

SPD \mathbb{Z}_{2^k} : Efficient MPC mod 2^k for Dishonest Majority

Ronald Cramer¹, Ivan Damgård², Daniel Escudero², Peter Scholl², and
Chaoping Xing³

¹ CWI, Amsterdam & Leiden University

² Aarhus University

³ Nanyang Technological University, Singapore

Abstract. Most multi-party computation protocols allow secure computation of arithmetic circuits over a finite field, such as the integers modulo a prime. In the more natural setting of integer computations modulo 2^k , which are useful for simplifying implementations and applications, no solutions with active security are known unless the majority of the participants are honest.

We present a new scheme for information-theoretic MACs that are homomorphic modulo 2^k , and are as efficient as the well-known standard solutions that are homomorphic over fields. We apply this to construct an MPC protocol for dishonest majority in the preprocessing model that has efficiency comparable to the well-known SPDZ protocol (Damgård et al., CRYPTO 2012), with operations modulo 2^k instead of over a field. We also construct a matching preprocessing protocol based on oblivious transfer, which is in the style of the MASCOT protocol (Keller et al., CCS 2016) and almost as efficient.

1 Introduction

In the context of secure multi-party computation (MPC) there are n parties P_1, \dots, P_n who want to compute a function $f : \mathcal{R}^n \rightarrow \mathcal{R}^n$ securely on an input (x_1, \dots, x_n) , where each party P_i holds x_i , without revealing the inputs to each other and only by exchanging messages between them. The main security guarantee we would like to achieve is that at the end of the interaction each party P_i only learns x_i and the i -th component of $f(x_1, \dots, x_n)$, and nothing else. This should hold even if an adversary corrupts some of the parties and, in case of active or malicious corruption, takes control of the corrupted parties and have them do what the adversary wants. These ideas are formalized by requiring that using the protocol should be essentially equivalent to having a trusted third party compute the function. For such a formalization see, for example, the Universal Composability Framework (UC) [4].

It is well known that the hardest case to handle efficiently is the dishonest majority case, where $t \geq n/2$ parties are actively corrupted. Here we cannot guarantee that the protocol terminates correctly, and we have to use computationally heavy public-key technology — unconditional security is not possible in

this scenario. However, in a recent line of work [2,9], it was observed that we can push the use of public-key tools into a preprocessing phase, where one does not need to know the inputs or even the function to be computed. This phase produces “raw material” (correlated randomness) that can be used later in an online phase to compute the function much more efficiently and with unconditional security (given the correlated randomness).

In all existing protocols that handle a dishonest majority and active corruptions, the function being computed must be expressed in terms of arithmetic operations (i.e. additions and multiplications) over a finite field, such as the integers modulo a prime. However, in many applications one would like to use numbers modulo some M that is chosen by the application and is not necessarily a prime. In particular, $M = 2^k$ is interesting because computation modulo 2^k matches closely what happens on standard CPUs and hence protocol designers can take advantage of the tricks found in this domain. For instance, functions containing comparisons and bitwise operations are typically easier to implement using arithmetic modulo 2^k ; these kinds of operations are expensive to emulate with finite field arithmetic, and also very common in applications of MPC such as secure benchmarking based on linear programming [6]. This has been done successfully by the team behind the Sharemind suite of protocols [3], which allows bitwise operations and integer arithmetic mod 2^{32} . However, in their basic setting, they could only get a passively secure solution: here, even corrupt players are assumed to follow the protocol. Also, the security of Sharemind completely breaks down if half (or more) of the players are corrupted, and the efficiency does not scale well beyond three parties.

To obtain active security over fields, the main idea of modern protocols is to use unconditionally secure message authentication codes (MACs) to prevent players from lying about the data they are given in the preprocessing phase. A typical example is the SPDZ protocol [9,7], where security reduces to the following game: we have a data value x , a random MAC key α and a MAC $m = \alpha x$, all in some finite field \mathbb{F} . The adversary is given x but not α or αx . He may now specify errors to be added to x , α and m , and we let x', α', m' be the resulting values. The adversary wins if $x \neq x'$ and $m' = \alpha' x'$. It is easy to see that the adversary must guess α to win, and so the probability of winning is $1/|\mathbb{F}|$. This authentication scheme is additively homomorphic, which is exploited heavily in the SPDZ protocol and is crucial for its efficiency.

However, the security proof depends on the fact that any non-zero value in \mathbb{F} is invertible, and it is easy to see that if we replace the field by a ring, say \mathbb{Z}_{2^k} , then the adversary can cheat with large probability. For instance, in the ring \mathbb{Z}_{2^k} he can choose $x' = x + 2^{k-1}$ and cheat with probability $1/2$. Up to now, it has been an open problem to design a homomorphic authentication scheme that would work over \mathbb{Z}_{2^k} or more generally \mathbb{Z}_M for any M , and is as efficient as the SPDZ scheme.

1.1 Our contributions

In this paper we solve the above question: we design a new additively homomorphic authentication scheme that works in \mathbb{Z}_{2^k} ⁴, and is as efficient as the standard solution over a field. The main idea is to choose the MAC key α randomly in \mathbb{Z}_{2^s} , where s is the security parameter, and compute the MAC αx in $\mathbb{Z}_{2^{k+s}}$. We explain below why this helps. We also design a method for checking large batches of MACs with a communication complexity that does not depend on the size of the batch. We believe that these techniques will be of independent interest.

We then use the MAC scheme to design a SPDZ-style online protocol that securely computes an arithmetic circuit over \mathbb{Z}_{2^k} with statistical security, assuming access to a preprocessing functionality that outputs multiplication triples in a suitable format. The total computational work done is dominated by $O(|C|n)$ elementary operations in the ring $\mathbb{Z}_{2^{k+s}}$, where C is the circuit to be computed. So if $k \geq s$, the work needed per player is equal to the work needed to compute C in the clear, up to a constant factor — as is the case for the SPDZ protocol. As in other protocols from this line of work, the overhead becomes more significant when k is small. Each player stores data from the preprocessing of size $O(|C|(k+s))$ bits. The communication complexity is $O(|C|(k+s))$ bits per player, plus an overhead that does not depend on C . This is due to the batch-checking of MACs mentioned above.

Our final result is an implementation of the preprocessing functionality to generate multiplication triples. It has communication complexity $O((k+s)^2)$ bits per multiplication gate, and is roughly as efficient as the MASCOT protocol [14], which is the state of the art for preprocessing over a field using oblivious transfer. Concretely, our triple generation protocol has around twice the communication cost of MASCOT, due to the overhead incurred when we have to work over larger rings in certain scenarios. However, this additional cost seems like a small price to pay for the potential benefits to applications from working modulo 2^k instead of in a field.

1.2 Overview of our techniques

For the authentication scheme, as mentioned, we have a data item $x \in \mathbb{Z}_{2^{k+s}}$, a key $\alpha \in \mathbb{Z}_{2^{k+s}}$ and we define the MAC as $m = \alpha x \pmod{2^{k+s}}$. Note that we want to authenticate k -bit values, so although $x \in \mathbb{Z}_{2^{k+s}}$, only the least significant k bits matter. The adversary is given x , and specifies errors e_x, e_α, e_m , which define modified values $x' = x + e_x, \alpha' = \alpha + e_\alpha, m' = m + e_m$. He wins if $m' = \alpha' x' \pmod{2^{k+s}}$, but note that since we store data in the least significant k bits only, this is only a forgery if $e_x \pmod{2^k} \neq 0$. As we show in detail in Section 3, if the adversary wins, he is able to compute $e_x \alpha \pmod{2^{k+s}}$. From this, and $e_x \pmod{2^k} \neq 0$, it follows that the adversary can effectively guess $\alpha \pmod{2^s}$, which is only possible with probability 2^{-s} .

We also want to batch-check many MACs using only a small amount of communication. The SPDZ protocol [9] uses a method that basically takes a random

⁴ We use modulus 2^k throughout, but the scheme easily extends to any modulus.

linear combination of all messages and MACs and checks only the resulting message and MAC. Unfortunately, applying the analysis we just sketched to this scenario does not give a negligible probability of cheating, unless we ‘lift’ again and compute MACs modulo 2^{k+2s} , but then our storage and preprocessing costs would become significantly bigger. We provide a more complicated but tighter analysis showing that we can still compute MACs mod 2^{k+s} and the batch checking works with $2^{-s+\log s}$ error probability, so we only need increase s by a few bits.

Using these MACs, we can create an information-theoretically secure MPC protocol over \mathbb{Z}_{2^k} in the preprocessing model, similar to the online phase of SPDZ from [7]. To implement the preprocessing phase, we follow the style of MASCOT [14], which uses oblivious transfer to produce shares of authenticated multiplication triples. We first design a protocol for authenticating values using correlated oblivious transfer, which allows creating the secret-shared MACs that will be added to the preprocessing data. This stage is similar to MASCOT, whereby first a passively secure protocol is used to compute shares of the MACs αx_i , for each value x_i that is to be authenticated, and then a random linear combination of these values is opened, and the resulting MAC checked for correctness. The main change we need to make here is that, depending on the size of the x_i ’s being authenticated, we may need to first compute the MACs over a larger ring in order to apply our analysis of taking random linear combinations.

Once the authentication scheme has been implemented, the main task is to create the multiplication triples needed in the online phase of our protocol. For this we also follow a similar approach to MASCOT, where the overall idea is that each party P_i chooses its shares (a^i, b^i) and then is engaged in an oblivious transfer subprotocol with P_j for each $j \neq i$, where shares of the cross products $a^i b^j$ and $a^j b^i$ are obtained. This yields shares of the product $(\sum_{i=1}^n a^i)(\sum_{j=1}^n b^j) = \sum_{i=1}^n a^i b^i + \sum_{i \neq j} (a^i b^j + a^j b^i)$, as required. Behind this simplification lies the problem that some information about the honest parties’ shares can be leaked to a cheating adversary. In MASCOT this potential leakage is mitigated by ‘spreading out’ the randomness by taking random linear combinations on correlated triples (with the same b value). When working over fields, the inner product yields a 2-universal hash function so the new distribution can be argued to be close to uniform using the Leftover Hash Lemma. However, this is not true anymore over rings like \mathbb{Z}_{2^k} . We overcome this issue by starting with triples where the shares of a are *bits* instead of ring elements, and then taking linear combinations over the bits. These combinations correspond to a subset sum over \mathbb{Z}_{2^k} , which is a 2-universal hash function, so allows for removing the leakage.

Additionally, random combinations are used in MASCOT to check the correctness of a triple by ‘sacrificing’ another one. The security argument is that if the adversary manages to authenticate an incorrect triple, then it will have to guess the randomness used in the sacrifice step, which is unlikely. This is argued by deriving an equation from which we can solve for the random value. In order

to extend this argument to the ring case, we use the technique sketched at the beginning of this section, working over $\mathbb{Z}_{2^{k+s}}$ to check correctness modulo 2^k .

Organization of this document. Section 2 introduces the notation we will use throughout this document. It also introduces the oblivious transfer and coin tossing functionalities, \mathcal{F}_{ROT} and $\mathcal{F}_{\text{Rand}}$, which constitute our most basic building blocks and will be used to implement the offline phase of our protocol. We then describe our information-theoretic MAC scheme in Section 3, and we show how to check correctness of several authenticated values assuming a functionality \mathcal{F}_{MAC} that generates keys and MACs. Next, in Section 4 we show how to use our scheme to realise the functionality $\mathcal{F}_{\text{Online}}$, i.e. to evaluate securely any arithmetic circuit modulo 2^k , in the preprocessing model.

The next two sections are concerned with the implementation of the preprocessing functionality $\mathcal{F}_{\text{Prep}}$. Section 5 deals with the implementation of the functionality \mathcal{F}_{MAC} , i.e. the distribution of the MAC key and the generation of MACs. Our construction is based on a primitive called vector Oblivious Linear Function Evaluation ($\mathcal{F}_{\text{VOLE}}$). This can be implemented using Correlated Oblivious Transfer ($\mathcal{F}_{\Delta\text{-OT}}$), which as we mention in that section can be implemented using our basic primitive \mathcal{F}_{ROT} . On the other hand, Section 6 builds on top of our MAC scheme and generates multiplication triples that will be used during the online phase of our protocol to evaluate multiplication gates. Finally, in Section 7 we provide an efficiency analysis of our protocol.

To help the reader, Fig. 1 illustrates the different dependencies among our functionalities, including the section where the protocol is described.

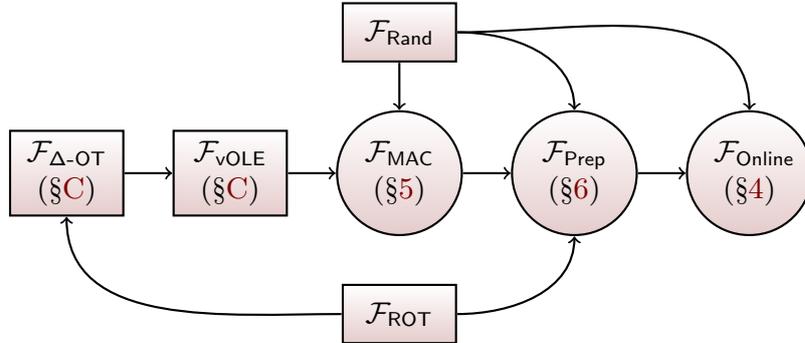


Fig. 1. Functionalities and their dependencies. An arrow from functionality A to functionality B means that B is realised in the A -hybrid model. Functionalities within circles represent the main contribution of our work.

Related work. There are only a few previous works that study MPC over rings, and none of these offer security against an active adversary who corrupts a dishonest majority of the parties. Cramer et al. showed how to construct actively secure MPC over black-box rings [5] using secret-sharing techniques for honest majority, but this is only a feasibility result and the concrete efficiency is not clear. As already mentioned, Sharemind [3] allows mixing of secure computation over the integers modulo 2^k with boolean computations, but is restricted to the three-party setting when at most one party is corrupted. In some settings Sharemind can also provide active security [18].

More recently, Damgård, Orlandi and Simkin [8] present a compiler that transforms a semi-honest secure protocol for t corruptions into a maliciously secure protocol that is secure against a smaller number of corruptions (approximately \sqrt{t}). This also works for protocols in the preprocessing model, but will always result in a protocol for honest majority, so they can tolerate a smaller number of corruptions. On the other hand, their compiler is perfectly secure, so it introduces no overhead that depends on the security parameter. Thus, their results are incomparable to ours.

2 Preliminaries

Notation We denote by \mathbb{Z}_M the set of integers x such that $0 \leq x \leq M - 1$. The congruence $x \equiv y \pmod{2^k}$ will be abbreviated as $x \equiv_k y$. We let $x \bmod M$ denote the remainder of x when divided by M , and we take this representative as an element of the set \mathbb{Z}_M . Given two vectors \mathbf{x} and \mathbf{y} of the same dimensions, $\mathbf{x} * \mathbf{y}$ denotes their component-wise product, $\langle \mathbf{x}, \mathbf{y} \rangle$ denotes their dot product and $\mathbf{x}[i]$ denotes the i -th entry of \mathbf{x} .

2.1 Oblivious Transfer and Coin Tossing Functionalities

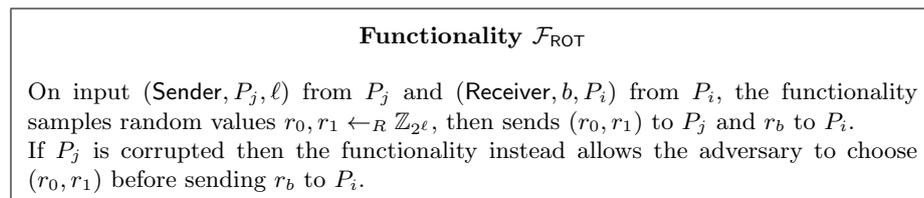


Fig. 2. Random Oblivious Transfer functionality between a sender and receiver

We use a standard functionality for oblivious transfer on random ℓ -bit strings, shown in Fig. 2. This can be efficiently realised using OT extension techniques with an amortized cost of κ bits per random OT, where κ is a computational security parameter [13]. We use the notation $\mathcal{F}_{\text{ROT}}^\tau$ to denote τ parallel copies of \mathcal{F}_{ROT} functionalities.

We also use a coin tossing functionality, which samples an element from a set \mathcal{R} uniformly at random. This can be implemented in the random oracle model by having each party P_i first commit to a random seed s_i with $H(i||s_i)$, then opening all commitments and using $\bigoplus_i s_i$ as a seed to sample from \mathcal{R} .

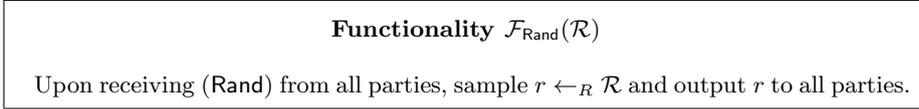


Fig. 3. Coin-tossing functionality

3 Information-Theoretic MAC Scheme

In this section we introduce our secret-shared, information-theoretic message authentication scheme. This forms the backbone of our MPC protocol over \mathbb{Z}_{2^k} . The scheme has two parameters, k , where 2^k is the size of the ring in which computations are performed, and a security parameter s . In the MAC scheme itself and the online phase of our MPC protocol there is no restriction on k , whilst in the preprocessing phase k also affects security.

There is a single, global key $\alpha = \sum_i \alpha^i \bmod 2^{k+s}$, where each party holds a random additive share $\alpha^i \in \mathbb{Z}_{2^s}$. For every authenticated, secret value $x \in \mathbb{Z}_{2^k}$, the parties will have additive shares on this value over the *larger ring* modulo 2^{k+s} , namely shares $x_i \in \mathbb{Z}_{2^{k+s}}$ such that $x' = \sum_i x_i \bmod 2^{k+s}$ and $x \equiv_k x'$. The parties will also have additive shares modulo 2^{k+s} of the MAC $m = \alpha \cdot x' \bmod 2^{k+s}$. We will denote this representation by $[x]$, so we have:

$$[x] = (x^i, m^i, \alpha^i)_{i=1}^n \in (\mathbb{Z}_{2^{k+s}} \times \mathbb{Z}_{2^{k+s}} \times \mathbb{Z}_{2^s})^n, \quad \sum_i m^i \equiv_{k+s} \left(\sum_i x^i \right) \cdot \left(\sum_i \alpha^i \right)$$

Notice that if the parties have $[x]$ and $[y]$, then it is straightforward to obtain by means of local operations $[x + y]$, $[c \cdot x]$ and $[x + c]$, where the arithmetic is modulo 2^{k+s} and c is a constant. For instance, additions and multiplication by a constant is obtained by adding each share (both the value and MAC shares) locally and multiplying by the constant, respectively. Addition by a constant can be achieved similarly. We state the procedures that allow the parties to do this in Fig. 4.

In Fig. 5 we define the functionality \mathcal{F}_{MAC} , which acts as a trusted dealer who samples and distributes shares of the MAC key, and creates secret-shared MACs of additively shared values input by the parties. As with previous works, it allows corrupt parties to choose their own shares instead of sampling them at random, since our protocols allow the adversary to influence the distribution of these. We will show how to implement this functionality in Section 5.

Procedure AffineComb

This procedure allows the parties to compute authenticated shares of $y = c + c_1 \cdot x_1 + \dots + c_t \cdot x_t \pmod{2^{k+s}}$ given $c, c_1, \dots, c_t, [x_1], \dots, [x_t]$. The input to this procedure are the constants $c, c_1, \dots, c_t \in \mathbb{Z}_{2^{k+s}}$, the shares of the values $\{x_i^j\}_{i=1}^t$, the shares of the MACs $\{m_i^j\}_{i=1}^t$, owned by each party P_j , and the shares of the MAC key $\{\alpha^j\}_j$.

1. Party P_1 sets $y^1 = c + c_1 \cdot x_1^1 + \dots + c_t \cdot x_t^1 \pmod{2^{k+s}}$;
2. Each party $P_j, j \neq 1$, sets $y^j = c_1 \cdot x_1^j + \dots + c_t \cdot x_t^j \pmod{2^{k+s}}$;
3. Each party P_j sets $m^j = \alpha_j \cdot c + c_1 \cdot m_1^j + \dots + c_t \cdot m_t^j \pmod{2^{k+s}}$.

At the end of the procedure $\{y^j\}_j$ are additive shares of y modulo 2^{k+s} and $\{m^j\}_j$ are shares of $\alpha \cdot y \pmod{2^{k+s}}$, the MAC of y . To simplify the exposition, we write

$$[c + c_1 \cdot x_1 + \dots + c_t \cdot x_t] = c + c_1 \cdot [x_1] + \dots + c_t \cdot [x_t]$$

whenever this procedure is called.

Fig. 4. Procedure for obtaining authenticated shares of affine combinations of shared values

Note that after running the **Authentication** command with input the shares (x_1^j, \dots, x_t^j) from P_j , the parties will obtain authenticated sharings $[x_1], \dots, [x_t]$, where $x_i = \sum_j x_i^j \pmod{2^k}$.

3.1 Opening Values and Checking MACs

Given an authenticated sharing $[x]$, a natural (but insufficient) approach to opening and reconstructing x is for each party to first broadcast the share x^i and then compute $x' = \sum_i x^i \pmod{2^{k+s}}$. The parties can then check the MAC relation $x' \cdot \alpha$ without revealing the key α using the method from [7]. Although this method guarantees *integrity* of the opened result modulo 2^k (by the same argument sketched in the introduction), it does not suffice for *privacy* when accounting for the fact that x may be a result of applying linear combinations on other private inputs. For example, suppose $x = y + z$ for some previous inputs y, z . When opening x modulo 2^{k+s} , although for correctness we only care about the lower k bits of x , to verify the MAC relation we have to reveal the *entire* shares modulo 2^{k+s} . This leaks whether or not the sum $y + z$ overflowed modulo 2^k .

To prevent this leakage we use an authenticated, random s -bit mask to hide the upper s bits of x when opening. The complete protocol for doing this is shown below.

Procedure SingleCheck($[x]$):

1. Generate a random, shared value $[r]$ using \mathcal{F}_{MAC} , where $r \in \mathbb{Z}_{2^s}$
2. Compute $[y] = [x + 2^k r]$

Functionality \mathcal{F}_{MAC}

The functionality generates shares of a global MAC key and, on input shares of a value, distributes shares of a tag of this value. Let A be the set of corrupted parties and s be a security parameter.

Initialize: On receiving (Init) from all parties, sample random values $\alpha^j \leftarrow_R \mathbb{Z}_{2^s}$ for $j \notin A$ and receive shares $\alpha^j \in \mathbb{Z}_{2^s}$, for $j \in A$, from the adversary. Store the MAC key $\alpha = \sum_{j=1}^n \alpha^j$ (over \mathbb{Z}) and output α^j to party P_j .

Macro Auth(ℓ, x^1, \dots, x^n) (this is an internal subroutine only)

1. Let $x = \sum_{j=1}^n x^j \bmod 2^\ell$ and $m = \alpha \cdot x \bmod 2^\ell$
2. Wait for input $\{m^j\}_{j \in A}$ from the adversary and sample $\{m^j\}_{j \notin A}$ at random conditioned on $m \equiv_\ell \sum_{j=1}^n m^j$. Output (m^1, \dots, m^n) .

Authentication: On input (MAC, $\ell, r, \{x_i^j\}_{i=1}^t$) from each party P_j , where $x_i^j \in \mathbb{Z}_{2^r}$ and $\ell \geq r$:

1. Wait for the adversary to send messages (guess, j, S_j), for every $j \notin A$, where S_j efficiently describes a subset of $\{0, 1\}^s$. If $\alpha^j \in S_j$ for all j then send (success) to \mathcal{A} . Otherwise, send \perp to all parties and abort.
2. Execute **Auth**($\ell, x_i^1, \dots, x_i^n$) for $i = 1, \dots, t$, and then wait for the adversary to send either OK or Abort. If the adversary sends OK then send the MAC shares $m_i^j \in \mathbb{Z}_{2^\ell}$ to party P_j , otherwise abort.

Fig. 5. Functionality for generating shares of global MAC key, distributing shares of inputs and tags

3. Each party broadcasts their shares y^i and reconstructs $y = \sum_i y^i \bmod 2^{k+s}$
4. P_i commits to $z^i = m^i - y \cdot \alpha^i \bmod 2^{k+s}$, where m^i is the MAC share on y
5. All parties open their commitments and check that $\sum_i z^i \equiv_{k+s} 0$
6. If the check passes then output $y \bmod 2^k$

Claim 1 *If the MAC check passes then $y \equiv_k x$, except with probability at most 2^{-s} .*

Proof. Suppose a corrupted party opens $[y]$ to some $y' = y + \delta$, where $\delta \in \mathbb{Z}_{2^{k+s}}$ can be chosen by \mathcal{A} , and $\delta \not\equiv_k 0$. To pass the MAC check, they must also come up with an additive error Δ in the committed values z^i such that $\sum_i z^i + \Delta$ is zero modulo 2^{k+s} . This simplifies to finding $\Delta \in \mathbb{Z}_{2^{k+s}}$ such that

$$\begin{aligned} \sum_i (m^i - (x + \delta) \cdot \alpha^i) + \Delta &\equiv_{k+s} 0 \\ \Leftrightarrow \delta \cdot \alpha &\equiv_{k+s} -\Delta \end{aligned}$$

Let v be the largest integer such that 2^v divides δ , and note that because $\delta \not\equiv_k 0$ we have $v < k$. This means that we can divide the above by 2^v , reducing the modulus from 2^{k+s} to 2^{k+s-v} accordingly:

$$\frac{\delta}{2^v} \cdot \alpha \equiv_{k+s-v} -\frac{\Delta}{2^v}$$

By definition of v , $\frac{\delta}{2^v}$ must be an odd integer, hence invertible modulo 2^{k+s-v} . Multiply by its inverse gives

$$\alpha \equiv_{k+s-v} \frac{\Delta}{2^v} \cdot \left(\frac{\delta}{2^v} \right)^{-1}$$

Note that $k + s - v > s$, since $v < k$, which implies that \mathcal{A} must have guessed $\alpha \bmod 2^s$ to come up with δ and Δ which pass the check. This requires guessing the s least significant bits of α , which are uniformly random, so the probability of success is at most 2^{-s} . \square

3.2 Batch MAC Checking with Random Linear Combinations

Procedure BatchCheck

Procedure for opening and checking the MACs on t shared values $[x_1], \dots, [x_t]$. Let x_i^j, m_i^j, α^j be P_j 's share, MAC share and MAC key share for $[x_i]$.

Open phase: To open a value $[x_i]$:

1. Each party P_j samples $r_i^j \leftarrow_R \mathbb{Z}_{2^s}$, and then calls \mathcal{F}_{MAC} on input $(s, s, r_i^j, \text{MAC})$ to obtain $[r_i]$.
2. Let $[\tilde{x}_i] = [x_i] + 2^k[r_i]$. Denote P_j 's share and MAC share on \tilde{x}_i by \tilde{x}_i^j and m_i^j .
3. Each party P_j broadcasts $\tilde{x}_i^j \in \mathbb{Z}_{2^{k+s}}$.
4. The parties compute $\tilde{x}_i = \sum_{j=1}^n \tilde{x}_i^j \bmod 2^{k+s}$.

MAC check phase:

5. The parties call $\mathcal{F}_{\text{Rand}}(\mathbb{Z}_{2^s}^t)$ to sample public random values $\chi_1, \dots, \chi_t \in \mathbb{Z}_{2^s}$ and then compute $\tilde{y} = \sum_{i=1}^t \chi_i \cdot \tilde{x}_i \bmod 2^{k+s}$.
6. Each party P_j computes $m^j = \sum_{i=1}^t \chi_i \cdot m_i^j \bmod 2^{k+s}$ and $z^j = m^j - \alpha^j \cdot \tilde{y}$. Then it commits to z^j , and then all parties open their commitments.
7. Finally, the parties verify that $\sum_{j=1}^n z^j \equiv_{k+s} 0$. If the check passes then the parties accept the values $x_i := \tilde{x}_i \bmod 2^k$, otherwise they abort.

Fig. 6. Procedure for checking a batch of MACs

The method described in the previous section allows the parties to open and then check one shared value $[x]$. However, in many settings it is useful to be able to check a large batch of values at the end of the protocol, rather than individually as they are opened. To do this, we present a batch MAC checking procedure for opening and checking t shared values $[x_1], \dots, [x_t]$, by checking a single linear combination of the MACs. By postponing the check to the end, this reduces both communication and round complexity compared with repeatedly running `SingleCheck`.

Technically speaking, our main contribution here is a new analysis of the distribution of random linear combinations of adversarially chosen errors modulo

2^k , when lifting these combinations to the larger ring $\mathbb{Z}_{2^{k+s}}$. If we naively apply the analysis from Claim 1 to this case, then we would have to lift to an even bigger ring $\mathbb{Z}_{2^{k+2s}}$ to prove security, adding extra overhead when creating and storing the MACs. With our more careful analysis in Lemma 1 below, we can still work over $\mathbb{Z}_{2^{k+s}}$ and obtain failure probability around $2^{-s+\log s}$, which gives a significant saving.

The procedure that achieves this is described in Fig. 6, and its guarantees are stated in the following theorem.

Theorem 1. *Suppose that the inputs $[x_1], \dots, [x_t]$ to the BatchCheck procedure are consistent sharings of x_1, \dots, x_t under the MAC key $\alpha = \sum_i \alpha^i \pmod{2^s}$, and the honest parties' shares $\alpha^j \in \mathbb{Z}_{2^s}$ are uniformly random in the view of an adversary corrupting at most $n-1$ parties. Then, if the procedure does not abort, the values \tilde{x}_i accepted by the parties satisfy $x_i \equiv_k \tilde{x}_i$ with probability at least $1 - 2^{-s+\log(s+1)}$.*

The following lemma will be used in the proof of this theorem. The lemma is very general, which will allow us to use it also when we prove the security of the preprocessing phase of our protocol. However, in the current context, this lemma will be used with $\ell = k+s$, $r = k$ and $m = s$, and the δ 's can be thought of as the errors introduced by the adversary during the opening phases.

Lemma 1. *Let ℓ, r and m be positive integers such that $\ell - r \leq m$. Let $\delta_0, \delta_1, \dots, \delta_t \in \mathbb{Z}$, and suppose that not all the δ_i 's are zero modulo 2^r , for $i > 0$. Let Y be a probability distribution on \mathbb{Z} . Then, if the distribution Y is independent from the uniform distribution sampling α below, we have*

$$\Pr_{\substack{\alpha, \chi_1, \dots, \chi_t \leftarrow_R \mathbb{Z}_{2^m}, \\ y \leftarrow_R Y}} \left[\alpha \cdot \left(\delta_0 + \sum_{i=1}^t \chi_i \cdot \delta_i \right) \equiv_{\ell} y \right] \leq 2^{-\ell+r+\log(\ell-r+1)},$$

Proof. Define $S := \delta_0 + \sum_{i=1}^t \chi_i \cdot \delta_i$, and define E to be the event that $\alpha \cdot S \equiv_{\ell} y$. Let W be the random variable defined as $\min(\ell, e)$, where 2^e is the largest power of two dividing S . We will use the following claims.

Proposition 1.

- i. $\Pr[E \mid W = r+c] \leq 2^{-(\ell-r-c)}$ for any $c \in \{1, \dots, \ell-r\}$
- ii. $\Pr[E \mid 0 \leq W \leq r] \leq 2^{-(\ell-r)}$
- iii. $\Pr[W = r+c] \leq 2^{-c}$ for any $c \in \{1, \dots, \ell-r\}$

Proof. For the first part, suppose that $0 < c < \ell - r$ (the case $c = \ell - r$ is trivial), in particular, $w = r + c$ is the largest exponent such that 2^w divides S and therefore $S/2^w$ is an odd integer. From the definitions of E and w we have that E holds if and only if $\alpha \cdot S \equiv_{\ell} y$, which in turn is equivalent to $\alpha \cdot \frac{S}{2^w} \equiv_{\ell-2^w} \frac{y}{2^w}$ and therefore to $\alpha \equiv_{\ell-2^w} \frac{y}{2^w} \cdot \left(\frac{S}{2^w}\right)^{-1}$. Since α is uniformly random in \mathbb{Z}_{2^m} and independent of the right-hand side, and also $\ell - w < m$

(as $r < w$ and $\ell - r \leq m$), we conclude that the event holds with probability $2^{-(\ell-w)} = 2^{-(\ell-r-c)}$, conditioned on $W = r + c$.

Similarly, if $0 \leq w \leq r$ then $\ell - w \geq \ell - r$ and so $\alpha \equiv_{\ell-r} \frac{y}{2^w} \cdot \left(\frac{S}{2^w}\right)^{-1}$. As $\ell - r \leq m$, the event holds with probability at most $2^{-(\ell-r)}$ if conditioned on $0 \leq W \leq r$. This proves the second part.

For the third part, we must also look at the randomness from the χ_i coefficients. Suppose without loss of generality that δ_t is non-zero modulo 2^r , and suppose that $W = r + c$ some $1 \leq c \leq \ell - r$. Since $2^W | S$, we have $S \equiv_{r+c} 0$, and so

$$\chi_t \cdot \delta_t \equiv_{r+c} -\delta_0 - \underbrace{\sum_{i \neq t} \chi_i \cdot \delta_i}_{=S'}$$

Let 2^v be the largest power of two dividing δ_t , and note that by assumption we have $v < r$ so $r + c - v > c$. Therefore,

$$\begin{aligned} \chi_t \cdot \frac{\delta_t}{2^v} &\equiv_{r+c-v} \frac{S'}{2^v} \\ \chi_t &\equiv_{r+c-v} \frac{S'}{2^v} \left(\frac{\delta_t}{2^v}\right)^{-1} \\ \chi_t &\equiv_c \frac{S'}{2^v} \left(\frac{\delta_t}{2^v}\right)^{-1} \end{aligned}$$

By the same argument as previously, and from the fact that $c \leq \ell - r \leq m$, this holds with probability 2^{-c} , over the randomness of $\chi_t \leftarrow_R \mathbb{Z}_{2^m}$, as required.

Putting things together, we apply the law of total probability over all possible values of w , obtaining:

$$\begin{aligned} \Pr[E] &= \Pr[E \mid 0 \leq W \leq r] \cdot \Pr[0 \leq W \leq r] + \sum_{c=1}^{\ell-r} \Pr[E \mid W = r+c] \cdot \Pr[W = r+c] \\ &\leq 2^{-\ell+r} \cdot 1 + \sum_{c=1}^{\ell-r} 2^{-\ell+r+c} \cdot 2^{-c} = 2^{-\ell+r} + \sum_{c=1}^{\ell-r} 2^{-\ell+r} \\ &= (\ell - r + 1) \cdot 2^{-\ell+r} \leq 2^{-\ell+r+\log(\ell-r+1)} \end{aligned}$$

where the first inequality comes from applying item **ii.** of Proposition 1 on the left, and items **i.** and **iii.** on the right. \square

Now we proceed with the proof of Theorem 1.

Proof (of Theorem 1). We first assume that \mathcal{A} sends no **Key Query** messages to \mathcal{F}_{MAC} , and later discuss how the claim still holds when this is not the case.

First of all notice that if no error is introduced by the adversary then the check passes, since we have

$$\sum_{j=1}^n z^j \equiv_{k+s} \sum_{j=1}^n m^j - \tilde{y} \cdot \sum_{j=1}^n \alpha^j \equiv_{k+s} 0$$

Now, consider the case in which the adversary does not open correctly to \tilde{x}_i in the protocol. Let $\tilde{x}_i + \delta_i \pmod{2^{k+s}}$ be the value opened in step 3, so the value computed in step 5 is equal to $\tilde{y}' = \tilde{y} + \delta \pmod{2^{k+s}}$, where $\delta = \sum_{i=1}^t \chi_i \cdot \delta_i \pmod{2^{k+s}}$. As a consequence, the share that an honest P_j should open in step 6 is $z^j - \alpha^j \cdot \delta \pmod{2^{k+s}}$. However, the adversary can open this value plus some errors that sum up to a value $\Delta \in \mathbb{Z}_{2^{k+s}}$. If the check passes, this means that

$$0 \equiv_{k+s} \sum_{j=1}^n (z^j - \alpha^j \cdot \delta) + \Delta \Leftrightarrow \alpha \cdot \delta \equiv_{k+s} \Delta.$$

Suppose that for some index it holds that $\delta_i \not\equiv_k 0$. By setting $\delta_0 = 0$, $\ell = k+s$, $r = k$, $m = s$ and Y to be the distribution of Δ produced by the adversary, we observe we are in the same setting as the hypothesis of Lemma 1. This allows us to conclude that the probability that the check passes is bounded by $2^{-\ell+r+\log(\ell-r+1)} = 2^{-s+\log(s+1)}$.

Handling key queries. We now show that this probability is the same for an adversary who makes some successful queries to an honest party's α^j using the (guess) command of \mathcal{F}_{MAC} . Let S be the set of possible keys guessed by \mathcal{A} (if there is more than one query then we take S to be the intersection of all sets). The probability that all these queries are successful is no more than $|S|/2^s$, and conditioned on this event, the min-entropy of the honest party's key share is reduced to $\log|S| \leq s$. Therefore, instead of success probability $2^{-s+\log(s+1)}$ as above, the overall probability of \mathcal{A} performing successful key queries *and* passing the check is bounded by

$$|S|/2^s \cdot 2^{-\log|S|+\log(\log|S|+1)} = 2^{-s+\log(\log|S|+1)} \leq 2^{-s+\log(s+1)}$$

as required. \square

3.3 Reducing the Number of Masks

One drawback of our MAC scheme is that with both the single and batch checking protocols, each opening requires an extra s -bit authenticated random value, to mask the upper s bits. This was not needed in the MAC scheme for SPDZ over fields. We observe, though, that throughout most of the online phase in our actual MPC protocol, this random mask is unnecessary. In our multiplication protocol for the online phase (Section 4), the two values being opened are masked by random multiplication triples. Since our triple generation functionality (Fig. 7) guarantees that the a and b components of triples are random modulo

2^{k+s} , they already serve to hide the upper s bits of the values being masked. This means we can omit the masks used in the **Open** phase of our batch MAC check procedure, for every multiplication gate in the circuit. This reduces the amount of preprocessing material needed, so leads to less communication in the preprocessing phase.

3.4 Previous Version of the Batch MAC Check

A previous version of the batch MAC checking protocol⁵ contained an additional optimization, where the masking was postponed until the **MAC check** phase of the protocol. This optimization was not secure, as we now explain. The idea was that, to initially open each value x_i , the parties can open their shares of x_i reduced modulo 2^k and reconstruct this, again modulo 2^k , avoiding any leakage on the upper s bits. When checking the MACs, the parties then opened the random linear combination applied to the upper s bits of the x_i 's, and used this to obtain the missing bits needed to carry out the final MAC check. This meant that only one s -bit mask was needed, namely, when opening this last linear combination, and also reduced the online communication (per party) from $O(k+s)$ to $O(k)$ bits per opening.

However, this check allows a corrupt party to cheat with probability $1/2$. If the adversary introduces some error ϵ into the final s -bit correction value, it can be seen that this translates into an error of

$$\alpha \cdot (2^k \cdot \epsilon + \sum_{i=1}^t \chi_i \cdot \delta_i)$$

which the adversary must try to guess (here, the δ_i 's are the original k -bit errors in the openings of each x_i). If, say, $\delta_1 = 2^{k-1}$ and all other δ_j are zero, and if χ_1 is even (which happens with probability $1/2$), then the adversary can choose $\epsilon = \chi_1/2$ to make the above go to zero. The main issue here, which was not considered in the previous version of this work, is that ϵ can be chosen by the adversary *after* the χ_i 's. In particular, it is not possible to apply Lemma 1 with $\delta_0 = 2^k \cdot \epsilon$, since in order to do that δ_0 should be fixed before the χ_i 's.

4 Online Phase

Our protocol is divided in two phases, a preprocessing phase and an online phase. The preprocessing, which is independent of each party's input, implements a functionality $\mathcal{F}_{\text{Prep}}$ which generates the necessary shared, authenticated values needed to compute the given function securely. This functionality is stated in Fig. 7.

The main difference, with respect to SPDZ, is that instead of generating the random input masks and multiplication triples over the same space as the inputs, we sample them over $\mathbb{Z}_{2^{k+s}}$, even though we are doing computations in \mathbb{Z}_{2^k} . In

⁵ <https://eprint.iacr.org/2018/482/20180523:024407>

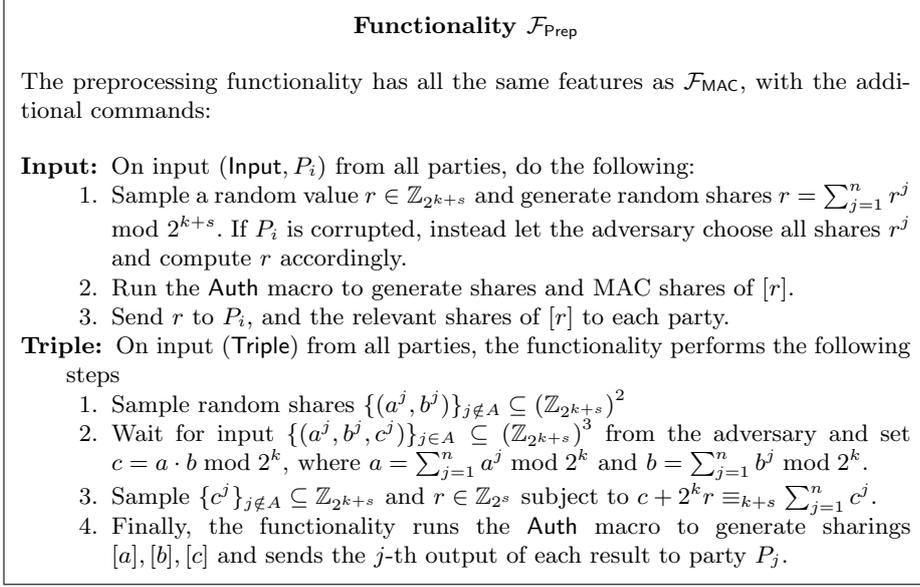


Fig. 7. Functionality for the preprocessing phase

the input phase, this is necessary to mask the parties' input whilst also obtaining a correct MAC over $\mathbb{Z}_{2^{k+s}}$. For the triples, we sample the shares and compute the MACs in $\mathbb{Z}_{2^{k+s}}$, but only care about *correctness* of the multiplication modulo 2^k , so the upper s bits of a triple are just random.

Modulo these differences, the online phase of our protocol, shown in Fig. 9, is similar to that in other secret sharing-based protocols like GMW, BeDOZa, SPDZ and MASCOT [11,2,9,14]. The overall idea is that the parties will have authenticated shares of the values of all wires in the circuit, beginning with the input wires until they get shares of the output values, which are then opened to reconstruct the result of the computation. If desired, the parties can then continue computing on shares and outputting more results, allowing for arbitrary *reactive* computations.

Shares of the inputs are distributed by means of the random shares provided by $\mathcal{F}_{\text{Prep}}$. When an addition gate is found, the parties obtain the output by adding their shares locally. On the other hand, multiplication triples are used for the multiplication gates, where the fact that $x \cdot y = c + \epsilon \cdot b + \delta \cdot a + \epsilon \cdot \delta$ for $c = a \cdot b$, $\epsilon = x - a$ and $\delta = y - b$ allows us to evaluate multiplications as affine operations on x and y , once the values of ϵ and δ are known. Finally, after checking correctness of all the values opened in multiplications using the batch MAC checking procedure from section 3, the values for the output wires are revealed.

The proof of the following theorem is quite straightforward, given the analysis of the MACs in Section 3, so we present it in Appendix A.

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

Initialization: The functionality receives input (Init, k) from all parties.

Input: On input $(\text{Input}, P_i, \text{vid}, x)$ from party P_i and input (Input, P_i) from the other parties, where vid is a fresh, valid identifier, the functionality stores $(\text{vid}, x \bmod 2^k)$.

Add: On input $(\text{add}, \text{vid}_1, \text{vid}_2, \text{vid}_3)$ from all parties, the functionality retrieves (if present in memory) the values (vid_1, x_1) , (vid_2, x_2) and stores $(\text{vid}_3, x_1 + x_2 \bmod 2^k)$.

Multiply: On input $(\text{multiply}, \text{vid}_1, \text{vid}_2, \text{vid}_3)$ from all parties, the functionality retrieves (if present in memory) the values (vid_1, x_1) , (vid_2, x_2) and stores $(\text{vid}_3, x_1 \cdot x_2 \bmod 2^k)$.

Output: On input $(\text{output}, \text{vid})$ from all honest parties, the functionality looks for (vid, y) in memory and if present, sends y to the adversary. The functionality then waits for a message **Abort** or **Proceed** from the adversary: if it sends **Abort** then the functionality aborts, otherwise the value y is delivered to all parties.

Fig. 8. Ideal functionality for the online phase

Protocol Π_{Online}

The protocol is parameterized by k , which specifies the word size on which the operations are to be performed, and a security parameter s .

Initialize: The parties call the functionality $\mathcal{F}_{\text{Prep}}$ as follows:

1. On input (Init) to get MAC key shares $\alpha^j \in \mathbb{Z}_{2^s}$.
2. On input (Input, P_i) for all parties to obtain random sharings $[r]$ where P_i learns r , for every input that P_i will provide.
3. On input (Triple) to get enough triples $([a], [b], [c])$.

Input: To share an input x^i held by P_i :

1. P_i broadcasts $\epsilon = x^i - r \bmod 2^{k+s}$, where $[r]$ is the next unused input mask.
2. The parties compute $[x^i] = [r] + \epsilon$.

Add: To add two values $[x]$ and $[y]$ the parties compute locally $[z] = [x] + [y]$.

Multiply: To multiply two values $[x]$ and $[y]$:

1. Open $[x] - [a]$ as ϵ and $[y] - [b]$ as δ using the **Open** phase of **BatchCheck**, where $([a], [b], [c])$ is the next unused triple.
2. Locally compute $[x \cdot y] = [c] + \epsilon \cdot [b] + \delta \cdot [a] + \epsilon \cdot \delta$.

Output: To output a value $[y]$:

1. Call the procedure **BatchCheck** to check the MACs on the values that have been opened so far in multiplications.
2. If this does not abort, the parties open and check the MAC on $[y]$ using the procedure **SingleCheck** from Section 3.1.

Fig. 9. Protocol for reactive secure multi-party computation over \mathbb{Z}_{2^k}

Theorem 2. *The protocol Π_{Online} implements $\mathcal{F}_{\text{Online}}$ in the $\mathcal{F}_{\text{Prep}}$ -hybrid model, with statistical security parameter s .*

5 Preprocessing: Creating the MACs

We now show how to authenticate additively shared values with the linear MAC scheme, realising the functionality \mathcal{F}_{MAC} from Section 3 (Fig. 5). Recall that after sampling shares of the MAC key $\alpha \in \mathbb{Z}_{2^s}$, the functionality takes as input secret-shared values $x \in \mathbb{Z}_{2^r}$, and produces shares of the MAC $x \cdot \alpha \bmod 2^\ell$. The input and output widths r and ℓ are parameters with $\ell \geq r$. In our protocol we actually require $\ell \geq 2s$ and $\ell \geq r + s$, where s is the security parameter, but if these do not hold then we work with $\ell' = \max(r + s, 2s)$ and reduce the outputs modulo 2^ℓ .

Building block: vector oblivious linear function evaluation. To create the MACs, we will use a functionality for random vector oblivious linear function evaluation (vector-OLE) over the integers modulo 2^ℓ . This is a protocol between two parties, P_A and P_B , that takes as input a fixed element $\alpha \in \mathbb{Z}_{2^s}$ from party P_A , a vector \mathbf{x} from party P_B , then samples a random vector $\mathbf{b} \in \mathbb{Z}_{2^\ell}^t$ as output to P_B , and sends $\mathbf{a} = \mathbf{b} + \alpha \cdot \mathbf{x} \bmod 2^\ell$ to P_A . In the specification of our ideal functionality in Fig. 10, \mathbf{x} is a vector of length $t + 1$, with the first t components from \mathbb{Z}_{2^r} and the final component from \mathbb{Z}_{2^ℓ} . This is because our MAC generation protocol will create a batch of t MACs at once on r -bit elements, but to do this securely we also need to authenticate an additional random mask of ℓ bits.

Notice that the functionality also allows a corrupted P_B to try to guess a subset of \mathbb{Z}_{2^s} in which α lies, but if the guess is incorrect the protocol aborts. This is needed in order to efficiently implement $\mathcal{F}_{\text{VOLE}}$ using oblivious transfer on correlated messages, based on existing oblivious transfer extension techniques, which we describe in Section C.

| Functionality $\mathcal{F}_{\text{VOLE}}^s$ |
|--|
| <p>Initialize: On receiving $(sid, \text{Init}, \alpha)$ from P_A, where $\alpha \in \mathbb{Z}_{2^s}$, and (sid, Init) from P_B, store α and ignore any subsequent (sid, Init) messages.</p> <p>Vector-OLE: On input $(sid, \ell, r, t, \mathbf{x})$ from P_B, where $\mathbf{x} \in \mathbb{Z}_{2^r}^t \times \mathbb{Z}_{2^\ell}$:</p> <ol style="list-style-type: none"> 1. Sample $\mathbf{b} \leftarrow_R \mathbb{Z}_{2^\ell}^{t+1}$. If P_B is corrupted, instead receive \mathbf{b} from \mathcal{A}. 2. Compute $\mathbf{a} = \mathbf{b} + \alpha \cdot \mathbf{x} \bmod 2^\ell$ 3. If P_A is corrupted, receive $\mathbf{a} \in \mathbb{Z}_{2^\ell}^t$ from \mathcal{A} and recompute $\mathbf{b} = \mathbf{a} - \alpha \cdot \mathbf{x}$. 4. If P_B is corrupted, wait for \mathcal{A} to input a message (guess, S), where S efficiently describes a subset of $\{0, 1\}^s$. If $\alpha \in S$ then send (success) to \mathcal{A}. Otherwise, send \perp to both parties and terminate. 5. Output \mathbf{a} to P_A and \mathbf{b} to P_B. |

Fig. 10. Random vector oblivious linear function evaluation functionality over $\mathbb{Z}_{2^{k+s}}$

MAC generation protocol. Each party samples a random MAC key share α^i , and uses this to initialize an instance of $\mathcal{F}_{\text{VOLE}}$ with every other party. On input

a vector of additive secret shares $\mathbf{x}^i = (x_1^i, \dots, x_t^i)$ from every P_i , each party samples a random ℓ' -bit mask x_{t+1}^i , and then uses \mathcal{F}_{OLE} to compute two-party secret-sharings of the products $\alpha^i \cdot (\mathbf{x}^j \| x_{t+1}^j)$ for all $j \neq i$. Each party can then obtain a share of the MACs $\alpha \cdot \mathbf{x}$ (where $\alpha = \sum \alpha^i$ and $\mathbf{x} = \sum \mathbf{x}^i$), by adding up all the two-party sharings together with the product $\alpha^i \cdot \mathbf{x}^i$.

So far, the protocol is only passively secure, since there is nothing to prevent a corrupt P_j from using inconsistent values of α^j or \mathbf{x}^j with two different honest parties, so the corrupt parties' inputs may not be well-defined. To prevent this issue, and ensure that in the security proof the simulator can correctly extract the adversary's inputs, we add a consistency check in steps 6–11: this challenges the parties to open a random linear combination of all authenticated values. This is where we need the additional random mask x_{t+1} , to prevent any leakage on the parties inputs from opening this linear combination. The check does not rule out *all* possible deviations in the protocol, however, in what follows we show that it ensures that the *sum* of all the errors directed towards any given honest party is zero, so these errors all cancel out. Intuitively, this suffices to realise \mathcal{F}_{MAC} because the functionality only adds a MAC to the sum of all parties' inputs, and not the individual shares themselves.

Protocol Π_{Auth}

Initialize: Each party P_i samples a MAC key share $\alpha^i \leftarrow_R \mathbb{Z}_{2^s}$. Every pair of parties (P_i, P_j) initializes an instance of \mathcal{F}_{OLE} , where P_i inputs α_i .

Authentication: To authenticate the values $\mathbf{x} = (x_1, \dots, x_t)$ over \mathbb{Z}_{2^ℓ} , where each party P_j inputs an additive share $\mathbf{x}^j \in \mathbb{Z}_{2^r}^t$:

1. Let $\ell' = \max(\ell, r + s, 2s)$.
2. Each party P_j samples a random mask $x_{t+1}^j \leftarrow_R \mathbb{Z}_{2^{\ell'}}$ and defines $\tilde{\mathbf{x}}^j := (\mathbf{x}^j, x_{t+1}^j) \in \mathbb{Z}_{2^r}^t \times \mathbb{Z}_{2^{\ell'}}$.
3. Every pair (P_i, P_j) (for $i \neq j$) calls their \mathcal{F}_{OLE} instance with input $(\ell', r, t, \tilde{\mathbf{x}}^j)$ from P_j .
4. P_j receives $\mathbf{b}^{j,i}$ and P_i receives $\mathbf{a}^{i,j}$, such that $\mathbf{a}^{i,j} = \mathbf{b}^{j,i} + \alpha^i \cdot \tilde{\mathbf{x}}^j \pmod{2^{\ell'}}$.
5. For $h = 1, \dots, t + 1$, each party P_j defines the MAC share

$$m_h^j = \alpha^j \cdot x_h^j + \sum_{i \neq j} (\mathbf{a}^{j,i} - \mathbf{b}^{j,i})[h] \pmod{2^{\ell'}}$$

Consistency check:

6. Sample $\chi_1, \dots, \chi_t \leftarrow_R \mathbb{Z}_{2^s}^t$ using $\mathcal{F}_{\text{Rand}}$.
7. Each party P_j computes and broadcasts $\hat{x}^j = \sum_{i=1}^t x_i^j \cdot \chi_i + x_{t+1}^j \pmod{2^{\ell'}}$.
8. Each party P_j defines $\hat{m}^j = \sum_{h=1}^t m_h^j \cdot \chi_h + m_{t+1}^j \pmod{2^{\ell'}}$ and $\hat{x} = \sum_i \hat{x}^i$.
9. Each party P_j commits to and then opens $z^j = \hat{m}^j - \hat{x} \cdot \alpha^j \pmod{2^{\ell'}}$.
10. All parties check that $\sum_i z^i = 0 \pmod{2^{\ell'}}$ and abort if the check fails.
11. Each party P_j outputs the MAC shares $m_1^j, \dots, m_t^j \pmod{2^\ell}$.

Fig. 11. Protocol for authenticating secret-shared values

5.1 Security

We now analyse the consistency check of the MAC creation protocol. There are two main types of deviations that a corrupt P_j can perform, namely (1) Input inconsistent values of α^j to the initialization phase of $\mathcal{F}_{\text{VOLUME}}$ with different honest parties, and (2) Input inconsistent shares \mathbf{x}^j in the authentication stage.

Let $\alpha^{j,i}$ and $\mathbf{x}^{j,i}$ be the actual values used by a corrupt P_j in its $\mathcal{F}_{\text{VOLUME}}$ instance with honest P_i . For both types of errors, we define the *correct* values α^j, \mathbf{x}^j to equal $\alpha^{j,i_0}, \mathbf{x}^{j,i_0}$, for some arbitrary, fixed honest party, P_{i_0} . We then define the errors

$$\gamma^{j,i} = \alpha^{j,i} - \alpha^j \quad \text{and} \quad \delta^{j,i} = \mathbf{x}^{j,i} - \mathbf{x}^j,$$

for each $j \in A$ and $i \notin A$. For an honest party P_i , we also define $\alpha^{i,j}, \mathbf{x}^{i,j}$ to be equal to α^i, \mathbf{x}^i for all $j \neq i$. Finally, for convenience, we define the errors between a corrupt P_j and corrupt P_i to be zero, i.e. $\alpha^{j,i} = \alpha^j$ and $\mathbf{x}^{j,i} = \mathbf{x}^j$ for $j, i \in A$.

In Claims 2 and 3 below we will show that, if the consistency check passes, then with overwhelming probability the *sum* of all corrupted parties' values is well-defined. That is, the values $\sum_{j \in A} \alpha^j$ and $\sum_{j \in A} \mathbf{x}^j$ would be exactly same even if they were defined using the inputs from P_j with a *different* honest party $P_{i_1} \neq P_{i_0}$. Since the MACs are computed based only on the sum of the MAC key shares and input shares, this suffices to prove security of the protocol.

Suppose that the corrupted parties compute the MAC shares \mathbf{m}^j as an honest P_j would, using the values α^j, \mathbf{x}^j we defined above, as well as the values $\mathbf{a}^{j,i}, \mathbf{b}^{j,i}$ sent to $\mathcal{F}_{\text{VOLUME}}$. Note that even though a corrupt P_j need not do this, any deviation here can be modelled by an additive error in the commitment to z^j in step 9, so we do not lose any generality.

The sum of the vector of MAC shares on \mathbf{x} is then given by

$$\begin{aligned} \sum_i \mathbf{m}^i &= \sum_i \alpha^i \cdot \mathbf{x}^i + \sum_i \sum_{j \neq i} (\mathbf{a}^{i,j} - \mathbf{b}^{j,i}) \\ &= \sum_i \alpha^i \cdot \mathbf{x}^i + \sum_i \sum_{j \neq i} \alpha^{i,j} \cdot \mathbf{x}^{j,i} \\ &= \alpha \cdot \mathbf{x} + \sum_{i \notin A} \mathbf{x}^i \cdot \underbrace{\sum_{j \in A} \gamma^{j,i}}_{=\gamma^i} + \sum_{i \notin A} \alpha_i \cdot \underbrace{\sum_{j \in A} \delta^{j,i}}_{=\delta^i} \end{aligned}$$

After taking random linear combinations with the vector $\boldsymbol{\chi} = (\chi_1, \dots, \chi_t)$ to compute the MAC on \hat{x} , the final MAC shares satisfy

$$\sum_i \hat{m}^i = \alpha \cdot \hat{x} + \sum_{i \notin A} (\langle \mathbf{x}^i, \boldsymbol{\chi} \rangle + x_{t+1}^i) \cdot \gamma^i + \sum_{i \notin A} \alpha^i \cdot (\langle \boldsymbol{\delta}^i, \boldsymbol{\chi} \rangle + \delta_{t+1}^i) \quad (1)$$

To pass the consistency check, the adversary must first open the random linear combination \hat{x} to some (possibly incorrect) value, say $\hat{x} + \varepsilon$, in step 7. Then they must come up with an error $\Delta \in \mathbb{Z}_{2^{\ell'}}$ such that

$$\begin{aligned}
0 &\equiv_{\ell'} \sum_i z^i + \Delta \\
&\equiv_{\ell'} \sum_i (\hat{m}^i - (\hat{x} + \varepsilon) \cdot \alpha^i) + \Delta \\
\Leftrightarrow -\Delta &\equiv_{\ell'} \sum_i \hat{m}^i - (\hat{x} + \varepsilon) \cdot \alpha \\
&\equiv_{\ell'} \alpha \cdot \varepsilon + \sum_{i \notin A} \underbrace{(\langle \mathbf{x}^i, \boldsymbol{\chi} \rangle + x_{t+1}^i)}_{=u^i} \cdot \gamma^i + \sum_{i \notin A} \alpha^i \cdot \langle \boldsymbol{\delta}^i, \boldsymbol{\chi} \rangle + \delta_{t+1}^i \\
-\Delta - \sum_{j \in A} \alpha^j \cdot \varepsilon &\equiv_{\ell'} \sum_{i \notin A} u^i \cdot \gamma^i + \sum_{i \notin A} \alpha^i \cdot (\langle \boldsymbol{\delta}^i, \boldsymbol{\chi} \rangle + \delta_{t+1}^i + \varepsilon)
\end{aligned}$$

where the last two congruences come from substituting (1) and moving information known by the adversary to the left-hand side.

When proving the two claims below we assume that the adversary does not send any (guess) messages to $\mathcal{F}_{\text{VOLE}}$. Similarly to the proof of Theorem 1, these can easily be extended to handle this case.

Claim 2 *If at least one $\gamma^i \neq 0$ then the probability of passing the check is no more than $2^{-s+\log n}$.*

Proof. Let i be an index for where $\gamma^i \neq 0$. Recall that $\gamma^i = \sum_{j \notin A} \gamma^{j,i}$, where each $\gamma^{j,i} < 2^s$, therefore $\gamma^i < 2^{s+\log n}$. Note that the distribution of u^i is uniform in $\mathbb{Z}_{2^{\ell'}}$ and independent of all other terms, due to the extra mask x_{t+1}^i , so we can write $u^i \cdot \gamma^i \equiv_{\ell'} \Delta'$, for some Δ' that is independent of u^i . Dividing by 2^v , the largest power of two dividing γ^i , we get

$$\begin{aligned}
u^i \cdot \frac{\gamma^i}{2^v} &\equiv_{\ell'-v} \frac{\Delta'}{2^v} \\
u^i &\equiv_{\ell'-v} \frac{\Delta'}{2^v} \cdot \left(\frac{\gamma^i}{2^v}\right)^{-1}
\end{aligned}$$

Since $v < s + \log n$, this holds with probability at most $2^{-\ell'+s+\log n} \leq 2^{-s+\log n}$ since $\ell' \geq 2s$.

Claim 3 *Suppose $\gamma^i = 0$ for all $i \notin A$, and $\boldsymbol{\delta}^i$ is non-zero modulo 2^r in at least one component for some $i \notin A$. Then, the probability of passing the check is no more than $2^{-s+\log(s)}$.*

Proof. Since i_0 is an honest index where $\boldsymbol{\delta}^{i_0}$ and $\delta_{t+1}^{i_0}$ are both zero, passing the check implies that

$$-\Delta - \sum_{j \in A} \alpha^j \cdot \varepsilon \equiv_{\ell'} \alpha^{i_0} \cdot \varepsilon + \sum_{i \neq i_0, i \notin A} \alpha^i \cdot (\langle \boldsymbol{\delta}^i, \boldsymbol{\chi} \rangle + \delta_{t+1}^i + \varepsilon) \quad (2)$$

Let 2^v be the largest power of two that divides ε . We can rewrite (2) as

$$\Delta' \equiv_{\ell'} \alpha^{i_0} \cdot \varepsilon,$$

for some Δ' that is independent of α^{i_0} . Dividing both sides by 2^v and multiplying by the inverse of $\varepsilon/2^v$, we get

$$\alpha^{i_0} \equiv_{\ell'-v} \Delta' \cdot (\varepsilon/2^v)^{-1} \quad (3)$$

Since α^{i_0} is uniformly random and independent of the right-hand side, the above only holds with probability $\rho_0 = 2^{-\ell'+v}$.

Clearly, if $v \leq r$ then $\rho_0 \leq 2^{-s}$ and we are done. We now assume that $v > r$. Since 2^v divides ε , we can write equation (2) in a different way, as

$$\Delta'' \equiv_v \alpha^{i_1} \cdot (\langle \delta^i, \chi \rangle + \delta_{t+1}^i) \quad (4)$$

for some Δ'' that is independent of α^{i_1} . Using Lemma 1 with $r = r, m = s, \ell = v$ and $\delta_0 = \delta_{t+1}^i$, the above can only hold with probability $\rho_1 = 2^{-v+r+\log(v-r+1)}$.

Combining Equations (3) and (4), and using the fact that $\alpha^{i_0}, \alpha^{i_1}$ are sampled independently, we have

$$\Pr[\mathcal{A} \text{ passes check}] \leq \rho_0 \cdot \rho_1 = 2^{-\ell'+v-v+r+\log(v-r+1)} \leq 2^{-s+\log(s)}$$

The above two claims show that, except with negligible probability in s , the sum of all errors directed towards any given honest party is zero, so all errors introduced by corrupt parties cancel out and the outputs form a correct MAC on the underlying shared value. In particular, for the security proof, this implies that in the ideal world the MAC shares seen by the environment (including those of honest parties') are identically distributed to the MAC shares output in the real world.

In Appendix D, we give a complete proof of the following.

Theorem 3. *The protocol Π_{Auth} securely realises \mathcal{F}_{MAC} in the $(\mathcal{F}_{\text{Vole}}, \mathcal{F}_{\text{Rand}})$ -hybrid model.*

6 Preprocessing: Creating Multiplication Triples

In this section we focus on developing a protocol that implements the Triple command in the preprocessing functionality. More precisely, let $\mathcal{F}_{\text{Triple}}$ be the functionality that has the same features as $\mathcal{F}_{\text{Prep}}$ (Fig. 7), but without the **Input** command. Our protocol, described in Fig. 12, implements the functionality $\mathcal{F}_{\text{Triple}}$ in the $(\mathcal{F}_{\text{ROT}}, \mathcal{F}_{\text{MAC}}, \mathcal{F}_{\text{Rand}})$ -hybrid model. In Appendix B we also present the complete protocol for implementing $\mathcal{F}_{\text{Prep}}$, which includes the method for generating random masks to authenticate parties' inputs.

The protocol itself is very similar to the one used in MASCOT [14], with several changes introduced in order to cope with the fact that our ring \mathbb{Z}_{2^k} has

Protocol Π_{Triple}

The integer parameter $\tau = 4s + 2k$ specifies the size of the input triple used to generate each output triple.

Multiply:

1. Each party P_i samples $\mathbf{a}^i = (a_1^i, \dots, a_\tau^i) \leftarrow_R (\mathbb{Z}_2)^\tau$, $b^i \leftarrow_R \mathbb{Z}_{2^{k+s}}$
2. Every ordered pair of parties (P_i, P_j) does the following:
 - (a) Both parties call $\mathcal{F}_{\text{ROT}}^\tau$ with P_i as the receiver and P_j as the sender. P_i inputs the bits $(a_1^i, \dots, a_\tau^i) \in (\mathbb{Z}_2)^\tau$.
 - (b) P_j receives $q_{0,h}^{j,i}, q_{1,h}^{j,i} \in \mathbb{Z}_{2^{k+s}}$ and P_i receives $s_h^{i,j} = q_{a_h^i, h}^{j,i}$ for $h = 1, \dots, \tau$.
 - (c) P_j sends $d_h^{j,i} = q_{0,h}^{j,i} - q_{1,h}^{j,i} + b^j \pmod{2^{k+s}}$, for $h = 1, \dots, \tau$.
 - (d) P_i sets $t_h^{i,j} = s_h^{i,j} + a_h^i \cdot d_h^{j,i} \pmod{2^{k+s}}$ for $h = 1, \dots, \tau$. In particular

$$\begin{aligned} t_h^{i,j} &\equiv_{k+s} s_h^{i,j} + a_h^i \cdot d_h^{j,i} \\ &\equiv_{k+s} q_{a_h^i, h}^{j,i} + a_h^i \cdot (q_{0,h}^{j,i} - q_{1,h}^{j,i} + b^j) \\ &\equiv_{k+s} q_{0,h}^{j,i} + a_h^i b^j. \end{aligned}$$

Therefore, the following equation holds modulo 2^{k+s} on each entry

$$\begin{pmatrix} t_1^{i,j} \\ t_2^{i,j} \\ \vdots \\ t_\tau^{i,j} \end{pmatrix} = \begin{pmatrix} q_{0,1}^{j,i} \\ q_{0,2}^{j,i} \\ \vdots \\ q_{0,\tau}^{j,i} \end{pmatrix} + b^j \begin{pmatrix} a_1^i \\ a_2^i \\ \vdots \\ a_\tau^i \end{pmatrix}$$

- (e) P_i sets $\mathbf{c}_{i,j}^i = (t_1^{i,j}, t_2^{i,j}, \dots, t_\tau^{i,j}) \in (\mathbb{Z}_{2^{k+s}})^\tau$.
- (f) P_j sets $\mathbf{c}_{i,j}^j = -(q_{0,1}^{j,i}, q_{0,2}^{j,i}, \dots, q_{0,\tau}^{j,i}) \in (\mathbb{Z}_{2^{k+s}})^\tau$.
- (g) The following congruence holds

$$\mathbf{c}_{i,j}^i + \mathbf{c}_{i,j}^j \equiv_{k+s} \mathbf{a}^i \cdot b^j,$$

where the modulo congruence is component-wise.

3. Each party P_i computes:

$$\mathbf{c}^i = \mathbf{a}^i \cdot b^i + \sum_{j \neq i} (\mathbf{c}_{i,j}^i + \mathbf{c}_{j,i}^i) \pmod{2^{k+s}}$$

Fig. 12. Triple generation protocol

non-invertible elements. Most of these changes involve taking the coefficients of random linear combinations in a different ring \mathbb{Z}_{2^s} , which is useful to argue that certain equations of the form $r \cdot a \equiv_{k+s} b$ are satisfied with low probability. This can be seen for example in the sacrifice step, where the random value t is chosen to have at least s random bits, instead of k . Additionally, in our protocol (like in MASCOT) random linear combinations must be used to extract randomness from partially leaked values a_1, \dots, a_t , which still have reasonably high entropy.

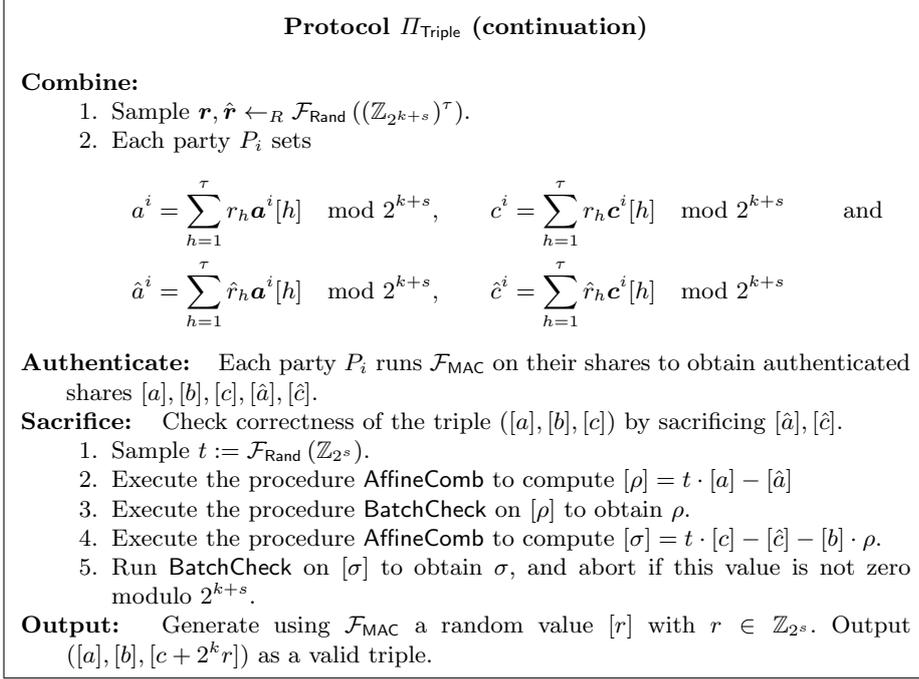


Fig. 13. Triple generation protocol (continuation)

In order to use the Leftover Hash Lemma in this context one needs to make sure that taking random linear combinations yields a universal hash function. However, in contrast to the field case it is not true in general that the function $r_1 \cdot a_1 + \dots + r_t \cdot a_t \pmod{2^k}$ is universal, unless we make some assumptions about the set the values a_i are picked from. In the case of our protocol, we force the a_i to be $-1, 0$ or 1 . With this additional condition it can be shown that the function above is universal.

The **Multiply** phase generates shares $\{(\mathbf{a}^i, b^i, \mathbf{c}^i)\}_{i=1}^n$ such that P_i has $(\mathbf{a}^i, b^i, \mathbf{c}^i)$, where \mathbf{a}^i is a vector of bits, b^i is a random element of $\mathbb{Z}_{2^{k+s}}$ and \mathbf{c}^i is a vector of random elements of $\mathbb{Z}_{2^{k+s}}$. These values satisfy $\mathbf{c} = \mathbf{a} \cdot b$, where $\mathbf{c} = \sum_{i=1}^n \mathbf{c}^i \pmod{2^{k+s}}$, $\mathbf{a} = \sum_{i=1}^n \mathbf{a}^i \pmod{2^{k+s}}$ and $b = \sum_{i=1}^n b^i \pmod{2^{k+s}}$. This is achieved by letting the parties choose their shares on \mathbf{a} and b , and using oblivious transfer to compute the cross products $\mathbf{a}^i \cdot b^j$. However, this is not a fully functional multiplication triple yet as it might not satisfy the right multiplicative relation (besides other technical issues like \mathbf{a} being a short vector, and not a value in $\mathbb{Z}_{2^{k+s}}$). To check that the triple is correct, the **Sacrifice** phase uses another triple to check correctness. As the name suggests, one triple is “sacrificed” (i.e. opened) so that we can check correctness of the other while keeping it secret.

On the other hand, we must also ensure that the triple looks random to all parties. As we will see shortly in the proof of Theorem 4, if the triple is correct this will reveal some partial information about the honest parties’ shares to the

adversary. This means that the adversary can guess a particular bit of these shares, which would allow him to distinguish in the simulation. This issue is addressed by the step **Combine**, which takes place before the Sacrifice step. Here the parties take a random linear combination of \mathbf{a} . Now, in order to pass the check, the adversary has to guess a random combination of the bits of \mathbf{a} , which is much harder.

At this point a triple $([a], [b], [c])$ has been created, with $c \equiv_k a \cdot b$. However, the s most significant bits of c have some information that could allow the adversary to guess the shares of a of the honest parties. Moreover, correctness of the triple is only required modulo 2^k , as this is the modulus in the circuit the parties want to compute. Therefore, in order to mitigate this issue the parties use a random authenticated mask to hide the s most significant bits of c . This mask is very similar to the one used in the procedure **SingleCheck** from Section 3.1. In fact, in an actual implementation we could ignore the mask on the triples, as these will be masked before opening in the MAC checking procedures. However, if we wish to apply the Composition Theorem to our final protocol, each subprotocol must be UC secure by itself, regardless of any further composition.

Now we proceed with the main theorem of the section, which states the security of the protocol in Fig. 12.

Theorem 4. *If $\tau \geq 4k + 2s$, then the protocol Π_{Triple} (Protocol 12) securely implements $\mathcal{F}_{\text{Triple}}$ in the $(\mathcal{F}_{\text{ROT}}, \mathcal{F}_{\text{MAC}}, \mathcal{F}_{\text{Rand}})$ -hybrid model, with statistical security parameter k .*

Proof. Let \mathcal{Z} be an environment, which we also refer to as adversary, corrupting a set A of at most $n - 1$ parties. We construct a simulator \mathcal{S} that has access to the ideal functionality $\mathcal{F}_{\text{Triple}}$ and interacts with \mathcal{Z} in such a way that the real interaction and the simulated interaction are indistinguishable to \mathcal{Z} . Our simulator \mathcal{S} proceeds as follows:

Simulating the Multiply phase The simulator emulates the functionality $\mathcal{F}_{\text{ROT}}^\tau$ and sends $q_{0,h}^{j,i}, q_{1,h}^{j,i} \in \mathbb{Z}_{2^k}$ for $h \in \{1, \dots, \tau\}$ to every $j \in A$ (on behalf of each honest party P_i). When a corrupted party P_j sends $d_h^{j,i}$ to an honest party P_i , $h \in \{1, \dots, \tau\}$, the simulator uses its knowledge on the q 's to extract the values of b used by the adversary as $b_h^j = d_h^{j,i} - q_{0,h}^{j,i} + q_{1,h}^{j,i} \pmod{2^k}$ (notice that if all the parties were honest we would have that all b_h^j for $h \in \{1, \dots, \tau\}$ are equal, however, the adversary can take any strategy and this may not be the case here). The simulator then emulates the multiplication procedure according to the protocol using a fixed consistent value b^j for each $j \in A$ (say the value of b_1^j used with a fixed honest party P_{i_0}). We let $\mathbf{b}^{j,i} \in (\mathbb{Z}_{2^{k+s}})^\tau$ denote the vector of values of b that P_j tried to use in interaction with the emulated honest party P_i in step (c) and we define $\delta_b = \mathbf{b}^{j,i} - \mathbf{b}^j$ (modulo 2^{k+s} on each entry) where \mathbf{b}^j is the vector (b^j, \dots, b^j) . In a similar way, we define $\delta_a^{j,i} = \mathbf{a}^{j,i} - \mathbf{a}^j$ where $\mathbf{a}^j = \mathbf{a}^{j,i_0}$ (these are the errors introduced by P_j when interacting with P_i with respect to the values used in the interaction with P_{i_0}) and $\mathbf{a}^{j,i}$ is the vector that the corrupt party

P_j used in the random OT when interacting with honest party P_i . Notice that $\delta_a^{j,i} \in \{-1, 0, 1\}^\tau$.

Simulating the Combining phase All the computations are local, so \mathcal{S} just emulates $\mathcal{F}_{\text{Rand}}$ and proceeds according to the protocol.

Simulating the Authentication phase Now \mathcal{S} emulates \mathcal{F}_{MAC} with inputs from the corrupt parties provided by \mathcal{Z} . Notice that \mathcal{S} can compute the actual values that each corrupt party should authenticate. The simulator authenticates these and defines e_{Auth} and \hat{e}_{Auth} to be the total error introduced by the adversary in this step. Note that here $e_{\text{Auth}}, \hat{e}_{\text{Auth}} \neq 0$ essentially means that the adversary authenticates values different from those computed in the previous phases. If \mathcal{Z} sends Abort to \mathcal{F}_{MAC} then \mathcal{S} sends Abort to $\mathcal{F}_{\text{Triple}}$.

Simulating the Sacrifice step The simulator opens a uniform value in $\mathbb{Z}_{2^{k+s}}$ as the value of ρ , and aborts if the triple that it has internally stored is incorrect modulo 2^k . Otherwise it stores this triple as a valid triple.

Now we argue that the environment \mathcal{Z} cannot distinguish between the hybrid execution and the simulated one. We begin by noticing that in the **Multiply** phase the adversary only learns the mask $d_h^{i,j}$ for each $i \in A$, but they look perfectly random as the values $q_{1-a_h^j, h}^{i,j}$ are uniformly random and never revealed to \mathcal{Z} . On the other hand, we still need to argue that the value ρ during the **Sacrifice** step has indistinguishable distributions in both executions, and that the triple $([a], [b], [c])$ obtained in the real execution is indistinguishable from the triple generated in the ideal execution (where a and b are uniformly random).

In order to analyze these distributions, we study what is the effect of the adversarial behavior in the final shared value \mathbf{c} , and we do this by considering what happens in the real execution at the end of step 2 when executed by a pair of parties (P_i, P_j) . If both j and i are honest, then the vectors $\mathbf{c}_{i,j}^i$ and $\mathbf{c}_{i,j}^j$ computed at the end of the execution satisfy $\mathbf{c}_{i,j}^i + \mathbf{c}_{i,j}^j \equiv_{k+s} \mathbf{a}^i \cdot b^j$. Also, if j and i are both corrupt then we can safely assume that $\mathbf{c}_{i,j}^i + \mathbf{c}_{i,j}^j \equiv_{k+s} \mathbf{a}^i \cdot b^j$ also holds, since any variation on this will result in an additive error term which depends only in adversarial values and therefore it will get absorbed by the authentication phase. Now suppose that j is corrupt and i is honest, then P_i uses \mathbf{a}^i and P_j uses $\mathbf{b}^{j,i}$, so the vectors $\mathbf{c}_{i,j}^i$ and $\mathbf{c}_{i,j}^j$ computed at the end of the execution satisfy

$$\mathbf{c}_{i,j}^i + \mathbf{c}_{i,j}^j \equiv_{k+s} \mathbf{a}^i \cdot \mathbf{b}^{j,i} \equiv_{k+s} \mathbf{a}^i \cdot \delta_b^{j,i} + \mathbf{a}^i \cdot b^j.$$

Similarly, if i is corrupt and j is honest, then P_i uses $\mathbf{a}^{i,j}$ and P_j uses b^j , so the vectors $\mathbf{c}_{i,j}^i$ and $\mathbf{c}_{i,j}^j$ computed at the end of the execution satisfy

$$\mathbf{c}_{i,j}^i + \mathbf{c}_{i,j}^j \equiv_{k+s} \mathbf{a}^{i,j} \cdot b^j \equiv_{k+s} \delta_a^{i,j} \cdot b^j + \mathbf{a}^i \cdot b^j.$$

Now, if \mathbf{c}^i is the vector obtained by party P_i at the end of the multiplication, then we have the following

$$\begin{aligned}
\mathbf{c} &\equiv_{k+s} \sum_{i=1}^n \mathbf{c}^i \equiv_{k+s} \sum_i \mathbf{a}^i \cdot b^i + \sum_{\substack{i \notin A, j \notin A \\ j \neq i}} (\mathbf{a}^i \cdot b^j) + \sum_{\substack{i \notin A \\ j \in A}} (\mathbf{a}^i * \delta_b^{j,i} + \mathbf{a}^i \cdot b^j) \\
&\quad + \sum_{\substack{i \in A \\ j \notin A}} (\delta_a^{i,j} \cdot b^j + \mathbf{a}^i \cdot b^j) + \sum_{\substack{i \in A, j \in A \\ j \neq i}} (\mathbf{a}^i \cdot b^j) \\
&\equiv_{k+s} \left(\sum_{i=1}^n \mathbf{a}^i \right) \cdot \left(\sum_{i=1}^n b^i \right) + \sum_{\substack{i \notin A \\ j \in A}} \mathbf{a}^i * \delta_b^{j,i} + \sum_{\substack{i \in A \\ j \notin A}} \delta_a^{i,j} \cdot b^j \\
&\equiv_{k+s} \mathbf{a} \cdot b + \underbrace{\sum_{i \notin A} \mathbf{a}^i * \delta_b^i}_{\mathbf{e}_a} + \underbrace{\sum_{j \notin A} \delta_a^j \cdot b^j}_{\mathbf{e}_b}
\end{aligned}$$

where $\mathbf{a} = \sum_{i=1}^n \mathbf{a}^i$, $b = \sum_{i=1}^n b^i$, $\delta_b^i = \sum_{j \in A} \delta_b^{j,i}$ and $\delta_a^j = \sum_{i \in A} \delta_a^{i,j}$, and all congruences are considered component-wise. Notice that each entry in δ_a^j is the sum of at most n bits and therefore it is upper bounded strictly by n , since we assume that $n \ll 2^{k+s}$ we can consider the sum $\mathbf{a} = \sum_{i=1}^n \mathbf{a}^i$ (without the modulus).

Assume all parties (including corrupt ones) take the right linear combination in the combine phase (every adversarial misbehavior will result in an additive error term that only depends on values that the adversary has, and this term will be absorbed by the error term in the authentication phase). Therefore, after the combination and authentication phases the parties obtain values $[b]$, $[a]$, $[c]$, $[\hat{a}]$, $[\hat{c}]$ where $b, a, c, \hat{a}, \hat{c} \in \mathbb{Z}_{2^{k+s}}$ satisfy

$$\begin{aligned}
c &\equiv_{k+s} a \cdot b + e_a + e_b + e_{Auth} \\
\hat{c} &\equiv_{k+s} \hat{a} \cdot b + \hat{e}_a + \hat{e}_b + \hat{e}_{Auth}
\end{aligned}$$

and

$$\begin{aligned}
c &\equiv_{k+s} \sum_{h=1}^{\tau} r_h \cdot \mathbf{c}[h], & \hat{c} &= \sum_{h=1}^{\tau} \hat{r}_h \cdot \mathbf{c}[h] \\
a &\equiv_{k+s} \sum_{h=1}^{\tau} r_h \cdot \mathbf{a}[h], & \hat{a} &\equiv_{k+s} \sum_{h=1}^{\tau} \hat{r}_h \cdot \mathbf{a}[h] \\
e_a &\equiv_{k+s} \sum_{h=1}^{\tau} r_h \cdot \mathbf{e}_a[h], & \hat{e}_a &\equiv_{k+s} \sum_{h=1}^{\tau} \hat{r}_h \cdot \mathbf{e}_a[h] \\
e_b &\equiv_{k+s} \sum_{h=1}^{\tau} r_h \cdot \mathbf{e}_b[h], & \hat{e}_b &\equiv_{k+s} \sum_{h=1}^{\tau} \hat{r}_h \cdot \mathbf{e}_b[h].
\end{aligned}$$

We prove the following two claims in Appendix E, using the same techniques as in the single and batch MAC checking protocols from Section 3, and Lemma 1.

Claim 4 *If the sacrifice step passes, then it holds that $e := e_a + e_b + e_{Auth} \equiv_k 0$ and $\hat{e} := \hat{e}_a + \hat{e}_b + \hat{e}_{Auth} \equiv_k 0$ with probability at least $1 - 2^{-s}$.*

Claim 5 *Suppose that the sacrificing step passes, then all the errors $\{\delta_a^i[h]\}_{h,i \notin A}$ are zero except with probability at most $2^{-k+\log(n \cdot (k+1-\log n))}$*

The previous claim allows us to conclude that $e_b = \hat{e}_b \equiv_{k+s} 0$, except with negligible probability. Now we would like to claim that the value $\rho \in \mathbb{Z}_{2^{k+s}}$ opened in the sacrifice step is indistinguishable from the one opened in the real execution. Since in the ideal execution the simulator opens a uniform value, what we actually need to show is that in a real execution ρ looks (close to) uniform. Given that $\rho = t \cdot a - \hat{a} \pmod{2^k}$, this can be accomplished by showing that \hat{a} looks uniform to the environment. In order to see that $\hat{a} \equiv_{k+s} \sum_{h=1}^{\tau} \hat{r}_h \cdot \mathbf{a}[h] \equiv_{k+s} \sum_{i=1}^n \left(\sum_{h=1}^{\tau} \hat{r}_h \cdot \mathbf{a}^i[h] \right)$ is uniformly distributed it suffices to show that at least for one $i_0 \notin A$ it holds that \hat{a}^{i_0} looks uniform to the environment, where $\hat{a}^i = \sum_{h=1}^{\tau} \hat{r}_h \cdot \mathbf{a}^i[h] \pmod{2^{k+s}}$, and that all these values are actually independent. This can be shown using the Leftover Hash Lemma by giving a good lower bound on the min-entropy of \mathbf{a}^{i_0} (see Appendix E for the precise variant of this lemma that we use). We proceed with the details below.

Using Claim 4 and Claim 5, we have that whenever the sacrifice step passes it holds that

$$-e_{Auth} \equiv_k e_a \equiv_k \sum_{h=1}^{\tau} r_h \cdot e_a[h] \equiv_k \sum_{h=1}^{\tau} r_h \sum_{i \notin A} \mathbf{a}^i[h] \cdot \delta_b^i[h].$$

and

$$-\hat{e}_{Auth} \equiv_k \hat{e}_a \equiv_k \sum_{h=1}^{\tau} \hat{r}_h \cdot e_a[h] \equiv_k \sum_{h=1}^{\tau} \hat{r}_h \sum_{i \notin A} \mathbf{a}^i[h] \cdot \delta_b^i[h].$$

Intuitively, the only information that the adversary has about the honest party's shares is that the sacrifice step passed, which in turn implies that the above equation holds. Ideally, the fact that this relation holds should not reveal so much information about $\{\mathbf{a}^i\}_{i \notin A}$ to the adversary. Indeed, this will be the case, which will be seen when we bound by below the entropy of this random variable. To this end, let $m = n - |A|$ be the number of honest parties and let $S \subseteq \mathbb{Z}_2^{m \cdot \tau}$ be the set of all possible honest shares $(\mathbf{a}^i)_{i \notin A}$ for which the sacrifice step would pass. Notice that in particular, these shares satisfy the equations above and therefore they are completely determined by the errors that are introduced by the adversary. Moreover, since the shares $(\mathbf{a}^i)_{i \notin A}$ are uniformly distributed in S , the min-entropy of these shares is $\log |S|$. Additionally, the vectors in $(\mathbf{a}^i)_{i \notin A}$ are independent one from each other, hence there is at least one honest party P_{i_0} such that the min entropy of \mathbf{a}^{i_0} is at least $\frac{\log |S|}{m}$. In the following we show that $a^{i_0} = \sum_{h=1}^{\tau} r_h \cdot \mathbf{a}^{i_0}[h] \pmod{2^{k+s}}$ and $\hat{a}^{i_0} = \sum_{h=1}^{\tau} \hat{r}_h \cdot \mathbf{a}^{i_0}[h] \pmod{2^{k+s}}$ look random to the environment.

Let β be the probability of passing the sacrifice step, i.e. $\beta = \frac{|S|}{2^{m\tau}} = 2^{-c}$ where $c = m\tau - \log |S|$. We get that

$$H_{\infty}(\mathbf{a}^{i_0}) \geq \frac{\log |S|}{m} = \tau - \frac{c}{m} \geq \tau - c.$$

Now consider the function $h_{\mathbf{r}, \hat{\mathbf{r}}} : (\mathbb{Z}_2)^\tau \rightarrow (\mathbb{Z}_{2^{k+s}})^2$ given by

$$h_{\mathbf{r}, \hat{\mathbf{r}}}(\mathbf{a}) = \left(\sum_{h=1}^{\tau} \mathbf{r}[h] \cdot \mathbf{a}[h] \pmod{2^{k+s}}, \sum_{h=1}^{\tau} \hat{\mathbf{r}}[h] \cdot \mathbf{a}[h] \pmod{2^{k+s}} \right),$$

We claim that this family of functions is 2-universal. Let $\mathbf{a}, \mathbf{a}' \in (\mathbb{Z}_2)^\tau$ such that $\mathbf{a} \neq \mathbf{a}'$, say $\mathbf{a}[h_0] \not\equiv_{k+s} \mathbf{a}'[h_0]$. If $h_{\mathbf{r}, \hat{\mathbf{r}}}(\mathbf{a}) = h_{\mathbf{r}, \hat{\mathbf{r}}}(\mathbf{a}')$ then $\sum_{h=1}^{\tau} \mathbf{r}[h] \cdot (\mathbf{a}[h] - \mathbf{a}'[h]) \equiv_{k+s} 0$ and $\sum_{h=1}^{\tau} \hat{\mathbf{r}}[h] \cdot (\mathbf{a}[h] - \mathbf{a}'[h]) \equiv_{k+s} 0$. Given that \mathbf{a} and \mathbf{a}' are vectors of bits, we have that $\mathbf{a}[h_0] - \mathbf{a}'[h_0] = \pm 1$, so we can solve for $\mathbf{r}[h_0]$ and $\hat{\mathbf{r}}[h_0]$ in the equations above. Therefore, these equations hold with probability at most $\frac{1}{2^{k+s}} \cdot \frac{1}{2^{k+s}} = \frac{1}{2^{2(k+s)}}$ over the choice of $(\mathbf{r}, \hat{\mathbf{r}})$, and hence the family is 2-universal.

According to the Leftover Hash Lemma (see Appendix E), even if the adversary knows \mathbf{r} and $\hat{\mathbf{r}}$, the statistical distance between $h_{\mathbf{r}, \hat{\mathbf{r}}}(X)$ and the uniform distribution in $(\mathbb{Z}_{2^{k+s}})^2$ is at most $2^{-\kappa}$, provided that $H_\infty(X) \geq 2\kappa + 2(k+s)$. This is satisfied if we take $\kappa = \frac{1}{2} \cdot (\tau - c - 2 \cdot (k+s))$.

Finally, ignoring the event in which the check passes with some non-zero errors, which happens with negligible probability, the distinguishing advantage of \mathcal{Z} is the multiplication between the probability of passing the sacrifice step and the probability of distinguishing the output distribution from random, given that the check passed. This is equal to

$$\beta \cdot 2^{-\kappa} = 2^{-c} \cdot 2^{-\frac{1}{2} \cdot (\tau - c - 2 \cdot (k+s))} = 2^{-\frac{\tau - 2 \cdot (k+s) - c}{2}}.$$

Since we want this probability to be bounded by 2^{-s} for any c , we take τ so that $s \leq \frac{\tau - 2 \cdot (k+s)}{2}$, which is equivalent to $\tau \geq 4s + 2k$. \square

7 Efficiency Analysis

We now turn to estimating the efficiency of our preprocessing protocol, focusing on the triple generation phase since this is likely to be the bottleneck in most applications. We emphasise that the costs presented here, compared with those of previous protocols, do not take into account the benefits to applications from working over \mathbb{Z}_{2^k} instead of a finite field with arithmetic modulo a prime. Supporting natural arithmetic modulo 2^k offers advantages on several levels: it simplifies implementations by avoiding the need for modular arithmetic, it reduces the complexity of compiling existing programs into arithmetic circuits, and we believe that it will also be beneficial in performing operations such as secure comparison and bit decomposition of shared values more efficiently than standard techniques using arithmetic modulo p .

Communication complexity of the preprocessing. When authenticating a secret-shared value $x \in \mathbb{Z}_{2^k}$, the main cost is running the vector OLEs, which have inputs over \mathbb{Z}_{2^k} and outputs over $\mathbb{Z}_{2^{k+s}}$, when the MAC key $\alpha \in \mathbb{Z}_{2^s}$. With the method described in Section C, each vector OLE requires s correlated

OTs on messages over \mathbb{Z}_{2^ℓ} , where $\ell = \max(k + s, 2s)$, which gives an amortized cost of $s \cdot \ell$ bits for each component of the vector OLE. We ignore the cost of the consistency check, since this is independent of the number of values being authenticated.

To generate a triple, we need τ random OTs on strings of length $k + s$ bits between every pair of parties, which cost $k + s$ bits of communication each using [13], followed by $\tau \cdot (k + s)$ bits to send the $d^{j,i}$ values. The parties then authenticate 5 values in $\mathbb{Z}_{2^{k+s}}$, which requires generating MACs modulo $\mathbb{Z}_{2^{k+2s}}$ for security. Generating these MACs costs $5s \cdot (k + 2s) \cdot n(n - 1)$ bits of communication using Π_{Auth} based on correlated OT, since the vector OLEs are performed with $\ell = k + 2s$. The costs of $\mathcal{F}_{\text{Rand}}$ and the sacrifice check are negligible compared to this, since the MAC check can be performed in a batch when producing many triples at once. This gives a total cost estimate of $n(n - 1)(5s(k + 2s) + 2\tau(k + s))$ bits per triple. Setting $\tau = 4s + 2k$ (to give failure probability 2^{-s}) this becomes $n(n - 1)(18s^2 + 4k^2 + 17sk)$.

Comparison with MASCOT. Table 1 shows the estimated communication complexity of our protocol for two parties creating a triple in different rings. Note that like MASCOT [14] — the most practical OT-based protocol for actively secure, dishonest majority MPC over finite fields — we expect that communication will be the bottleneck, since the protocol has very simple computational costs. In the table we fix the computational security parameter to 128, and set the statistical security parameter to $s = 64$ in a 64 or 128-bit ring, or $s = 32$ in the 32-bit ring, giving the claimed security bounds (cf. Theorem 1 and Claim 5). Compared with MASCOT, our protocol needs around twice as much communication for 64 or 128-bit triples, with roughly the same level of statistical security. Over the integers modulo 2^{32} , the overhead reduces to around 50% more than MASCOT, although here the statistical security parameters of 26 and 32 bits may be too low for some applications. Note that many applications will not be possible with MASCOT or SPDZ over a 32-bit field, since here integer overflow (modulo p) occurs more easily, and emulating operations such as secure comparison and bit decomposition over a field requires working with a much larger modulus to avoid overflow. When working over $\mathbb{Z}_{2^{32}}$ instead, this should not be necessary.

These overheads for triple generation, compared with MASCOT, come from the fact that our protocol sometimes needs to work in larger rings to ensure security. For example, for the triple check to be secure, our protocol authenticates shares of triples modulo 2^{k+s} , even though the triples are only ever used modulo 2^k in the online phase. This means that when creating these MACs with the protocol from Section 5, we need to work over $\mathbb{Z}_{2^{k+2s}}$ to ensure security. We leave it to future work to try to avoid these costs and improve efficiency.

Comparison with SPDZ using homomorphic encryption. In very recent work [15], Keller, Pastro and Rotaru presented a new variant of the SPDZ protocol that improves upon the performance of MASCOT. In the two-party setting, they show that an optimized implementation of the original SPDZ [9]

| Protocol | Message space | Stat. security | Input cost (kbit) | Triple cost (kbit) |
|----------|------------------------|----------------|-------------------|--------------------|
| Ours | $\mathbb{Z}_{2^{32}}$ | 26 | 3.17 | 79.87 |
| | $\mathbb{Z}_{2^{64}}$ | 57 | 12.48 | 319.49 |
| | $\mathbb{Z}_{2^{128}}$ | 57 | 16.64 | 557.06 |
| MASCOT | 32-bit field | 32 | 1.06 | 51.20 |
| | 64-bit field | 64 | 4.16 | 139.26 |
| | 128-bit field | 64 | 16.51 | 360.44 |

Table 1. Communication cost of our protocol and previous protocols for various rings and fields, and statistical security parameters

runs around twice as fast as MASCOT, and give a new variant that performs 6 times as fast in 64-bit fields; this would probably be around 12 times as fast as our protocol for 64-bit rings. The original SPDZ uses somewhat homomorphic encryption based on the ring-LWE assumption, while their newer variant uses additively homomorphic encryption, and the conjecture that ring-LWE based additively homomorphic encryption has “linear-only” homomorphism. It seems likely that both of these protocols could be adapted to generate triples over \mathbb{Z}_{2^k} using our techniques. One challenge, however, is to adapt the ciphertext packing techniques used in SPDZ for messages over \mathbb{F}_p to the case of \mathbb{Z}_{2^k} , to allow parallel homomorphic operations on ciphertexts; it was shown how this can be done in [10], but it’s not clear how efficient this method is in practice.

Acknowledgements

This work has been supported by the European Research Council (ERC) under the European Unions’s Horizon 2020 research and innovation programme under grant agreement No 669255 (MPCPRO); the European Union’s Horizon 2020 research and innovation programme under grant agreement No 731583 (SODA); the European Union’s Horizon 2020 research and innovation programme under grant agreement No 74079 (ALGSTRONGCRYPTO); and the Danish Independent Research Council under Grant-ID DFF-6108-00169 (FoCC).

References

1. ASHAROV, G., LINDELL, Y., SCHNEIDER, T., AND ZOHNER, M. More efficient oblivious transfer extensions with security for malicious adversaries. In *EUROCRYPT 2015, Part I* (Apr. 2015), E. Oswald and M. Fischlin, Eds., vol. 9056 of *LNCS*, Springer, Heidelberg, pp. 673–701.
2. BENDLIN, R., DAMGÅRD, I., ORLANDI, C., AND ZAKARIAS, S. Semi-homomorphic encryption and multiparty computation. In *EUROCRYPT 2011* (May 2011), K. G. Paterson, Ed., vol. 6632 of *LNCS*, Springer, Heidelberg, pp. 169–188.

3. BOGDANOV, D., LAUR, S., AND WILLEMSON, J. Sharemind: A framework for fast privacy-preserving computations. In *ESORICS 2008* (Oct. 2008), S. Jajodia and J. López, Eds., vol. 5283 of *LNCS*, Springer, Heidelberg, pp. 192–206.
4. CANETTI, R. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS* (Oct. 2001), IEEE Computer Society Press, pp. 136–145.
5. CRAMER, R., FEHR, S., ISHAI, Y., AND KUSHILEVITZ, E. Efficient multi-party computation over rings. In *EUROCRYPT 2003* (May 2003), E. Biham, Ed., vol. 2656 of *LNCS*, Springer, Heidelberg, pp. 596–613.
6. DAMGÅRD, I., DAMGÅRD, K., NIELSEN, K., NORDHOLT, P. S., AND TOFT, T. Confidential benchmarking based on multiparty computation. In *FC 2016* (Feb. 2016), J. Grossklags and B. Preneel, Eds., vol. 9603 of *LNCS*, Springer, Heidelberg, pp. 169–187.
7. DAMGÅRD, I., KELLER, M., LARRAIA, E., PASTRO, V., SCHOLL, P., AND SMART, N. P. Practical covertly secure MPC for dishonest majority - or: Breaking the SPDZ limits. In *ESORICS 2013* (Sept. 2013), J. Crampton, S. Jajodia, and K. Mayes, Eds., vol. 8134 of *LNCS*, Springer, Heidelberg, pp. 1–18.
8. DAMGÅRD, I., ORLANDI, C., AND SIMKIN, M. Yet another compiler for active security or: Efficient MPC over arbitrary rings. Cryptology ePrint Archive, Report 2017/908, 2017. <http://eprint.iacr.org/2017/908>.
9. DAMGÅRD, I., PASTRO, V., SMART, N. P., AND ZAKARIAS, S. Multiparty computation from somewhat homomorphic encryption. In *CRYPTO 2012* (Aug. 2012), R. Safavi-Naini and R. Canetti, Eds., vol. 7417 of *LNCS*, Springer, Heidelberg, pp. 643–662.
10. GENTRY, C., HALEVI, S., AND SMART, N. P. Better bootstrapping in fully homomorphic encryption. In *PKC 2012* (May 2012), M. Fischlin, J. Buchmann, and M. Manulis, Eds., vol. 7293 of *LNCS*, Springer, Heidelberg, pp. 1–16.
11. GOLDBREICH, O., MICALI, S., AND WIGDERSON, A. How to play any mental game or A completeness theorem for protocols with honest majority. In *19th ACM STOC* (May 1987), A. Aho, Ed., ACM Press, pp. 218–229.
12. IMPAGLIAZZO, R., LEVIN, L. A., AND LUBY, M. Pseudo-random generation from one-way functions (extended abstracts). In *21st ACM STOC* (May 1989), ACM Press, pp. 12–24.
13. KELLER, M., ORSINI, E., AND SCHOLL, P. Actively secure OT extension with optimal overhead. In *CRYPTO 2015, Part I* (Aug. 2015), R. Gennaro and M. J. B. Robshaw, Eds., vol. 9215 of *LNCS*, Springer, Heidelberg, pp. 724–741.
14. KELLER, M., ORSINI, E., AND SCHOLL, P. MASCOT: Faster malicious arithmetic secure computation with oblivious transfer. In *ACM CCS 16* (Oct. 2016), E. R. Weippl, S. Katzenbeisser, C. Kruegel, A. C. Myers, and S. Halevi, Eds., ACM Press, pp. 830–842.
15. KELLER, M., PASTRO, V., AND ROTARU, D. Overdrive: Making SPDZ great again. In *EUROCRYPT* (2018), LNCS. <https://eprint.iacr.org/2017/1230>.
16. NIELSEN, J. B., NORDHOLT, P. S., ORLANDI, C., AND BURRA, S. S. A new approach to practical active-secure two-party computation. In *CRYPTO 2012* (Aug. 2012), R. Safavi-Naini and R. Canetti, Eds., vol. 7417 of *LNCS*, Springer, Heidelberg, pp. 681–700.
17. NIELSEN, J. B., SCHNEIDER, T., AND TRIFILETTI, R. Constant round maliciously secure 2pc with function-independent preprocessing using LEGO. In *24th NDSS Symposium* (2017), The Internet Society. <http://eprint.iacr.org/2016/1069>.
18. PETTAI, M., AND LAUD, P. Automatic proofs of privacy of secure multi-party computation protocols against active adversaries. In *IEEE 28th Computer Security*

Foundations Symposium, CSF 2015, Verona, Italy, 13-17 July, 2015 (2015), IEEE, pp. 75–89.

19. SCHOLL, P. Extending oblivious transfer with low communication via key-homomorphic PRFs. In *Public-Key Cryptography (PKC)* (2018), Lecture Notes in Computer Science. <https://eprint.iacr.org/2018/036>.

A Proof of the Online Phase

Theorem 5 (Theorem 2, restated). *The protocol Π_{Online} implements $\mathcal{F}_{\text{Online}}$ in the $\mathcal{F}_{\text{Prep}}$ -hybrid model, with statistical security parameter s .*

Proof. We prove security using the variant of the universal composability framework [4] with a *dummy adversary*, who simply forwards messages sent and received by corrupted parties in the protocol as directed by the environment, \mathcal{Z} . This means we can assume that \mathcal{Z} plays the role of both the distinguisher and the adversary.

Let \mathcal{Z} be an environment corrupting a set A of at most $n - 1$ parties. We construct a simulator \mathcal{S} that has access to the ideal functionality $\mathcal{F}_{\text{Prep}}$ and interacts with \mathcal{Z} in such a way that the real interaction and the simulated interaction are indistinguishable to \mathcal{Z} . Recall that \mathcal{Z} chooses the inputs of all parties, and at the end of the execution also sees the outputs of all parties, including the honest parties. The simulator \mathcal{S} works as follows

Simulating the Initialization and Input phases The simulator simply emulates the functionality $\mathcal{F}_{\text{Prep}}$ honestly. \mathcal{S} also emulates virtual honest parties with dummy inputs. Notice that \mathcal{S} knows all parties' shares of the MAC key α and also the randomness distributed to each party for the input masks and multiplication triples. In particular, \mathcal{S} can respond to any adversary's key guesses to the **Authentication** stage of $\mathcal{F}_{\text{Prep}}$ using the α^i shares it knows, and will abort if any guess fails.

In the **Input** phase, when a corrupted party P_i broadcasts ϵ , \mathcal{S} extracts its input as $x^i = \epsilon + r$, where r is the random value that P_i should have used.

The simulator now uses these values as input to the $\mathcal{F}_{\text{Online}}$ functionality.

Simulating additions This only consists of local computations, which \mathcal{S} carries out honestly on behalf of the virtual honest parties.

Simulating multiplications When the values ϵ and δ for the multiplication are opened, \mathcal{S} opens random shares on behalf of the virtual honest parties.

Simulating the Output and MAC-Checking phases In the output phase, \mathcal{S} first calls $\mathcal{F}_{\text{Online}}$ to obtain the correct output y . Next, to simulate the MAC checking procedure, \mathcal{S} executes `BatchCheck` with the adversary, on behalf of the virtual honest parties. If the check fails then \mathcal{S} sends `Abort` to $\mathcal{F}_{\text{Online}}$.

If the above check passed, \mathcal{S} modifies the honest parties' shares it holds to be consistent with the output y , and also modifies the MAC shares to be a random sharing of $\alpha \cdot y$, which it can do since it knows α . It then runs the `SingleCheck` procedure with the adversary, on behalf of the honest parties. If this aborts then \mathcal{S} sends `Abort` to $\mathcal{F}_{\text{Online}}$, otherwise it sends `Proceed`.

Now we argue that \mathcal{Z} cannot distinguish between the real and ideal executions. This will follow as these are statistically close, as we now show. This is clear for the Initialization phase, where \mathcal{Z} gets random values in both executions. This is also the case in the Input and Multiplication phases, where \mathcal{Z} only see values of the form $x' = x + r$ where r is a fresh uniformly random mask in $\mathbb{Z}_{2^{k+s}}$.

It is only in the output phase where care needs to be taken. First notice that the probability that the Batch and Single MAC-Checking phases result in Abort is the same in both executions. On the other hand, if the first MAC check passes, then the honest parties reveal their shares in both executions. In the real execution these shares are conditioned on adding up (plus the internal shares of the adversary) to the value computed in the protocol, whereas in the simulated execution this sum is equal to the value output by the functionality. However, due to Theorem 1, if this check passes then these two values are the same, except with probability $2^{-s+\log(s+1)}$. Hence the shares revealed have the same distribution in both cases with overwhelming probability. On the other hand, if the single MAC check passes, then due to Claim 1 we have that the final value output by the honest parties in both executions is the same, except with probability 2^{-s} .

Summing up, the overall distinguishing advantage of \mathcal{Z} can be bounded by $2^{-s+\log(s+1)} + 2^{-s}$, which is negligible in s . \square

B Complete Preprocessing Phase

So far we have shown how to create authenticated, secret-shared values and multiplication triples, but are still missing the random input masks created by $\mathcal{F}_{\text{Prep}}$. This can be done using essentially the same protocol as Π_{Auth} from Section 5, except only a single party, P_i , needs to authenticate their input. In more detail, P_i samples the random masks r_1, \dots, r_t (plus one extra for security) and authenticates them using $\mathcal{F}_{\text{VOLE}}$ with every other party. P_i also distributes random shares of the masks to all other parties, so they end up with authenticated sharings $[r_1], \dots, [r_t]$, where only P_i knows r_j , as required. The parties can then use the same MAC checking procedure as in Section 5 to ensure that this was done correctly.

This gives us the complete protocol for preprocessing shown in Fig. 14, which calls the previous MAC and triple generation protocols to realise those parts of $\mathcal{F}_{\text{Prep}}$. Security of the initialize, authentication and triple generation stages follows from the proofs of security of Π_{Auth} and Π_{Triple} . Regarding security of the input phase, note that a corrupted P_i may choose the masks r_j and their shares from a different distribution, which is why $\mathcal{F}_{\text{Prep}}$ allows a corrupt P_i to specify these. This does not pose a problem for the online phase, since the randomness of the masks is only used to protect P_i 's input and not that of any other party. The consistency check can be analysed just as in the Π_{Auth} protocol, since we use exactly the same MAC checking technique there, and this ensures that P_i inputs the same values of \mathbf{r} into the vector-OLE instances with every other party. This gives us the following theorem.

Theorem 6. *The protocol Π_{Prep} securely realises $\mathcal{F}_{\text{Prep}}$ in the $(\mathcal{F}_{\text{VOLE}}, \mathcal{F}_{\text{Rand}}, \mathcal{F}_{\text{ROT}})$ -hybrid model.*

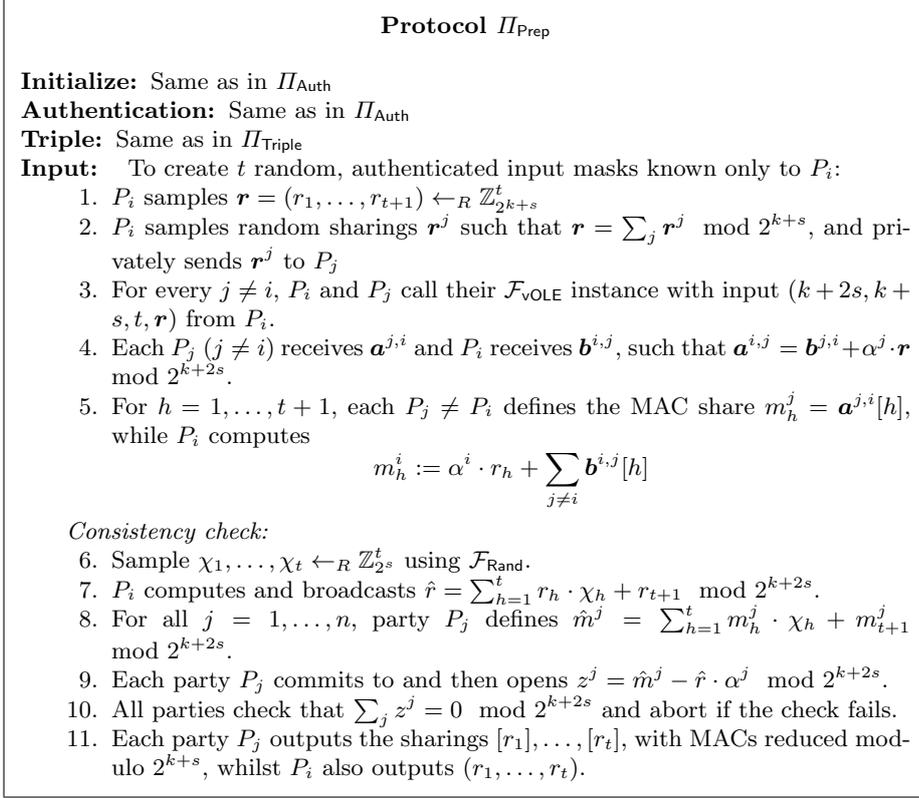


Fig. 14. Complete preprocessing protocol — creating authenticated input masks

C Implementing Vector-OLE mod 2^ℓ

We now describe how to implement the random vector oblivious linear function evaluation functionality $\mathcal{F}_{\text{VOLE}}$ (Fig. 10), which was used to create MAC shares in Section 5, using oblivious transfer. Suppose we have a functionality for OT on correlated strings, where the sender's messages are vectors over \mathbb{Z}_{2^ℓ} of the form $(\mathbf{b}_i, \mathbf{b}_i + \mathbf{x})$, where we are also guaranteed that the correlation \mathbf{x} is the same in every pair.

The receiver, who has an input $\alpha \in \mathbb{Z}_{2^s}$ for the vector-OLE protocol, bit-decomposes α into $\alpha_1, \dots, \alpha_s$ and uses these as choice bits for the correlated OT. The messages received can be written as

$$\mathbf{a}_i = \mathbf{b}_i + \alpha_i \cdot \mathbf{x}, \quad \text{for } i = 1, \dots, s$$

To compute the vector-OLE outputs, the receiver and sender then respectively compute

$$\mathbf{a} = \sum_{i=1}^s \mathbf{a}_i \cdot 2^{i-1} \pmod{2^\ell}, \quad \mathbf{b} = \sum_{i=1}^s \mathbf{b}_i \cdot 2^{i-1} \pmod{2^\ell}$$

and output these vectors, which satisfy $\mathbf{a} = \mathbf{b} + \alpha \cdot \mathbf{x} \pmod{2^\ell}$ as required.

The main challenge in doing this is to implement the correlated OT protocol over \mathbb{Z}_{2^ℓ} with active security. Although most previous protocols [16,17] only support correlations over $GF(2)$, Scholl [19, Sec. 4] recently showed how to create correlated OTs over any finite abelian group using a generalisation of the technique of Asharov et al. [1]. Specifically, the functionality $\mathcal{F}_{\Delta\text{-ROT}}$ from that work can be instantiated over the additive group $\mathbb{Z}_{2^r}^t \times \mathbb{Z}_{2^\ell}$ to obtain the correct form of correlated OTs we need. This requires a setup phase consisting of s oblivious transfers on random strings, which can be realised using \mathcal{F}_{ROT} , and a pseudorandom function. However, this protocol only suffices to realise a *leaky* version of correlated OT, where a corrupt sender may try to guess a few bits of the receiver’s choice bits α_i — if the guess fails then the protocol aborts. This is the reason why we define the $\mathcal{F}_{\text{VOLUME}}$ functionality to allow an adversary to try to guess the honest party’s input α , and abort if the guess fails.

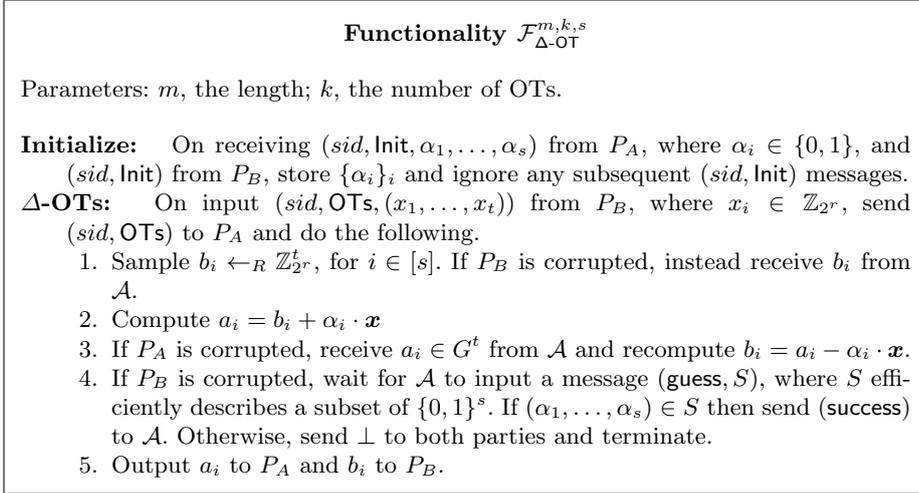


Fig. 15. Correlated random oblivious transfer functionality over $\mathbb{Z}_{2^{k+s}}$

D Proof of the MAC Generation Protocol

Theorem 7 (Theorem 3, restated). *The protocol Π_{Auth} securely realises \mathcal{F}_{MAC} in the $(\mathcal{F}_{\text{VOLUME}}, \mathcal{F}_{\text{Rand}})$ -hybrid model.*

Proof. We construct a simulator \mathcal{S} , such that no environment \mathcal{Z} can distinguish between an execution of the real protocol interacting with the honest parties, and an ideal execution where \mathcal{Z} interacts with \mathcal{S} , who also has access to the ideal functionality \mathcal{F}_{MAC} .

The overall strategy of the simulator is simply to run a copy of all the honest parties “in its head” using all-zero inputs, and use this to simulate the interaction

with the dummy adversary \mathcal{A} , controlled by \mathcal{Z} . \mathcal{S} also extracts the corrupted parties' inputs based on the messages sent to $\mathcal{F}_{\text{VOLE}}$ with a single honest party, as described in the text earlier.

Simulating the Initialize phase. Let $\alpha^{j,i}$ be the input by a corrupt P_j into the $\mathcal{F}_{\text{VOLE}}$ instance with an honest P_i . \mathcal{S} picks an honest party, say P_{i_0} , and sends $\alpha^j := \alpha^{j,i_0}$ to \mathcal{F}_{MAC} for $j \in A$.

Simulating the Authentication phase.

1. For all $j \in A, i \notin A$, \mathcal{S} receives $\hat{\mathbf{x}}^{j,i} = (\mathbf{x}^{j,i}, x_{t+1}^{j,i}) \in \mathbb{Z}_{2^r}^t \times \mathbb{Z}_{2^{\ell'}}$, as well as $\mathbf{a}^{j,i}, \mathbf{b}^{j,i} \in \mathbb{Z}_{2^{\ell'}}^{t+1}$ from \mathcal{A} as inputs to $\mathcal{F}_{\text{VOLE}}$ with the honest P_i .
2. If \mathcal{A} sends any (guess, S) message to $\mathcal{F}_{\text{VOLE}}$ with some honest party P_i , then forward the guess to \mathcal{F}_{MAC} . If \mathcal{F}_{MAC} aborts then abort, otherwise store the set $S_i = S_i \cap S$ (where initially $S_i = \mathbb{Z}_{2^s}$).
3. Sample $\alpha^i \leftarrow_R S_i, x_{t+1}^i \leftarrow_R \mathbb{Z}_{2^{\ell'}}$, for $i \notin A$, on behalf of the honest parties, and then honestly compute their MAC shares m_h^i , for $h = 1, \dots, t+1$, using zero-valued share inputs.
4. Pick an honest party P_{i_0} and let $\hat{\mathbf{x}}^j := \hat{\mathbf{x}}^{j,i_0}$. Write $\hat{\mathbf{x}}^j = (\mathbf{x}^j, x_{t+1}^j) \in \mathbb{Z}_{2^r}^t \times \mathbb{Z}_{2^{\ell'}}$.
5. Sample and send to \mathcal{A} the random values $\chi_1, \dots, \chi_t \leftarrow_R \mathbb{Z}_{2^s}^t$.
6. Send to \mathcal{A} the honestly computed shares \hat{x}^i , for $i \notin A$, and receive back shares that reconstruct to x' .
7. Receive the corrupt parties' commitments z^j , for $j \in A$, and then open the honest parties' commitments as $z^i = \hat{m}^i - x' \cdot \alpha^i$.
8. Carry out the consistency check. If this fails then send \perp to \mathcal{F}_{MAC} and terminate.
9. If the check passes then define the corrupted parties' MACs m_i^j , for $j \in A, i = 1, \dots, t$ using $\alpha^j, \mathbf{x}^j, \mathbf{a}^{j,i}, \mathbf{b}^{j,i}$ according to the protocol, then send (x_i^j, m_i^j) to \mathcal{F}_{MAC} .

Indistinguishability. It is easy to see that the transcript of messages seen by an adversary during the protocol is identically distributed in the real and ideal executions. The probability of passing the consistency check is also identical in both executions, since the previous analysis showed that it only depends on the honest parties' MAC key shares and the random masks x_{t+1}^i , which are identically distributed in the view of the environment in both worlds, until the end of the protocol. It only remains to argue that the MAC shares output by all parties are distributed the same in both worlds.

From Claims 2 and 3, we know that if the consistency check passed in the protocol then the MAC shares form a consistency sharing of a correctly computed MAC, just as the shares output by the ideal functionality do, except with negligible probability. Moreover, the shares of honest parties are uniformly random (subject to this constraint) in both executions, since in the protocol they are obtained by summing up the random outputs of $\mathcal{F}_{\text{VOLE}}$ instances with all

other parties, which serve as random masks. We conclude that the simulation is statistically close. \square

E Preliminaries and Proofs for Triple Generation Protocol

E.1 Min-entropy preliminaries

Definition 1. Let R be any set and $H = \{h_r\}_{r \in R}$ be a family of (keyed) hash functions $h_r : \{0, 1\}^m \rightarrow \{0, 1\}^k$. Then H is a 2-universal hash function family, if for every $x, y \in \{0, 1\}^m$, if $x \neq y$ then

$$\Pr_{r \leftarrow R} [h_r(x) = h_r(y)] \leq 2^{-k}.$$

Definition 2. The min-entropy of a discrete random variable X is defined as

$$H_\infty(X) = -\log \left(\max_x \Pr[X = x] \right)$$

In the proof of Theorem 4, we use the following facts about min-entropy.

- If U is the uniform distribution on a set R then $H_\infty(U) = \log |R|$.
- If $X = (X_1, \dots, X_n)$ is a joint distribution then there is an index i for which $H_\infty(X_i) \geq H_\infty(X)/n$.
- If X and Y are independent distributions over an additive group, then $H_\infty(X + Y) \geq \max(H_\infty(X), H_\infty(Y))$,

Lemma 2 (Leftover Hash Lemma). [12] Let S be a set and ℓ a positive integer. Let X be a random variable over S and $H = \{h_r\}_{r \in R}$, $h_r : S \rightarrow \{0, 1\}^\ell$, a 2-universal hash function. Let U_S be the uniform distribution over S . If

$$H_\infty(X) \geq 2\kappa + \ell$$

then for $t \leftarrow_R \{0, 1\}^\ell$ (independent of X), we have

$$(h_r(X), t) \stackrel{s}{\approx} (U_S, t)$$

for statistical security parameter κ .

E.2 Proofs of claims

Claim 6 (Claim 4, restated) If the sacrifice step passes, then it holds that $e := e_a + e_b + e_{Auth} \equiv_k 0$ and $\hat{e} := \hat{e}_a + \hat{e}_b + \hat{e}_{Auth} \equiv_k 0$ with probability at least $1 - 2^{-s}$.

Proof. Suppose that the protocol does not abort in the sacrifice step, then $\sigma \equiv_{k+s} 0$ which means that $t \cdot c - \hat{c} - b \cdot \rho \equiv_{k+s} 0$. Since $c \equiv_{k+s} a \cdot b + e$ we have that

$$\begin{aligned} 0 &\equiv_{k+s} t \cdot c - \hat{c} - b \cdot \rho \\ &\equiv_{k+s} t \cdot (a \cdot b + e) - (\hat{a} \cdot b + \hat{e}) - b \cdot (t \cdot a - \hat{a}) \\ &\equiv_{k+s} t \cdot e - \hat{e}. \end{aligned}$$

Now suppose that $e \not\equiv_k 0$. Let 2^v be the exact power of 2 that divides e . Since by hypothesis e is not divisible by 2^k we have that $v < k$. Therefore $\frac{e}{2^v}$ is odd (so it has a multiplicative inverse modulo 2^{k-v+s}) which allows us to write

$$t \equiv_{k+s-v} \left(\frac{e}{2^v}\right)^{-1} \left(\frac{\hat{e}}{2^v}\right).$$

Since $s < k+s-v$, this shows that the value of $t \bmod 2^s$ is completely determined, but since t is chosen at random in \mathbb{Z}_{2^s} we conclude that the probability of this particular event is at most 2^{-s} .

Finally, if $e \equiv_k 0$ then $\hat{e} \equiv_k t \cdot e \equiv_k 0$.

Claim 7 (Claim 5, restated) *Suppose that the sacrificing step passes, then all the errors $\{\delta_a^i[h]\}_{h,i \notin A}$ are zero except with probability at most $2^{-k+\log(n \cdot (k+1-\log n))}$*

Proof. Due to the previous claim we have that if the sacrificing step passes then it holds that

$$e_a + e_b + e_{Auth} \equiv_k 0,$$

in particular

$$\begin{aligned} -e_{Auth} &\equiv_k e_a + e_b \equiv_k \sum_{h=1}^{\tau} r_h \cdot e_a[h] + \sum_{h=1}^{\tau} r_h \cdot e_b[h] \equiv_k \sum_{h=1}^{\tau} r_h (e_a[h] + e_b[h]) \\ &\equiv_k \sum_{h=1}^{\tau} r_h \left(\sum_{i \notin A} a^i[h] \cdot \delta_b^i[h] + \sum_{i \notin A} \delta_a^i[h] \cdot b^i \right) \\ &\equiv_k \sum_{i \notin A} b^i \underbrace{\left(\sum_{h=1}^{\tau} r_h \delta_a^i[h] \right)}_{S_i} + \sum_{h=1}^{\tau} r_h \left(\sum_{i \notin A} a^i[h] \cdot \delta_b^i[h] \right). \end{aligned}$$

Suppose that $\delta_a^{i_0}[h_0]$ is not zero, as $0 < |\delta_a^{i_0}[h_0]| < n$ this means that $\delta_a^{i_0}[h_0] \not\equiv_{\log n} 0$. We can write the equation above as $b^{i_0} \cdot S_{i_0} \equiv_k \Delta$ where

$$\Delta = - \sum_{i \notin A, i \neq i_0} b^i \underbrace{\left(\sum_{h=1}^{\tau} r_h \delta_a^i[h] \right)}_{S_i} - \sum_{h=1}^{\tau} r_h \left(\sum_{i \notin A} a^i[h] \cdot \delta_b^i[h] \right).$$

Notice that Δ is a random variable which is independent from b^{i_0} . If we set $\ell = k$, $r = \log n$ and $m = k$ in Lemma 1, we get that if the sacrifice step passes then the probability of having a non-zero error in $\{\delta_a^i[h]\}_{h,i \notin A}$ is bounded by $2^{-\ell+r+\log(\ell-r+1)} = 2^{-k+\log(n)+\log(k-\log(n)+1)}$.