \in [N]$. Then for all $i \in [N]$, $\mathcal{F}_{\text{FanOut}}$ computes $\mathbf{x}_j^{(i)} := f_j^{(i)} \cdot \mathbf{z}^{(i)}$ for all $j = 1, \dots, m^{(i)}$.
4. $\mathcal{F}_{\text{FanOut}}$ receives from the adversary a set of even number of parties $S \subset \mathcal{P}_{\text{active}}$ such that $|S|/2 \leq |S \cap \mathcal{C}_{\text{active}}|$. $\mathcal{F}_{\text{FanOut}}$ updates $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S$. $\mathcal{F}_{\text{FanOut}}$ then sends S to all active honest parties.
5. $\mathcal{F}_{\text{FanOut}}$ receives from the adversary the shares of active corrupted parties of $[\phi(\mathbf{x}_j^{(i)})]_t$ for all $i \in [N]$ and for all $j \in [m^{(i)}]$.
6. $\mathcal{F}_{\text{FanOut}}$ samples the whole sharings $[\phi(\mathbf{x}_j^{(i)})]_t$ so that they are compatible with the active corrupted parties' shares for all $i \in [N]$ and for all $j \in [m^{(i)}]$. For each active honest party P_h , $\mathcal{F}_{\text{FanOut}}$ sends P_h 's shares of $[\phi(\mathbf{x}_j^{(i)})]_t$ for all $i \in [N]$ and for all $j \in [m^{(i)}]$ to P_h .

Note that each fan-out operation can be viewed as a linear transformation of the original secret sharing, so we will implement $\mathcal{F}_{\text{FanOut}}$ using the functionality $\mathcal{F}_{\text{LinearTrans}}$. Let $\{[\phi(\mathbf{z}^{(i)})]_t\}_{i=1}^N$ denote the input Shamir sharings of the protocol, and let $n_j^{(i)}$ denote the number of times the j -th wire needs to be copied in the fan-out sharings. We require that $\sum_{j=1}^{\ell} n_j^{(i)}$ is a multiple of ℓ greater than 0 for all $i \in [N]$. Recall that $m^{(i)} := \frac{\sum_{j=1}^{\ell} n_j^{(i)}}{\ell}$ denote the number of fan-out operations needed for the output sharing $[\phi(\mathbf{z}^{(i)})]_t$. All parties first need to agree on the sequence of all entries in the fan-out results. For all $i \in [N]$, let $f_j^{(i)}$ denote the agreed fan-out operations for $j = 1, \dots, m^{(i)}$. Note that $f_j^{(i)}$ is a linear transformation. The description of Π_{FanOut} that implements $\mathcal{F}_{\text{FanOut}}$ appears in Protocol 14.

Protocol 14: $\Pi_{\text{FanOut}}(N, \{([\phi(\mathbf{z}^{(i)})]_t, (n_j^{(i)})_{j=1}^{\ell})_{i=1}^N\})$

1. All parties locally compute the desired fan-out linear transformations $f_j^{(i)}$ on $\mathbf{z}^{(i)}$ for $j = 1, \dots, m^{(i)}$ for all $i \in [N]$.
2. All parties invoke $\mathcal{F}_{\text{LinearTrans}}$ with $[\phi(\mathbf{z}^{(1)})]_t, \dots, [\phi(\mathbf{z}^{(N)})]_t$ and the fan-out operations $\{f_j^{(i)} : j \in [m^{(i)}], i \in [N]\}$. Then all parties get the set of eliminated parties S from $\mathcal{F}_{\text{LinearTrans}}$ and the output sharings

$$\{[\phi(f_j^{(i)} \cdot \mathbf{z}^{(i)})]_t : j \in [m^{(i)}], i \in [N]\}.$$

All parties output S and all sharings in $\{[\phi(f_j^{(i)} \cdot \mathbf{z}^{(i)})]_t : j \in [m^{(i)}], i \in [N]\}$.

Let $M := \sum_{i=1}^N m^{(i)}$ denote the total number of fan-out sharings which are the output of the protocol. In Step 2, the communication complexity of calling $\mathcal{F}_{\text{LinearTrans}}(M)$ is $O(M \cdot n \cdot m + n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$. Therefore, the total communication complexity of Π_{FanOut} is $O(M \cdot n \cdot m + n^3 \cdot m^2)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

Lemma 14. *The protocol Π_{FanOut} computes $\mathcal{F}_{\text{FanOut}}$ with perfect security in the $\mathcal{F}_{\text{LinearTrans}}$ -hybrid model when $|\mathcal{C}_{\text{active}}| < |\mathcal{P}_{\text{active}}|/3$.*

Proof. Let \mathcal{A} denote the adversary. We will construct a simulator Sim to simulate the behaviors of honest parties for the adversary. In Step 2, Sim faithfully emulates $\mathcal{F}_{\text{LinearTrans}}$. Because there is no other communication, the distribution of the real world is the same as the distribution of the ideal world.

Instantiating $\mathcal{F}_{\text{Permute}}$. Let $[\phi(\mathbf{x}_1)]_t, \dots, [\phi(\mathbf{x}_N)]_t$ denote the input sharings, and let p_1, \dots, p_N denote the desired permutations. The description of $\mathcal{F}_{\text{Permute}}$ appears in Functionality 19.

Functionality 19: $\mathcal{F}_{\text{Permute}}(N)$

1. Let $[\phi(\mathbf{x}_1)]_t, \dots, [\phi(\mathbf{x}_N)]_t$ denote the input sharings. Let p_1, \dots, p_N denote the desired permutations. $\mathcal{F}_{\text{Permute}}$ receives from active honest parties their shares of $[\phi(\mathbf{x}_1)]_t, \dots, [\phi(\mathbf{x}_N)]_t$. Then $\mathcal{F}_{\text{Permute}}$ reconstructs the secrets $\phi(\mathbf{x}_1), \dots, \phi(\mathbf{x}_N)$. $\mathcal{F}_{\text{Permute}}$ further computes the shares of $([\phi(\mathbf{x}_1)]_t, \dots, [\phi(\mathbf{x}_N)]_t)$ held by active corrupted parties and send the shares to the adversary.
2. $\mathcal{F}_{\text{Permute}}$ locally computes $\mathbf{x}_i = \psi(\phi(\mathbf{x}_i))$ for all $i \in [N]$. Then $\mathcal{F}_{\text{Permute}}$ computes $\mathbf{y}_i := p_i \cdot \mathbf{x}_i$ for all $i \in [N]$.
3. $\mathcal{F}_{\text{Permute}}$ receives from the adversary a set of even number of parties $S \subset \mathcal{P}_{\text{active}}$ such that $|S|/2 \leq |S \cap \mathcal{C}_{\text{active}}|$. $\mathcal{F}_{\text{Permute}}$ updates $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S$. $\mathcal{F}_{\text{Permute}}$ then sends S to all active honest parties.
4. $\mathcal{F}_{\text{Permute}}$ receives from the adversary the shares of active corrupted parties of $[\phi(\mathbf{y}_1)]_t, \dots, [\phi(\mathbf{y}_N)]_t$.
5. $\mathcal{F}_{\text{Permute}}$ samples the whole sharings $[\phi(\mathbf{y}_1)]_t, \dots, [\phi(\mathbf{y}_N)]_t$ so that they are compatible with the active corrupted parties' shares. For each active honest party P_h , $\mathcal{F}_{\text{Permute}}$ sends P_h 's shares of $[\phi(\mathbf{y}_1)]_t, \dots, [\phi(\mathbf{y}_N)]_t$ to P_h .

The description of Π_{Permute} that implements $\mathcal{F}_{\text{Permute}}$ appears in Protocol 15. The communication complexity of Π_{Permute} with N input sharings is $O(N \cdot n \cdot m + n^3 \cdot m)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

Protocol 15: $\Pi_{\text{Permute}}(N, \{[\phi(\mathbf{x}_i)]_t\}_{i=1}^N, \{p_i\}_{i=1}^N)$

1. All parties invoke $\mathcal{F}_{\text{LinearTrans}}(N)$ with the input $[\phi(\mathbf{x}_1)]_t, \dots, [\phi(\mathbf{x}_N)]_t$ and the desired permutations p_1, \dots, p_N , and get a set of eliminated parties S and the sharings $[\phi(p_1 \cdot \mathbf{x}_1)]_t, \dots, [\phi(p_N \cdot \mathbf{x}_N)]_t$ as output.
2. All parties output S and the sharings $[\phi(p_1 \cdot \mathbf{x}_1)]_t, \dots, [\phi(p_N \cdot \mathbf{x}_N)]_t$.

Lemma 15. *The protocol Π_{Permute} computes $\mathcal{F}_{\text{Permute}}$ with perfect security in the $\mathcal{F}_{\text{LinearTrans}}$ -hybrid model when $|\mathcal{C}_{\text{active}}| < |\mathcal{P}_{\text{active}}|/3$.*

Proof. Let \mathcal{A} denote the adversary. We will construct a simulator Sim to simulate the behaviors of honest parties for the adversary. In Step 1, Sim faithfully emulates $\mathcal{F}_{\text{LinearTrans}}$. Because there is no other communication, the distribution of the real world is the same as the distribution of the ideal world.

Instantiating $\mathcal{F}_{\text{Collect}}$. The description of $\mathcal{F}_{\text{Collect}}$ appears in Functionality 20.

Functionality 20: $\mathcal{F}_{\text{Collect}}(N)$

1. Let $[\phi(\mathbf{x}_1)]_t, \dots, [\phi(\mathbf{x}_M)]_t$ denote the input Shamir sharings, and let $[\phi(\mathbf{z}_1)]_t, \dots, [\phi(\mathbf{z}_N)]_t$ denote the target output Shamir sharings. Suppose that the j -th wire of \mathbf{z}_i is collected from the $u_{i,j}$ -th wire of $\mathbf{x}_{w_{i,j}}$.
2. $\mathcal{F}_{\text{Collect}}$ receives from active honest parties their shares of $[\phi(\mathbf{x}_1)]_t, \dots, [\phi(\mathbf{x}_M)]_t$. $\mathcal{F}_{\text{Collect}}$ reconstructs the secrets $\phi(\mathbf{x}_1), \dots, \phi(\mathbf{x}_M)$. $\mathcal{F}_{\text{Collect}}$ further computes the shares of $([\phi(\mathbf{x}_1)]_t, \dots, [\phi(\mathbf{x}_M)]_t)$ held by active corrupted parties and send the shares to the adversary.
3. $\mathcal{F}_{\text{Collect}}$ locally computes $\mathbf{x}_i = \psi(\phi(\mathbf{x}_i))$ for all $i \in [M]$. Then $\mathcal{F}_{\text{Collect}}$ computes

$$\mathbf{z}_i := \sum_{j=1}^{\ell} e_{u_{i,j}} \star \mathbf{x}_{w_{i,j}}$$

- for all $i \in [N]$.
4. $\mathcal{F}_{\text{Collect}}$ receives from the adversary a set of even number of parties $S \subset \mathcal{P}_{\text{active}}$ such that $|S|/2 \leq |S \cap \mathcal{C}_{\text{active}}|$. $\mathcal{F}_{\text{Collect}}$ updates $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S$. $\mathcal{F}_{\text{Collect}}$ then sends S to all active honest parties.
5. $\mathcal{F}_{\text{Collect}}$ receives from the adversary the shares of active corrupted parties of $[\phi(\mathbf{z}_1)]_t, \dots, [\phi(\mathbf{z}_N)]_t$.
6. $\mathcal{F}_{\text{Collect}}$ samples the sharings $[\phi(\mathbf{z}_1)]_t, \dots, [\phi(\mathbf{z}_N)]_t$ so that they are compatible with the active corrupted parties' shares. For each active honest party P_h , $\mathcal{F}_{\text{Collect}}$ sends P_h 's shares of $[\phi(\mathbf{z}_1)]_t, \dots, [\phi(\mathbf{z}_N)]_t$ to P_h .

The protocol Π_{Collect} implements $\mathcal{F}_{\text{Collect}}$. The description of Π_{Collect} appears in Protocol 16. The communication complexity of Π_{Collect} to generate N output sharings is $O(N \cdot n \cdot m + n^3 \cdot m)$ elements in $\mathbb{Z}/2^k\mathbb{Z}$.

Protocol 16: $\Pi_{\text{Collect}}(N, \{[\phi(\mathbf{x}_i)]_t\}_{i=1}^M, \{(u_{i,j}, w_{i,j})\}_{i \in [N], j \in [\ell]})$

1. For all $i \in [N]$, all parties locally compute

$$[s_i]_t := \sum_{j=1}^{\ell} \phi(\mathbf{e}_{u_{i,j}}) \cdot [\phi(\mathbf{x}_{w_{i,j}})]_t.$$

We note that $\psi(s_i) = \mathbf{x}_i$ for all $i \in [N]$.

2. All parties call the functionality $\mathcal{F}_{\text{ReEncode}}$ with the input $[s_1]_t, \dots, [s_N]_t$, and get a set of eliminated parties S and the sharings $[\phi(\mathbf{z}_1)]_t, \dots, [\phi(\mathbf{z}_N)]_t$.
3. All parties locally update $\mathcal{P}_{\text{active}} := \mathcal{P}_{\text{active}} - S$, and then output the eliminated party set S and the sharings $[\phi(\mathbf{z}_1)]_t, \dots, [\phi(\mathbf{z}_N)]_t$.

Lemma 16. *The protocol Π_{Collect} computes $\mathcal{F}_{\text{Collect}}$ with perfect security in the $\mathcal{F}_{\text{ReEncode}}$ -hybrid model when $|\mathcal{C}_{\text{active}}| < |\mathcal{P}_{\text{active}}|/3$.*

Proof. Let \mathcal{A} denote the adversary. We will construct a simulator Sim to simulate the behaviors of honest parties for the adversary. In Step 2, Sim faithfully emulates $\mathcal{F}_{\text{ReEncode}}$. Because there is no other communication, the distribution of the real world is the same as the distribution of the ideal world.

F Security Proofs

F.1 Proof of Lemma 3

Proof. Suppose by contradiction that in some scenario the semi-corrupted pair E selected by the protocol consists of two honest parties. We discuss two cases based on whether the dealer is corrupted or not.

When D is honest, we have that the sharing $\langle \hat{o} \rangle$ must be an invalid additive sharing of 0. To see this, if $\langle o'' \rangle$ is invalid D will broadcast $\langle o'' \rangle$ is invalid, and if $\langle o'' \rangle$ is valid then $\langle o \rangle - \langle o' \rangle + \langle o'' \rangle$ must be an invalid additive sharing of 0 since $\langle o \rangle$ is invalid and $\langle o' \rangle$ is derived robustly due to the robust reconstruction of the Shamir sharings. Also, D can broadcast at least one of the messages in Step 6. Otherwise, all parties have sent correct sharings and all parties' communication is consistent, so $\sum_{i=1}^{n'} \hat{o}_j^{(i)} = 0$ for all $j \in [n']$, which is contradictory to $\sum_{j=1}^{n'} \sum_{i=1}^{n'} \hat{o}_j^{(i)} \neq 0$. By the assumption, D should not have broadcast the message $(P_i, \text{incorrect})$ in Step 6, or otherwise P_i must be corrupted. It follows that D broadcast the message $(P_{j_1}, P_{j_2}, \hat{o}_{j_1}^{(j_2)}, (\hat{o}_{j_1}^{(j_2)})', \text{inconsistent})$ in Step 6. Because P_{j_1} and P_{j_2} cannot both be honest for D to broadcast this message, and by the assumption at least one of them is honest, so exactly one of P_{j_1} and P_{j_2} is corrupted.

- If P_{j_1} is corrupted, then $E = \{D, P_{j_2}\}$ according to the assumption. So P_{j_1} agrees with the broadcast message, and P_{j_2} does not agree with the broadcast message. However, P_{j_2} will disagree only if the message $(\hat{o}_{j_1}^{(j_2)})'$ broadcast by D is not the message that P_{j_2} has sent to D , which is not possible since D is honest.
- If P_{j_2} is corrupted, then $E = \{D, P_{j_1}\}$ according to the assumption. So P_{j_1} does not agree with the broadcast message. However, P_{j_1} will disagree only if the message $\hat{o}_{j_1}^{(j_2)}$ broadcast by D is not the message that P_{j_1} has sent to D , which is not possible since D is honest.

When D is corrupted, by the assumption D must broadcast the message $(P_{j_1}, P_{j_2}, \hat{o}_{j_1}^{(j_2)}, (\hat{o}_{j_1}^{(j_2)})', \text{inconsistent})$ in Step 6. Also, we must have P_{j_1} and P_{j_2} are honest, and they all agree with the broadcast message according to the assumption. However, this is not possible because honest parties $\hat{o}_{j_1}^{(j_2)}$ that P_{j_1} claims to have sent to P_{j_2} will always equal to $(\hat{o}_{j_1}^{(j_2)})'$ that P_{j_2} claims to have received from P_{j_1} .

Above all, in all cases, the protocol is always able to find a semi-corrupted pair.

F.2 Proof of Lemma 10

Proof. In Step 2, if the functionality $\mathcal{F}_{\text{VerifyBC}}$ outputs **consistent**, all parties will hold consistent values of $\mathbf{x}^{(i)} + \mathbf{a}^{(i)}$ and $\mathbf{y}^{(i)} + \mathbf{b}^{(i)}$ for all $i \in [N]$. If the output of $\mathcal{F}_{\text{VerifyBC}}$ is a semi-corrupted pair, it is guaranteed by the functionality that at least one party of the semi-corrupted pair is corrupted.

If all parties do not halt on Step 2, and if all the parties hold the correct $\mathbf{x}^{(i)} + \mathbf{a}^{(i)}$ and $\mathbf{y}^{(i)} + \mathbf{b}^{(i)}$ for all $i \in [N]$, then it is clear the protocol will generate random output Shamir sharings for all gate groups correctly, and all parties will pass the verification.

Next we need only show that if all parties does not halt on Step 2, and if there exists some $\mathbf{x}^{(i)} + \mathbf{a}^{(i)}$ or $\mathbf{y}^{(i)} + \mathbf{b}^{(i)}$ is not correct, then all parties will output a semi- where at least one is corrupted.

Consider the incorrect entry of $\mathbf{x}^{(i)} + \mathbf{a}^{(i)}$ or $\mathbf{y}^{(i)} + \mathbf{b}^{(i)}$ that has the smallest topological order, W.L.O.G. we suppose the entry is $x_{j_0}^{(i_0)} + a_{j_0}^{(i_0)}$. Since from Step 3 to Step 6, all the computation that derives $\tilde{\mathbf{x}}^{(i_0)} + \mathbf{a}^{(i_0)}$ is robust and only involves the correct Shamir secret sharings, we have that $\tilde{x}_{j_0}^{(i_0)} + a_{j_0}^{(i_0)}$ should be the correct value of $x_{j_0}^{(i_0)} + a_{j_0}^{(i_0)}$. Therefore, in Step 7 all parties can detect this incorrect entry and they will proceed to Step 7.b.

If there already exist some eliminated parties, *i.e.* $S \neq \emptyset$ in Step 10.(b), it is guaranteed by the previous functionalities that at least half of S are corrupted. Otherwise, all active parties will run $\Pi_{\text{FaultDetection}}$.

Note that the extended additive sharing of $x_{j_0}^{(i_0)} + a_{j_0}^{(i_0)}$ is calculated correctly. Also note that all parties that holds the shares of $\langle x_{j_0}^{(i_0)} + a_{j_0}^{(i_0)} \rangle$ from the beginning of Π_{Eval} remain active in this step. Since the dealer D has all parties' shares of $\langle x_{j_0}^{(i_0)} \rangle + \langle a_{j_0}^{(i_0)} \rangle + \langle o_{j_0}^{(i_0)} \rangle$, and $\langle x_{j_0}^{(i_0)} \rangle + \langle a_{j_0}^{(i_0)} \rangle$ is derived from the extended additive sharing $\langle x_{j_0}^{(i_0)} + a_{j_0}^{(i_0)} \rangle$, by Lemma 3 all parties can call $\Pi_{\text{FaultDetection}}$ to output a set of eliminated parties where at least half of the set are corrupted.

F.3 Proof of Lemma 4

Proof. Let \mathcal{A} denote the adversary. We will construct a simulator Sim to simulate the behaviors of honest parties for the adversary. Recall that $\mathcal{C}_{\text{active}}$ denotes the active corrupted parties and that $\mathcal{H}_{\text{active}}$ denotes the active honest parties.

The correctness of Π_{Main} follows the correctness of the protocol Π_{Verify} (Lemma 10) and the correctness of network routing (Theorem 2).

Simulation for Π_{Main} . Sim simulates Π_{Main} as follows:

1. In Step 1, Sim does nothing.
2. In Step 2, Sim simulates Π_{Input} as follows:
 - (a) In Step 1.(b) of Π_{Input} , Sim faithfully emulates $\mathcal{F}_{\text{InputShamir}}$.
 - (b) In Step 1.(c) of Π_{Input} , Sim faithfully emulates $\mathcal{F}_{\text{RandShamir}}$.
 - (c) In Step 1.(d) of Π_{Input} , if Client is corrupted, Sim learns y_1, \dots, y_N from \mathcal{A} . Then Sim samples $\mathbf{r}_i \in (\mathbb{Z}/2^k\mathbb{Z})^\ell$, and then send $y_i + \phi(\mathbf{r}_i)$ to the active corrupted parties on behalf of $\mathcal{F}_{\text{OpenPub}}$ for all $i \in [N]$. If Client is honest, Sim samples random $\tilde{\mathbf{r}}_i \in (\mathbb{Z}/2^k\mathbb{Z})^\ell$, and sends $\phi(\tilde{\mathbf{r}}_i)$ to \mathcal{A} on behalf of $\mathcal{F}_{\text{OpenPub}}$ for all $i \in [N]$.
 - (d) In Step 1.(e) of Π_{Input} , Sim does nothing.
 - (e) In Step 2 of Π_{Input} , Sim faithfully emulates $\mathcal{F}_{\text{FanOut}}$.
 - (f) In Step 3 of Π_{Input} , Sim faithfully emulates $\mathcal{F}_{\text{Permute}}$.
3. In Step 3.(a), Sim simulates Π_{Eval} for **seg** as follows:
 - (a) In Step 2 of Π_{Eval} , Sim faithfully emulates $\mathcal{F}_{\text{RandZeroAdd}}$.
 - (b) In Step 3 of Π_{Eval} , Sim faithfully emulates $\mathcal{F}_{\text{TripleAdd}}$ and $\mathcal{F}_{\text{TripleMult}}$.
 - (c) In Step 5.(b) of Π_{Eval} , if D is corrupted, Sim samples uniformly random values in $\mathbb{Z}/2^k\mathbb{Z}$ as the shares of $\langle x \rangle + \langle a_j \rangle + \langle o_1 \rangle$ and $\langle y \rangle + \langle b_j \rangle + \langle o_2 \rangle$ held by active honest parties whose indices are in $\{1, 2, \dots, t+1\}$. If D is honest, Sim follows the protocol.
 - (d) In step 5.(c) of Π_{Eval} , if D is corrupted, Sim follows the protocol. If D is honest, Sim samples two uniformly random elements in $\mathbb{Z}/2^k\mathbb{Z}$ as u and v sent to active corrupted parties.
4. In Step 3.(b), Sim simulates Π_{Verify} for **seg** as follows:
 - (a) In Step 2 of Π_{Verify} , Sim faithfully emulates $\mathcal{F}_{\text{VerifyBC}}$.

- (b) In Step 4 of Π_{Verify} , Sim faithfully emulates $\mathcal{F}_{\text{Mult}}$ and then faithfully emulates $\mathcal{F}_{\text{ReEncode}}$.
- (c) In Step 5 of Π_{Verify} , Sim faithfully emulates the functionalities $\mathcal{F}_{\text{FanOut}}$, $\mathcal{F}_{\text{Permute}}$, $\mathcal{F}_{\text{Collect}}$ and $\mathcal{F}_{\text{Permute}}$ in the protocol $\Pi_{\text{NetworkRouting}}$.
- (d) In Step 6 of Π_{Verify} , recall that $\mathbf{x}^{(i)} + \mathbf{a}^{(i)}$ and $\mathbf{y}^{(i)} + \mathbf{b}^{(i)}$ are computed from the values sent by D in Π_{Eval} . Let $\tilde{u}_j^{(i)} := \tilde{x}_j^{(i)} + a_j^{(i)}$ denote the j -th wire of $\tilde{\mathbf{x}}^{(i)} + \mathbf{a}^{(i)}$, and let $u_j^{(i)} := \tilde{x}_j^{(i)} + a_j^{(i)}$ denote the j -th wire of $\mathbf{x}^{(i)} + \mathbf{a}^{(i)}$. Similarly, let $\tilde{y}_j^{(i)} + b_j^{(i)}$ denote the j -th wire of $\tilde{\mathbf{y}}^{(i)} + \mathbf{b}^{(i)}$. Note that $u_j^{(i)}$ is reconstructed by D at Step 5.(c) of Π_{Eval} . Also note that $u_j^{(i)}$ should equal $\tilde{u}_j^{(i)}$ if the reconstruction is correct, because of the correctness of network routing.
 If D is corrupted, Sim simulates $\tilde{\mathbf{x}}^{(i)} + \mathbf{a}^{(i)}$ and $\tilde{\mathbf{y}}^{(i)} + \mathbf{b}^{(i)}$ as follows. Since D is corrupted, we must have that \mathcal{A} locally possesses this correct value $\tilde{u}_j^{(i)}$ at Step 5.(b) of Π_{Eval} . Therefore, Sim gets $\tilde{u}_j^{(i)} = \tilde{x}_j^{(i)} + a_j^{(i)}$ from \mathcal{A} for all $i \in [N]$ and for all $j \in [\ell]$. Similarly, Sim can get $\tilde{y}_j^{(i)} + b_j^{(i)}$ from \mathcal{A} for all $i \in [N]$ and for all $j \in [\ell]$. Then Sim sends $\phi(\tilde{\mathbf{x}}^{(i)} + \mathbf{a}^{(i)})$ and $\phi(\tilde{\mathbf{y}}^{(i)} + \mathbf{b}^{(i)})$ to the active corrupted parties on behalf of $\mathcal{F}_{\text{OpenPub}}$ for all $i \in [N]$.
 If D is honest, Sim simulates $\tilde{\mathbf{x}}^{(i)} + \mathbf{a}^{(i)}$ and $\tilde{\mathbf{y}}^{(i)} + \mathbf{b}^{(i)}$ as follows. At Step 5.(b) of Π_{Eval} , \mathcal{A} possesses the additive error introduced to $\tilde{u}_j^{(i)}$, which we suppose is $d_j^{(i)} \in \mathbb{Z}/2^k\mathbb{Z}$. Therefore, Sim receives $d_j^{(i)}$ from \mathcal{A} , and sets $\tilde{u}_j^{(i)} = u_j^{(i)} - d_j^{(i)}$ for all $i \in [N]$ and for all $j \in [\ell]$. Similarly, Sim can get $\tilde{y}_j^{(i)} + b_j^{(i)}$ from \mathcal{A} for all $i \in [N]$ and for all $j \in [\ell]$. Then Sim sends $\phi(\tilde{\mathbf{x}}^{(i)} + \mathbf{a}^{(i)})$ and $\phi(\tilde{\mathbf{y}}^{(i)} + \mathbf{b}^{(i)})$ to the active corrupted parties on behalf of $\mathcal{F}_{\text{OpenPub}}$ for all $i \in [N]$.
- (e) In Step 7.(a) of Π_{Verify} , Sim faithfully emulates $\mathcal{F}_{\text{FanOut}}$, and then faithfully emulates $\mathcal{F}_{\text{Permute}}$.
- (f) In Step 7.(b) of Π_{Verify} , it is clear that Sim can locate the first wire that is incorrect, denoted by $\langle s \rangle$ and $\langle o \rangle$. Then Sim simulates $\Pi_{\text{FaultDetection}}$ as follows:
- i. In Step 1 of $\Pi_{\text{FaultDetection}}$, if P_i is corrupted, Sim receives from \mathcal{A} the shares of $\langle o_i'' \rangle$ for active honest parties, and sends them to active honest parties. If P_i is honest, Sim samples uniformly random elements in $\mathbb{Z}/2^k\mathbb{Z}$ as the shares of $\langle o_i'' \rangle$ held by the active corrupted parties whose indices are in $\{1, 2, \dots, t+1\}$.
 - ii. In Step 2 of $\Pi_{\text{FaultDetection}}$, if D is corrupted, Sim samples a random additive zero sharing as the sharing $\sum_{i \in \mathcal{H}_{\text{active}}} \langle o_i'' \rangle$. Then Sim sets $\langle o'' \rangle = \sum_{i \in \mathcal{H}_{\text{active}}} \langle o_i'' \rangle + \sum_{i \in \mathcal{C}_{\text{active}}} \langle o_i'' \rangle$ as the communication from active honest parties to D . Afterwards, Sim simulates all the communication in the byzantine agreement protocol.
 If D is honest, Sim receives from \mathcal{A} the shares of $\langle o'' \rangle$ sent by active corrupted parties. Sim follows the protocol to send active corrupted parties' shares of $\langle o'' \rangle$ to D . Then Sim checks if $\sum_{i \in \mathcal{C}_{\text{active}}} \langle o_i'' \rangle$ sent by active corrupted parties equals $\sum_{i \in \mathcal{C}_{\text{active}}} \langle o_i'' \rangle$ actually held by active corrupted parties. If so, Sim learns that D broadcasts that $\langle o'' \rangle$ is valid. Otherwise, Sim learns that D broadcasts that $\langle o'' \rangle$ is invalid. Next Sim simulates the byzantine agreement protocol according to what is broadcast by D .
 - iii. In Step 4 of $\Pi_{\text{FaultDetection}}$, Sim faithfully emulates $\mathcal{F}_{\text{RandParity}}$.
 If D is corrupted, note that Sim has already sample s in Step 3.(c). Then Sim samples the shares of $\langle s \rangle + [p]_t$ held by active honest parties that are compatible with the shares of $\langle s \rangle + [p]_t$ held by active corrupted parties and with s . Sim sends the shares of $\langle s \rangle + [p]_t$ held by active honest parties to \mathcal{A} .
 If D is honest, Sim follows the protocol.
 - iv. In Step 5 of $\Pi_{\text{FaultDetection}}$, if D is corrupted, Sim samples random additive zero sharings for $\{\langle \hat{o}_i \rangle\}_{i \in \mathcal{H}_{\text{active}}}$, such that $\sum_{i \in \mathcal{H}_{\text{active}}} \langle \hat{o}_i \rangle = \langle \hat{o} \rangle - \sum_{i \in \mathcal{C}_{\text{active}}} \langle \hat{o}_i \rangle$, and $\langle \hat{o}_i \rangle$ is compatible with the simulated active corrupted parties' shares. For all $P_i \in \mathcal{H}_{\text{active}}$, Sim gets $\hat{o}_i^{(j)}$ from the sampled $\langle \hat{o}_i \rangle$, and Sim gets $(\hat{o}_j^{(i)})'$ from the sampled sharing $\langle \hat{o}_j \rangle$ if $j \in \mathcal{H}_{\text{active}}$, or from \mathcal{A} if $j \in \mathcal{C}_{\text{active}}$.
 If D is honest, Sim follows the protocol.
 - v. In Step 6 of $\Pi_{\text{FaultDetection}}$, if D is corrupted, since Sim has sampled all sharings, Sim can simply follow the protocol.
 If D is honest, Sim receives from \mathcal{A} all the incorrect or inconsistent messages. From these messages, Sim can determine D 's broadcast message. Then Sim follows the byzantine agreement protocol to broadcast the message.

- vi. In Step 7 of $\Pi_{\text{FaultDetection}}$, if $(P_i, \text{incorrect})$ is broadcast by D , Sim follows the protocol. If $(P_{j_1}, P_{j_2}, \hat{o}_{j_1}^{(j_2)}, (\hat{o}_{j_1}^{(j_2)})', \text{inconsistent})$ is broadcast by D , Sim can perfectly simulate if P_{j_1} or P_{j_2} agree as follows. If D is corrupted, since Sim has sampled all sharings, Sim follows the protocol to simulate if P_{j_1} or P_{j_2} agrees. If D is honest, Sim knows that the honest party between P_{j_1} and P_{j_2} will always agree, and receives from \mathcal{A} whether the corrupted party between P_{j_1} and P_{j_2} agrees. Then Sim follows the byzantine agreement protocol.
5. In Step 4, Sim simulates Π_{Output} as follows:
- In Step 1 of Π_{Output} , Sim faithfully emulates $\mathcal{F}_{\text{Collect}}$.
 - In Step 2 of Π_{Output} , Sim faithfully emulates $\mathcal{F}_{\text{Permute}}$.
 - In Step 3 of Π_{Output} , if Client is corrupted, Sim receives from $\mathcal{F}_{\text{Main}}$ the output towards Client. Then Sim samples the shares of the output Shamir sharings held by active honest parties that are compatible with the secrets and the active corrupted parties' shares, and then Sim sends them to Client. If Client is honest, Sim follows the protocol.

Hybrid Arguments. We show that Sim perfectly simulates the behaviors of honest parties.

Hybrid₀: The execution in the real world.

Hybrid₁: In this hybrid Sim faithfully emulates the functionalities $\mathcal{F}_{\text{InputShamir}}$, $\mathcal{F}_{\text{RandShamir}}$ in Step 1.(b) and Step 1.(c) of Π_{Input} . The distribution is the same as **Hybrid₀**.

Hybrid₂: In this hybrid Sim simulates Step 1.(d) of Π_{Input} . If Client is corrupted, the distribution is the same since Sim receives the true value of $\{y_i\}_{i=1}^N$ from \mathcal{A} , and samples from the same distribution for the secrets $\{\phi(\mathbf{r}_i)\}_{i=1}^N$. If Client is honest, since Client's input will always be valid and has the form of $y_i = \phi(\mathbf{x}_i)$ for all $i \in [N]$, we have the distribution of $y_i + \phi(\mathbf{r}_i)$ is the same as the distribution that Sim samples from for all $i \in [N]$. Therefore, the distribution is the same as **Hybrid₁**.

Hybrid₃: In this hybrid Sim faithfully emulates the functionalities $\mathcal{F}_{\text{FanOut}}$ and $\mathcal{F}_{\text{Permute}}$ in Step 2 and Step 3 of Π_{Input} . The distribution is the same as **Hybrid₂**.

Hybrid₄: Sim faithfully emulates $\mathcal{F}_{\text{RandZeroAdd}}$ in Step 2 of Π_{Eval} , and $\mathcal{F}_{\text{TripleAdd}}$ and $\mathcal{F}_{\text{TripleMult}}$ in Step 3 of Π_{Eval} . The distribution is the same as **Hybrid₃**.

Hybrid₅: In this hybrid, Sim simulates Step 5.(b) and 5.(c) of Π_{Eval} when D is corrupted. Note that the shares of $\langle x \rangle + \langle a_j \rangle + \langle o_1 \rangle$ and $\langle y \rangle + \langle b_j \rangle + \langle o_2 \rangle$ held by active corrupted parties are independent of the shares of $\langle x \rangle + \langle a_j \rangle + \langle o_1 \rangle$ and $\langle y \rangle + \langle b_j \rangle + \langle o_2 \rangle$ held by active honest parties and the secrets $x + a_j, y + b_j$. This is because $x + a_j$ and $y + b_j$ are uniformly random given the shares of $[\phi(\mathbf{a})]_t$ and $[\phi(\mathbf{b})]_t$ held by active corrupted parties, and because additive sharings have threshold t . Note that in Step 5.(c) Sim follows the protocol, which will not change the distribution. Therefore, the distribution is the same as **Hybrid₄**.

Hybrid₆: In this hybrid, Sim faithfully emulates the functionalities among Step 2 to Step 5 of Π_{Verify} when D is corrupted. Specifically, Sim faithfully emulates $\mathcal{F}_{\text{VerifyBC}}$ in Step 2 of Π_{Verify} , $\mathcal{F}_{\text{Mult}}$ and $\mathcal{F}_{\text{ReEncode}}$ in Step 4 of Π_{Verify} , and $\mathcal{F}_{\text{FanOut}}$, $\mathcal{F}_{\text{Permute}}$, $\mathcal{F}_{\text{Collect}}$, and $\mathcal{F}_{\text{Permute}}$ in Step 5 of Π_{Verify} . The distribution is the same as **Hybrid₅**.

Hybrid₇: In this hybrid, Sim simulates Step 9 of Π_{Verify} when D is corrupted. Because Sim has sampled the entire sharing of $\langle x \rangle + \langle a_j \rangle + \langle o_1 \rangle$ and $\langle y \rangle + \langle b_j \rangle + \langle o_2 \rangle$ when simulating Step 3.(c) of Π_{Eval} , Sim can locally generate the correct values $\tilde{\mathbf{x}}^{(i)} + \mathbf{a}^{(i)}$ and $\tilde{\mathbf{y}}^{(i)} + \mathbf{b}^{(i)}$ for all $i \in [N]$. Therefore, the distribution of $\phi(\tilde{\mathbf{x}}^{(i)} + \mathbf{a}^{(i)})$ and $\phi(\tilde{\mathbf{y}}^{(i)} + \mathbf{b}^{(i)})$ is the same as receiving the reconstructed value from $\mathcal{F}_{\text{OpenPub}}$ for all $i \in [N]$. It follows that the distribution of this hybrid is the same as **Hybrid₆**.

Hybrid₈: In this hybrid, Sim faithfully emulates $\mathcal{F}_{\text{FanOut}}$ and $\mathcal{F}_{\text{Permute}}$ in Step 10.(a) of Π_{Verify} when D is corrupted. The distribution is the same as **Hybrid₇**.

Hybrid₉: In this hybrid, Sim simulates Step 1 of $\Pi_{\text{FaultDetection}}$ when D is corrupted. Because additive zero sharings have threshold t , the shares of $\langle o_i'' \rangle$ held by active corrupted parties are uniformly random when P_i is honest, Since Sim samples active corrupted parties' shares uniformly randomly, the distribution of this hybrid is the same as **Hybrid₈**.

Hybrid₁₀: In this hybrid, Sim simulates Step 2 of $\Pi_{\text{FaultDetection}}$ when D is corrupted. Consider the zero additive sharing $\langle o_i'' \rangle$ for an honest party P_i . Its shares held by active honest parties are independent of its shares held by active corrupted parties. It follows that the shares of $\sum_{i \in \mathcal{H}_{\text{active}}} \langle o_i'' \rangle$ held by active honest parties are independent of the shares of $\langle o_i'' \rangle$ held by active corrupted parties

for all $i \in \mathcal{H}_{\text{active}}$. So the distribution of $\sum_{i \in \mathcal{H}_{\text{active}}} \langle o_i'' \rangle$ sampled by **Sim** is identical to the distribution of $\sum_{i \in \mathcal{H}_{\text{active}}} \langle o_i' \rangle$ sampled by active honest parties. Therefore, the distribution of this hybrid is the same as **Hybrid₉**.

Hybrid₁₁: In this hybrid, **Sim** simulates Step 4 of $\Pi_{\text{FaultDetection}}$ when D is corrupted. **Sim** first faithfully emulates $\mathcal{F}_{\text{RandParity}}$ in Step 4 of $\Pi_{\text{FaultDetection}}$. After this step, the distribution is the same as **Hybrid₁₀**.

Suppose the sharing $\langle s \rangle + [p]_t$ sampled by **Sim** is denoted by $[w]_t$. We argue that $[w]_t$ has the same distribution as the real sharing $\langle s \rangle + [p]_t$ as follows.

- Given the secret s , the sharing $\langle s \rangle + \langle o \rangle$ is independent of the sharing $\langle s \rangle$ and is therefore independent of the sharing of $\langle s \rangle + [p]_t$. It follows that the sharing $\langle s \rangle + \langle o \rangle$ sampled by **Sim** is independent of the distribution of $\langle s \rangle + [p]_t$.
- Note that $[p]_t$ is a uniformly random parity sharing given active corrupted parties' shares, and the sampled sharing $[w]_t$ is valid if and only if $[w]_t - \langle s \rangle$ is a parity sharing. So $\langle s \rangle + [p]_t$ has the same distribution as $[w]_t$.

Therefore, the distribution of this hybrid is the same as **Hybrid₁₀**.

Hybrid₁₂: In this hybrid, **Sim** simulates Step 5, Step 6, and Step 7 of $\Pi_{\text{FaultDetection}}$ when D is corrupted. If $\langle \hat{o} \rangle = \langle o'' \rangle$, **Sim** has sampled $\sum_{i \in \mathcal{H}_{\text{active}}} \langle o_i'' \rangle$. If $\langle \hat{o} \rangle = \langle o \rangle - \langle o' \rangle + \langle o'' \rangle$, note that **Sim** has sampled $\langle s \rangle + [p]_t$, **Sim** can derive $\sum_{i \in \mathcal{H}_{\text{active}}} (\langle o_i \rangle - \langle o_i' \rangle + \langle o_i'' \rangle)$ from the sampled sharings. In either case, **Sim** has $\sum_{i \in \mathcal{H}_{\text{active}}} \langle \hat{o}_i \rangle$, and it is uniformly random given the shares of $\sum_{i \in \mathcal{H}_{\text{active}}} \langle \hat{o}_i \rangle$ held by active corrupted parties.

For $i \in \mathcal{H}_{\text{active}}$, $\langle \hat{o}_i \rangle$ is a uniformly random additive zero sharing given the shares held by active corrupted parties. Therefore, when **Sim** uniformly samples $\langle \hat{o}_i \rangle$ for all $i \in \mathcal{H}_{\text{active}}$ given $\sum_{i \in \mathcal{H}_{\text{active}}} \langle \hat{o}_i \rangle$, the distribution is the same as **Hybrid₁₁**.

Since **Sim** follows the rest of Steps of $\Pi_{\text{FaultDetection}}$, and the sharing distribution is the same, we have that the distribution of this hybrid is the same as **Hybrid₁₁**.

Hybrid₁₃: In this hybrid, **Sim** simulates Step 5.(b) and Step 5.(c) of Π_{Eval} when D is honest. We have the secret $x + a_j$ and $y + b_j$ are uniformly random, and are independent of the shares of $\langle x \rangle + \langle a_j \rangle + \langle o_1 \rangle$ and $\langle y \rangle + \langle b_j \rangle + \langle o_2 \rangle$ held by active corrupted parties. Since **Sim** samples $x + a_j$ and $y + b_j$ uniformly randomly, we have the distribution of this hybrid is the same as **Hybrid₁₂**.

Hybrid₁₄: In this hybrid, **Sim** faithfully emulates the functionalities among Step 2 to Step 8 of Π_{Verify} when D is honest. Specifically, **Sim** faithfully emulates $\mathcal{F}_{\text{VerifyBC}}$ in Step 2 of Π_{Verify} , $\mathcal{F}_{\text{Mult}}$ and $\mathcal{F}_{\text{ReEncode}}$ in Step 4 of Π_{Verify} , $\mathcal{F}_{\text{FanOut}}$ in Step 5 of Π_{Verify} , $\mathcal{F}_{\text{Permute}}$ in Step 6 and Step 8 of Π_{Verify} , and $\mathcal{F}_{\text{Collect}}$ in Step 7 of Π_{Verify} . The distribution is the same as **Hybrid₁₃**.

Hybrid₁₅: In this hybrid, **Sim** simulates Step 9 of Π_{Verify} when D is honest. Because the network routing is correct, and **Sim** can receive from \mathcal{A} the additive errors introduced to $\tilde{u}_j^{(i)}$, i.e. $u_j^{(i)} - \tilde{u}_j^{(i)}$, for all $i \in [N]$ and for all $j \in [\ell]$, **Sim** correctly generates $\tilde{u}_j^{(i)}$ from the $u_j^{(i)}$ that it sampled when simulating Step 5.(d) of Π_{Eval} . Therefore, the distribution of $\phi(\tilde{\mathbf{x}}^{(i)} + \mathbf{a}^{(i)})$ and $\phi(\tilde{\mathbf{y}}^{(i)} + \mathbf{b}^{(i)})$ is the same as receiving the reconstructed value from $\mathcal{F}_{\text{OpenPub}}$ for all $i \in [N]$. It follows that the distribution of this hybrid is the same as **Hybrid₁₄**.

Hybrid₁₆: In this hybrid, **Sim** faithfully emulates $\mathcal{F}_{\text{FanOut}}$ and $\mathcal{F}_{\text{Permute}}$ in Step 10.(a) of Π_{Verify} when D is honest. The distribution is the same as **Hybrid₁₅**.

Hybrid₁₇: In this hybrid, **Sim** simulates Step 1 of $\Pi_{\text{FaultDetection}}$ when D is honest. Because additive zero sharings have threshold t , the shares of $\langle o_i'' \rangle$ held by active corrupted parties are uniformly random when P_i is honest, Since **Sim** samples active corrupted parties' shares uniformly randomly, the distribution of this hybrid is the same as **Hybrid₁₆**.

Hybrid₁₈: In this hybrid, **Sim** simulates Step 2, Step 4, Step 5, Step 6, and Step 7 of $\Pi_{\text{FaultDetection}}$ when D is honest. In Step 2 of $\Pi_{\text{FaultDetection}}$, because **Sim** can perfectly simulate the value broadcast by D , the distribution is the same. In Step 4 and Step 5 of $\Pi_{\text{FaultDetection}}$, because **Sim** follows the protocol, the distribution remains the same. In Step 6 of $\Pi_{\text{FaultDetection}}$, because **Sim** receives from \mathcal{A} the incorrect or inconsistent message, **Sim** can perfectly simulate D 's broadcast message. In Step 7 of $\Pi_{\text{FaultDetection}}$, if $(P_{j_1}, P_{j_2}, \hat{o}_{j_1}^{(j_2)}, (\hat{o}_{j_1}^{(j_2)})', \text{inconsistent})$ is broadcast, the honest party between P_{j_1} and P_{j_2} will always agree, and the other party is the corrupted party. So **Sim** can perfectly simulate the message broadcast by P_{j_1} and P_{j_2} . Above all, the distribution of this hybrid is the same as **Hybrid₁₇**.

Hybrid₁₉: In this hybrid, **Sim** faithfully emulates the functionalities $\mathcal{F}_{\text{Collect}}$ and $\mathcal{F}_{\text{Permute}}$ in Step 1 and Step 2 of Π_{Output} . The distribution is the same as **Hybrid₁₈**.

Hybrid₂₀: In this hybrid, **Sim** simulates Step 3 of Π_{Output} . Suppose the output sharings for **Client** are $[\phi(\mathbf{z}_1)]_t, \dots, [\phi(\mathbf{z}_N)]_t$. If **Client** is corrupted, it is guaranteed by $\mathcal{F}_{\text{Permute}}$ that the shares of $[\phi(\mathbf{z}_1)]_t, \dots, [\phi(\mathbf{z}_N)]_t$ held by active honest parties are uniformly sampled given the secrets $\phi(\mathbf{z}_1), \dots, \phi(\mathbf{z}_N)$ and the shares held by active corrupted parties. Therefore, the shares held by active honest parties sampled by **Sim** have the same distribution as the real shares held by active honest parties in the view of \mathcal{A} . If **Client** is honest, since **Sim** follows the protocol, the distribution remains the same. Therefore, the distribution of this hybrid is the same as **Hybrid₁₉**.

Note that this hybrid is the execution in the ideal world.

G List of Functionalities and Protocols

Functionalities

Functionality 1: $\mathcal{F}_{\text{Main}}(C)$	9
Functionality 2: $\mathcal{F}_{\text{OpenPub}}(N)$	24
Functionality 3: $\mathcal{F}_{\text{Mult}}(N)$	24
Functionality 4: $\mathcal{F}_{\text{ReEncode}}(N)$	24
Functionality 5: $\mathcal{F}_{\text{VerifyBC}}(N)$	24
Functionality 6: $\mathcal{F}_{\text{InputShamir}}(N)$	25
Functionality 7: $\mathcal{F}_{\text{Rand}}(N)$	27
Functionality 8: $\mathcal{F}_{\text{RandRobust}}(N)$	30
Functionality 9: $\mathcal{F}_{\text{RandShamir}}(N)$	31
Functionality 10: $\mathcal{F}_{\text{RandShPacked}}(d, h, N)$	31
Functionality 11: $\mathcal{F}_{\text{RandZeroAdd}}(N)$	32
Functionality 12: $\mathcal{F}_{\text{RandZeroSh}}(d, h, N)$	33
Functionality 13: $\mathcal{F}_{\text{RandParity}}(N)$	34
Functionality 14: $\mathcal{F}_{\text{TripleAdd}}(N)$	34
Functionality 15: $\mathcal{F}_{\text{TripleMult}}(N)$	35
Functionality 16: $\mathcal{F}_{\text{RandSharTrans}}(N)$	40
Functionality 17: $\mathcal{F}_{\text{LinearTrans}}(N)$	44
Functionality 18: $\mathcal{F}_{\text{FanOut}}(N)$	46
Functionality 19: $\mathcal{F}_{\text{Permute}}(N)$	47
Functionality 20: $\mathcal{F}_{\text{Collect}}(N)$	47

Protocols

Protocol 1: $\Pi_{\text{Eval}}(\text{seg})$	14
Protocol 2: $\Pi_{\text{NetworkRouting}}(\text{seg})$	17
Protocol 3: $\Pi_{\text{FaultDetection}}(D, \langle s \rangle + \langle o \rangle, \langle s \rangle)$	18
Protocol 4: Π_{Input}	19
Protocol 5: Π_{Output}	20
Protocol 6: $\Pi_{\text{Main}}(C)$	20
Protocol 7: $\Pi_{\text{Rand}}(N)$	28
Protocol 8: $\Pi_{\text{RandRobust}}(N)$	30
Protocol 9: $\Pi_{\text{TripleAdd}}(N)$	35
Protocol 10: $\Pi_{\text{TripleMult}}(N)$	35
Protocol 11: $\Pi_{\text{Verify}}(\text{seg}, S_0)$	36
Protocol 12: $\Pi_{\text{RandSharTrans}}(\Pi_1, \dots, \Pi_{\ell_r})$	41
Protocol 13: $\Pi_{\text{LinearTrans}}(N, \{[\phi(\mathbf{x}_i)]_t\}_{i=1}^N, \{L_i\}_{i=1}^N)$	45
Protocol 14: $\Pi_{\text{FanOut}}(N, \{([\phi(\mathbf{z}^{(i)})]_t, (n_j^{(i)})_{j=1}^{m^{(i)}})\}_{i=1}^N)$	46
Protocol 15: $\Pi_{\text{Permute}}(N, \{[\phi(\mathbf{x}_i)]_t\}_{i=1}^N, \{p_i\}_{i=1}^N)$	47
Protocol 16: $\Pi_{\text{Collect}}(N, \{[\phi(\mathbf{x}_i)]_t\}_{i=1}^M, \{(u_{i,j}, w_{i,j})\}_{i \in [N], j \in [\ell]})$	48