Zaphod: Efficiently Combing LSSS and Garbled Circuits in SCALE

Abdelrahman Aly

Emmanuela Orsini

abdelrahman.aly@esat.kuleuven.be

emmanuela.orsini@kuleuven.be

imec-COSIC, KU Leuven, Leuven, Belgium.

Dragos Rotaru

Nigel P. Smart

Tim Wood

dragos.rotaru@esat.kuleuven.be

nigel.smart@kuleuven.be

t.wood@kuleuven.be

imec-COSIC, KU Leuven, Leuven, Belgium.

Dept. Comp. Sci., University of Bristol, Bristol, UK.

ABSTRACT

We present modifications to the MPC system SCALE-MAMBA to enable the evaluation of garbled circuit (GC) based MPC functionalities and Linear Secret Sharing (LSSS) based MPC functionalities along side each other. This allows the user to switch between different MPC paradigms to achieve the best performance. To do this we present modifications to the GC-based MPC protocol of Hazay et al. (Asiacrypt 2017) (to enable it to support reactive computation), and combine different aspects of their pre-processing phase with those of Wang et al. (CCS 2017), in order to optimize our pre-processing protocols. We also give a more efficient method for producing daBits (double authenticated Bits) than that presented in the work of Rotaru and Wood (ePrint 2019). Finally, we examine how the functionality can be integrated within the existing MPC framework SCALE-MAMBA.

CCS CONCEPTS

• **Security and privacy** → **Security protocols**; *Privacy-preserving protocols*; • **Theory of computation** → **Cryptographic protocols**.

KEYWORDS

Multi-Party Computation; Secret Sharing; Garbled Circuits

ACM Reference Format:

Abdelrahman Aly, Emmanuela Orsini, Dragos Rotaru, Nigel P. Smart, and Tim Wood. 2019. Zaphod: Efficiently Combing LSSS and Garbled Circuits in SCALE. In *7th Workshop on Encrypted Computing & Applied Homomorphic Cryptography (WAHC'19)*, November 11, 2019, London, United Kingdom. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3338469.3358943>

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

WAHC'19, November 11, 2019, London, United Kingdom

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6829-2/19/11...\$15.00

<https://doi.org/10.1145/3338469.3358943>

1 INTRODUCTION

Multi-Party Computation (MPC) is a set of techniques that enable different parties to compute a function on their joint (and privately held) data, without disclosing anything about the data itself, bar what can be derived from the result of the function. It was first introduced in the 1980s and it has, over the last ten years, seen a remarkable improvement in terms of performance. Indeed, MPC has reached a level where companies are starting commercial deployments of MPC based solutions.

MPC comes in essentially two flavours (paradigms). One flavour operates via linear secret-sharing schemes (LSSSs) over some finite field $\mathbb{F}_p$, or sometimes a finite ring. The function to be computed is expressed as an arithmetic circuit (this is not quite true, but this abstraction is good enough for us). Secure computation proceeds by evaluating the circuit via a series of messages being exchanged. We denominate each exchange a *communication round*. The number of rounds depends on the multiplicative depth of the circuit, with the total amount of communication (volume) depending on the number of multiplication gates. Additions do not require message exchanges, hence we consider them to be “free” in terms of communication. Protocols in this family date back to [4, 10], with modern efficient protocols (for various access structures) including those in [5, 11, 12, 27].

The second flavour operates via so-called garbled circuits (GC). In the GC paradigm the function to be evaluated is expressed as a boolean combinatorial circuit. Each AND-gate in the circuit is “encrypted” via a garbled table. In the first phase of computation the parties produce these encrypted tables, then in the second phase the parties evaluate the function using the tables. In this case, the total amount of communication depends on the number of AND gates, but the round complexity is constant. In the passively secure two party case this idea goes back to Yao [29], which was generalised to the honest-majority multi-party setting by Beaver et al. [3]. Recently, there has been renewed interest in n -party actively-secure protocols in this paradigm, such as [16, 28]. The focus in these newer protocols is almost always on full-threshold access structures.

It is clear that these two flavours of MPC are mutually complementary, i.e. evaluation of some functions is more suited to the LSSS paradigm, and for others it is more suited to the GC paradigm. For example, we would prefer to use boolean circuits to perform computation over the integers where the operation is better expressed

as a boolean circuit, and arithmetic circuits where the computation to be performed is best expressed as an arithmetic circuit. Thus we want to combine $\mathbb{Z}_p$ operations and Boolean operations, and indeed this is the goal of this work.

We call our system for doing this Zaphod, after the two headed character **Zaphod Beeblebrox** from Hitchhiker’s Guide to the Galaxy.

Let us consider the following example, where a set of parties wants to compute the function $F : \mathbb{F}_p^{n+m} \rightarrow \{0, 1\}$ defined by

$$F(x_1, \dots, x_n, y_1, \dots, y_m) := \left(\sum_{i=1}^n x_i^2 \right) > \left(\sum_{i=1}^m y_i^2 \right).$$

In this case, we might perform the summations of squares using arithmetic circuits over a field $\mathbb{F}_p$, where p is large enough prime to ensure that the summation is valid, i.e. no overflow occurs. We might nonetheless prefer to perform the comparison using a boolean circuit, since comparison is hard in the LSSS paradigm, but integer addition is free in the LSSS paradigm and integer multiplication is also cheaper in the LSSS paradigm.

However, a major drawback has been that almost all systems have required one to choose a specific paradigm for one’s MPC computations. While significant work has been done on rectifying this in the two-party setting with passive security [6, 7, 9, 13, 17, 22] and some work in the three-party honest-majority active security setting [24], such protocols are often tailored to specific choices of GC and LSSS protocols.

By contrast, Rotaru and Wood [26] recently introduced a mechanism, called daBits (for doubly-authenticated Bits), to combine the GC and LSSS paradigms for general actively secure n -party computation. The idea there is to make use of the GC and LSSS protocols as black boxes to the fullest possible extent to allow greater flexibility in choosing the two protocols. This is done by the parties generating bits which are authenticated both in the GC and LSSS world, and then using these bits to transfer data from one world to the other, and then back again. Whilst this is good for general purpose computation, there are significant engineering issues to be addressed to apply this idea within a working and practical system. In particular, how such constructs are to be exposed to the user of the system? What data-types one will support on either side? How the system will be integrated?

Our Contribution. In this paper we answer these questions, and explain an integration of this idea into the LSSS-based SCALE-MAMBA system [1]. SCALE-MAMBA implements state of the art MPC protocols over finite fields $\mathbb{F}_p$, for various access structures. A number of issues need to be addressed in such an endeavour:

- SCALE-MAMBA supports reactive computation (i.e. data can be processed so that the system finishes a specific computation to obtain an intermediate disclosed result, before continuing with another secure computation using the same data). For this, we need to modify existing GC based MPC protocols to also support reactive computation. This in turn influences our selection of the underlying n -party GC protocol. Thus, in this work, we adapt the BMR-style protocol of [16]. The reason for selecting [16] over, say [28] is that the former is role symmetric and it is easier to turn into a reactive MPC protocol compared to [28].

Figure 1: Our functionalities

- The GC-based protocol of [16] makes use pre-processing of so-called authenticated bits and authenticated AND-triples, via OT-extension. There are a number of different ways of producing such triples e.g., Oblivious Transfer (OT). Thus we need to select the ones which are most efficient from the existing literature.
- The original method of generating daBits was very costly, involving expensive cut-and-choose and bucketing techniques. In this work we give a method that deals explicitly with shares of secrets (rather than handles as described in [26]), which allows the generation of daBits with total communication cost that is up to one quarter less and with a throughput of up to twice as much.
- SCALE-MAMBA comes with a programming language which enables a user to easily program the system. Producing boolean circuit level abstractions is a cumbersome task. Instead of users producing circuit descriptions of functions, it is more desirable to have a suitable high level abstraction in which a user can program her functions on the GC side of the computation. For this, we allow the user to interact with GC computations via an abstraction of signed 64-bit integers, much like any high level programming language would do, for access to real circuits in silicon.

Despite the LSSS operations being relatively general in regard to which access structure they support, the GC operations are set to be secure for full threshold access structures only. Our overall protocol is to be secure for the access structure supported by the underlying LSSS, assuming output to parties is always performed in the LSSS scheme.

2 PRELIMINARIES

Here we define the basic notation and constructs needed to describe our protocols. First, we explain the abstract LSSS-based computation which we will support; here we are totally general and allow n -party secure computation over a number of access structures (all of which are supported in the SCALE-MAMBA system). We also specify our notation for authenticated secret-shared values on the $\mathbb{F}_p$ -side of the computation. Then, we introduce the equivalent authenticated sharing for bits on the GC side of the computation. We conclude this section by giving the abstract functionality which our modification to SCALE-MAMBA aims to achieve.

Notation. We let κ (resp. sec) denote the computational (resp. statistical) security parameter and $\mathcal{P} = \{P_1, \dots, P_n\}$ be the set of parties. In practice we set $\kappa = 128$ (so as to use AES in our PRF) and, in our experiments we utilized $\text{sec} = 40$ and $\text{sec} = 80$. We denote by $[d]$ the set of integers $\{1, \dots, d\}$ and use bold lower-case letters for vectors of elements over both binary and arithmetic fields.

The protocols presented in this work are all secure in the Universal Composability framework of Canetti [8]. We consider security against a static, malicious adversary. For the most part, the protocols presented here are secure for any full-threshold or Q2 access structure. In some of our protocols we will need a coin-tossing functionality $\mathcal{F}_{\text{Rand}}$, which given a set $\mathcal{D}$, outputs a uniformly random element r from $\mathcal{D}$.

2.1 LSSS-Based MPC for Arithmetic Circuits

In this section we detail the type of secret-sharing-based MPC we will be utilizing, and the associated access structures.

We let $\langle \cdot \rangle_p$ denote a linear secret sharing scheme (LSSS) over the finite field $\mathbb{F}_p$ which realizes either a full threshold access structure or a Q2-access structure. There are various MPC protocols that enable actively secure MPC with abort to be carried out using such an LSSS. For example in the full-threshold case, one can define an LSSS $\langle \cdot \rangle_p$ and an associated MPC protocol, using the SPDZ [12] system. For general Q2 (actually *multiplicative*) LSSSs one can utilize the protocol by Smart and Wood [27] to produce an MPC protocol directly using the given LSSS. Whereas using a modified LSSS one can do the same using the protocol given by Chida et al. [11]. The difference between the last two protocols is that [27] provides an online phase with lower communication complexity than [11], at the expense of having to utilize a slightly more expensive offline phase.

We now describe what is meant to share secrets using the LSSS given by $\langle \cdot \rangle_p$, using the well-known description of monotone span programs (MSPs) first given by Karchmer and Wigderson [18]. Let $M \in \mathbb{F}_p^{m \times k}$ be a matrix, choose a non-zero “target” vector $\mathbf{t} \in \mathbb{F}_p^k$ and an index function $\iota : \{1, \dots, m\} \rightarrow \{1, \dots, n\}$. To share a secret s using these data, the dealer samples a vector $\mathbf{k} \leftarrow \mathbb{F}_p^k$ such that $\mathbf{t} \cdot \mathbf{k}^\top = s \in \mathbb{F}_p$, sets $\mathbf{s} = (s_1, \dots, s_m) = M \cdot \mathbf{k}$, and for each $j \in [m]$ computes $i = \iota(j)$ and sends s_i to party P_i over a secure channel. The matrix is chosen in such a way that for any qualified set of parties, there is a (public) recombination vector $\mathbf{r}$ that given the share vector $\mathbf{s}$ (i.e. the concatenation of shares held by the qualified set of parties) can recover the secret by computing $s = \mathbf{r} \cdot \mathbf{s}^\top$.

To ensure correctness in the case of a malicious adversary, we can add authentication to LSSSs in different ways, depending on the access structure. In the SPDZ protocol, every secret has a linear information-theoretic message authentication code (MAC) attached to it: the parties hold a secret-shared global MAC key $\langle \alpha \rangle_p$, and for every secret x , the parties hold $\langle x \rangle_p$ and $\langle \alpha \cdot x \rangle_p$. At the end of a computation, the MACs are checked to reveal whether such additive errors were introduced. This means that in order to change a secret in the scheme (without detection) by adding error ε to a share of the secret, the adversary has to modify a share on the corresponding MAC share by $\alpha \cdot \varepsilon$ – i.e. it must correctly guess the MAC key α . For a Q2 access structure, it has been shown [27] that

error-detection comes essentially “for free” using a kind of internal redundancy of the LSSS.

Both of these approaches to authentication are linear, so in fact they can be described in essentially the same way, as follows. First, note that the MAC’d SPDZ sharing can be described as above by setting the target vector as $\mathbf{t} = (1, \dots, 1, 0, \dots, 0) \in \mathbb{F}_p^{2 \cdot n-1}$, the index function as $\iota(i) = \iota(n+i) = i$, for $i \in [n]$, and $M = (I_{2n-1} | \mathbf{m}^\top)^\top$, where $\mathbf{m} = (\alpha, \dots, \alpha, -1, \dots, -1)$. There is a matrix H (the parity-check matrix of the code generated by M) which takes the full set of shares $\mathbf{s}$ (a vector consisting of both shares of the secret and shares of the corresponding MAC) and checks whether they are *valid*. A set of shares is the output of the above procedure if and only if $H \cdot \mathbf{s}^\top = \mathbf{0}$.

In MPC based on Q2 access structures, the error-detection process involves the same multiplication by a parity-check matrix H of the LSSS matrix, where in this case the matrix H is public, whereas for SPDZ the matrix H is secret. The difference between the protocols in the Q2 and full threshold case are (essentially) in how one uses H to check shares are valid, i.e. to detect whether the protocol should abort or not. We will denote by $\llbracket \cdot \rrbracket_p$ values in $\mathbb{F}_p$ that are secret shared according to an LSSS and authenticated.

2.2 BMR-style MPC for Binary Circuits

Our multiparty GC-based protocol is a combination, with some extensions, of protocols described in [14, 16, 19, 28]. We let $\mathbb{K} = \mathbb{F}_{2^\kappa}$ be a finite field, recall we use $\mathbb{K} = \mathbb{F}_{2^{128}}$ in our implementation. For each element $x \in \mathbb{F}_2$ (resp. $\mathbb{K}$), we denote $\langle x \rangle_2$ (resp. $\langle x \rangle_{\mathbb{K}}$) the *unauthenticated* additive sharing of x over $\mathbb{F}_2$ (resp. $\mathbb{K}$), where $x = \sum_{i \in [n]} x_i$, with party P_i holding $x_i \in \mathbb{F}_2$ (resp. $\mathbb{K}$).

Similarly to the $\mathbb{F}_p$ case, we denote by $\llbracket b \rrbracket_2$ a bit that is linearly secret shared according to $\langle \cdot \rangle_2$ and *authenticated* according to the pairwise BDOZ-style MAC introduced by Bendlin et al. [5]. This means that every party P_i authenticates their share b_i towards party P_j , for each $j \neq i$, by holding a MAC $M_i^j \in \mathbb{K}$, such that

$$M_i^j = K_j^i + b_i \cdot \Delta_j \in \mathbb{K},$$

where P_j holds the local key $K_j^i \in \mathbb{K}$ and the fixed global MAC key $\Delta_j \in \mathbb{K}$. In what follows, we let $\underline{\Delta}$ denote the vector $(\Delta_1, \dots, \Delta_n)$, and denote by $\llbracket b_i \rrbracket_2^{ij}$ this 2-party authenticated sharing of the bit b_i . We can also define an n -party representation of a bit $b = \sum_{i=1}^n b_i$, where each P_i holds the bit-share b_i , $n-1$ MACs M_i^j , $n-1$ local keys K_i^j and Δ_i , i.e.

$$\llbracket b \rrbracket_2 = \{b_i, \Delta_i, \{M_i^j, K_i^j\}_{j \neq i}\}_{i \in [n]}.$$

Clearly, two sharings can be added together and multiplied by a constant in the trivial way, so that also the binary $\llbracket \cdot \rrbracket_2$ -representation is linear. Addition of constants, $\llbracket a \rrbracket_2 \leftarrow \llbracket b \rrbracket_2 + c$ for $c \in \{0, 1\}$, is done by setting

$$a_1 \leftarrow b_1 + c \quad \text{and} \quad a_i \leftarrow b_i \quad \text{for } i \neq 1,$$

with the MAC values and keys being updated by setting

$$\begin{aligned} M_1^j[a] &\leftarrow M_1^j[b], & K_j^1[a] &= K_j^1[b] \oplus c \cdot \Delta_j \quad \text{for } 1 \leq j \leq n, \\ M_i^j[a] &\leftarrow M_i^j[b], & K_i^j[a] &\leftarrow K_i^j[b] \quad \text{for } i \neq 1, 1 \leq j \leq n. \end{aligned}$$

2.3 Zaphod - Online Functionality

Our MPC protocol is defined in the so-called pre-processing model, with an offline phase that provides all the correlated randomness needed for the online computation. In this section we explain the online functionality we aim to provide. Our goal is to describe a reactive MPC protocol which combines *both* arithmetic and boolean circuits. When using the mapping from binary to arithmetic circuits we will assume the data being transferred has bit length bounded by ℓ . In other words every data item x can be represented using in ℓ bits or is equivalently an element of $\mathbb{Z}_{\langle 2^\ell \rangle} = \{-2^{\ell-1} + 1, \dots, 2^{\ell-1}\}$. In our implementation we select $\ell = 64$, but of course in practice ℓ can vary depending on the application. We select $\ell = 64$ so as to easily map onto the *unsigned int* data-type known to programmers.

As mentioned before, we let sec denote a security parameter related to the statistical security of our protocols. To enable more efficient protocols we will make one of the following three assumptions.

- (1) Either $\ell + \text{sec} < \log_2 p$,
- (2) Or $(p - 2^{\lceil \log_2 p \rceil})/p < 2^{-\text{sec}}$,
- (3) Or $(2^{\lceil \log_2 p \rceil} - p)/2^{\lceil \log_2 p \rceil} < 2^{-\text{sec}}$.

In each of these cases we guarantee that selecting $\log_2 p$ bits b_i at random and then forming, for $x \in \mathbb{Z}_{\langle 2^\ell \rangle}$,

$$x + \sum_{i=0}^{\lceil \log_2 p \rceil} b_i \cdot 2^i \pmod{p}$$

will statistically hide the value of x . The assumptions are needed as the value $\sum b_i \cdot 2^i \pmod{p}$ is not necessarily a uniformly random element from $\mathbb{F}_p$. In Figure 2 and Figure 3 we provide the functionality for our MPC black box. Each value in $\mathcal{F}_{\text{MPC}}$ is uniquely identified by an identifier $\text{varid} \in I$, where I is a set of valid identifiers, and a domain set $\text{domain} \in \{\mathbb{F}_p, \mathbb{F}_2^\ell\}$. We model the functionality in such a way that it is independent of the details of the authentication technique used in the protocol, and captures all the MPC computations enabled by the SCALE-MAMBA system. This abstraction, when instantiated with $\mathbb{F}_p$, captures both the SPDZ protocol and the Q2 protocol from [27], but not the protocol from [11]. The modifications needed to also encompass the protocol from [11] are straightforward, but would lead to a more complex description from which we spare the reader. When the functionality is instantiated with $\mathbb{F}_2^\ell$, it is intended to implement the BMR-protocol. Also note that our abstract MPC functionality provides actively secure MPC with abort.

Hence, one can see $\mathcal{F}_{\text{MPC}}$ - **Zaphod Evaluation** as a combination of two MPC black boxes, specified by the set assigned to domain , along with two conversion routines, *ConvertToField* and *ConvertToBinary*, given in $\mathcal{F}_{\text{MPC}}$ - **Zaphod Conversion**. If $\text{domain} = \mathbb{F}_p$, the MPC black box provides arithmetic operations over the finite field $\mathbb{F}_p$, whereas if $\text{domain} = \mathbb{F}_2^\ell$, it enables one to execute arbitrary binary circuits over binary vectors of length ℓ , i.e. functions with arguments and results in the set $\mathbb{F}_2^\ell$. Note that the execution of **Output** with $\text{type} = -1$ is a command required for simulation in reactive functionalities.

$\mathcal{F}_{\text{MPC}}$ - **Zaphod Conversion** permits parties to switch between the two MPC black boxes. Note that we have defined *ConvertToBinary* and *ConvertToField* to ensure that they are mutual inverses of each

other (if the $\mathbb{F}_p$ -input element is fewer than ℓ bits in length when in the centred interval $(-p/2, \dots, p/2]$).

Functionality $\mathcal{F}_{\text{MPC}}$ - Zaphod Evaluation

The functionality runs with parties $P_1, \dots, P_n$ and an ideal adversary $\mathcal{A}$. Let A be the set of corrupt parties. Given a set I of valid identifiers, all values are stored in the form $(\text{varid}, \text{domain}, x)$, where $\text{varid} \in I$, $\text{domain} \in \{\mathbb{F}_2^\ell, \mathbb{F}_p\}$ and $x \in \text{domain}$. We assume p is restricted as in the main text.

Initialize: On input (Init) from all parties, the functionality activates. If (Init) was received before, do nothing.

Input: On input $(\text{Input}, P_i, \text{varid}, \text{domain}, x)$ from P_i and $(\text{input}, P_i, \text{varid}, \text{domain})$ from all other parties, with varid a fresh identifier, store $(\text{varid}, \text{domain}, x)$.

Evaluate: Upon receiving $(\{\text{varid}_j\}_{j \in m}, \text{varid}, \text{domain}, C_{\hat{f}})$, from all parties, where $\hat{f} : \{\text{domain}\}^m \rightarrow \text{domain}$ and varid is a fresh identifier, if $\{\text{varid}_j\}_{j \in [m]}$ were previously stored, proceed as follows:

- (1) Retrieve $(\text{varid}_j, \text{domain}, x_j)$, for each $j \in [m]$
- (2) Store $(\text{varid}, \text{domain}, x_{m+1} \leftarrow \hat{f}(x_1, \dots, x_m))$

Output: On input $(\text{Output}, \text{varid}, \text{domain}, \text{type})$, from all parties (if varid is present in memory):

- (1) If $\text{type} = 0$ (**Public Output**): Retrieve (varid, y) and send it to $\mathcal{A}$. If the adversary sends Deliver, send y to all parties.
- (2) If $\text{type} = -1$ (**No Output**): Retrieve (varid, y) and send it to $\mathcal{A}$. If the adversary sends Deliver, store (varid, y) and continue.
- (3) Otherwise (**Private Output**): Send (varid) to $\mathcal{A}$. Upon receiving Deliver from $\mathcal{A}$, send y to P_i

Abort: The adversary can at any time send *abort*, upon which send *abort* to all honest parties and halt.

Figure 2: The ideal functionality for MPC with Abort over $\mathbb{F}_p$ and $\mathbb{F}_2^\ell$ - Evaluation

Functionality $\mathcal{F}_{\text{MPC}}$ - Zaphod Conversion

Convert To Field: On input $(\text{Convert}, \text{varid}_1, \mathbb{F}_2^\ell, \text{varid}_2, \mathbb{F}_p)$:

- (1) Retrieve $(\text{varid}_1, \mathbb{F}_2^\ell, x)$ and convert x to an element $y \in \mathbb{F}_p$ by setting $y \leftarrow -x_{\ell-1} \cdot 2^{\ell-1} + \sum_{i=0}^{\ell-2} x_i \cdot 2^i$.
- (2) Store $(\text{varid}_2, \mathbb{F}_p, y)$.

Convert To Binary On input $(\text{Convert}, \text{varid}_1, \mathbb{F}_p, \text{varid}_2, \mathbb{F}_2^\ell)$:

- (1) Retrieve $(\text{varid}_1, \mathbb{F}_p, x)$ as an integer in the range $(-p/2, \dots, p/2)$.
- (2) Express $y = x \pmod{2^\ell}$ as $y = \sum_{i=0}^{\ell-1} y_i \cdot 2^i$ for $y_i \in \{0, 1\}$
- (3) Consider the values y_i as elements in $\mathbb{F}_2$ and pack them into a vector $y \in \mathbb{F}_2^\ell$.

Figure 3: Extending the ideal functionality for MPC over $\mathbb{F}_p$ and $\mathbb{F}_2^\ell$ to allow conversion

3 PRE-PROCESSING

Here we describe the entire pre-processing phase of Zaphod. At a high level, we need to produce all the correlated randomness needed to implement both the arithmetic and binary MPC functionalities described in the previous section, and the conversion between them.

3.1 SPDZ and TinyOT-style Pre-processing

We recall here that the main task of the offline phase in our LSSS-based MPC in the pre-processing model with active security is

to produce random authenticated secret-shared values and random authenticated multiplication triples to allow the input and the multiplication steps during the online evaluation.

This is formalized by the ideal functionality $\mathcal{F}_{\text{Prep}}[\text{AuthLSSS}]$ described in Figure 4. It is a general pre-processing functionality for LSSS-based MPC, that allows parties to choose between arithmetic and binary fields using the **Initialize** command.

Again, we model our functionality so that it is independent of the authentication method used in the protocol that implements it. However, when the domain set is fixed we concretely instantiate the handle on values with $\llbracket \cdot \rrbracket_p$ if $\text{domain} = \mathbb{F}_p$, and with $\llbracket \cdot \rrbracket_2$ otherwise. We denote the $\mathbb{F}_p$ -instantiations by $\mathcal{F}_{\text{Prep}}[\text{SPDZ}]$.

After the initialization, $\mathcal{F}_{\text{Prep}}[\text{AuthLSSS}]$ allows parties to provide private inputs (**Input**), perform linear operations (**LinearComb**), produce random secret-shared bits (**RandomBit**) and triples (**Triple**), and reveal values either publicly to every party (**Open**) or privately to a single party (**OpenTo**). These open commands permit the adversary to output inconsistent values or shares, so to ensure correctness at a later point in time, the functionality provides the **Check** command. Since we aim for security with abort, the adversary can at any point send the message *abort* to the functionality (**Abort**).

This general pre-processing functionality can be implemented in different ways [12, 20, 21, 25, 27] depending on the access structure. In particular, in the SCALE-MAMBA system (when $\mathbb{F} = \mathbb{F}_p$) the functionality $\mathcal{F}_{\text{Prep}}[\text{SPDZ}]$ is implemented according to the Overdrive-High Gear variant of Keller et al. [21], in the full threshold case, and by the protocol in [27], in the Q2 case. We call both these instantiations $\Pi_{\text{Prep}}[\text{SPDZ}]$. We refer the reader to these papers for a complete description of the protocols and the corresponding security proofs.

3.2 BMR Preprocessing

In the pre-processing phase of modern variants of the BMR protocol, all parties contribute to the garbling of each gate individually. We now describe our implementation of the BMR pre-processing protocol $\Pi_{\text{Prep}}[\text{BMR}]$ which is fully described in Figure 6.

3.2.1 $\mathcal{F}_{\text{n-TinyOT}}$ functionality. The protocol $\Pi_{\text{Prep}}[\text{BMR}]$ assumes the ideal functionality $\mathcal{F}_{\text{Prep}}[\text{AuthLSSS}]$ described in the previous section with $\mathbb{F} = \mathbb{F}_2$. However, to achieve concrete efficiency and ease the description, we consider a slightly modified version of $\mathcal{F}_{\text{Prep}}[\text{AuthLSSS}]$ that is better suited for the $\llbracket \cdot \rrbracket_2^{ij}$ and $\llbracket \cdot \rrbracket_2$ representations described in Section 2.2, similarly to what was done by Hazay et al. [16]. We call this modified functionality $\mathcal{F}_{\text{n-TinyOT}}$. (See [16] for a complete description of it.) We summarize here the main differences compared to $\mathcal{F}_{\text{Prep}}[\text{AuthLSSS}]$:

- $\mathcal{F}_{\text{n-TinyOT}}$.Initialize: The functionality receives from the honest parties and the adversary the MAC key shares $\Delta_i \in \mathbb{K}$, for $i \in [n]$.
- $\mathcal{F}_{\text{n-TinyOT}}$.Get: On input (Get, $\mathbb{F}_2$, *varid*, *i*) from all parties, the functionality samples a random $b_i \in \mathbb{F}_2$ if $i \notin A$, otherwise receives b_i from the adversary. Store (*varid*, $\mathbb{F}_2$, b_i). This command is used by $\mathcal{F}_{\text{n-TinyOT}}$ instead of the private input command in $\mathcal{F}_{\text{Prep}}[\text{AuthLSSS}]$.
- $\mathcal{F}_{\text{n-TinyOT}}$.Triple: On input (Triple, $\mathbb{F}_2$, $\{ \text{varid}_i \}_{i \in [3n+3]}$) the functionality selects $a, b \in \mathbb{F}_2$ at random and set $c \leftarrow a \cdot b$. It receives shares a_i, b_i, c_i from $\mathcal{A}$, for $i \in A$, and samples a_i, b_i, c_i , $i \notin A$,

Functionality $\mathcal{F}_{\text{Prep}}[\text{AuthLSSS}]$

The functionality runs with parties $P_1, \dots, P_n$ and an ideal adversary $\mathcal{A}$. Let A be the set of corrupt parties. Given a set I of identifiers, all values are stored in the form (*varid*, *domain*, *x*), where *varid* $\in I$ and $\text{domain} = \{\mathbb{F}_2, \mathbb{F}_p\}$.

Initialize: On input (*Init*, $\mathbb{F}$) from all parties, the functionality stores (*domain*, $\mathbb{F}$), where $\mathbb{F}$ is either $\mathbb{F}_2$ or $\mathbb{F}_p$.

Input: On input (*Input*, P_i , *varid*, *domain*, *x*) from party P_i and (*input*, P_i , *varid*) from all other parties, with *varid* a fresh identifier and $x \in \text{domain}$, store (*varid*, *x*).

RandomBit: On input (*RandomBit*, *domain*, *varid*) from all parties, with *varid* a fresh identifier, sample $b \in \{0, 1\}$, store (*varid*, *b*).

LinearComb: On input (*LinearComb*, *domain*, *varid*₁, *varid*₂, *varid*₃, *a*, *b*, *c*) from all parties, where *a*, *b*, *c* $\in \text{domain}$, retrieve (*varid*₁, *x*), (*varid*₂, *y*) and store (*varid*₃, *domain*, $a \cdot x + b \cdot y + c$).

Triple: On command (*Triple*, *domain*, *varid*₁, *varid*₂, *varid*₃), for fresh identifiers *varid*₁, *varid*₂, *varid*₃, select *a*, *b* $\in \text{domain}$ at random and set $c \leftarrow a \cdot b$. Store (*varid*₁, *domain*, *a*), (*varid*₂, *domain*, *b*), (*varid*₃, *domain*, *c*).

Open: On input (*Open*, *domain*, *varid*) from all parties (if *varid* is present in memory), functionality retrieve (*varid*, *y*) and output *y* to the adversary. Wait for an input *y'* from $\mathcal{A}$. If this input is *Deliver* then output *y'* to all parties. Otherwise send *abort* to honest parties and halt.

OpenTo: On input (*Open*, *domain*, *varid*, *i*) from all parties (if *varid* is present in memory), retrieve (*varid*, *y*). If $i \in A$, send *y* to $\mathcal{A}$ and wait for a reply, otherwise just wait for an input from $\mathcal{A}$. Output according to the adversary response.

Check: On input (*Check*, *varid*₁, $\dots$, *varid*_{*t*}, *x*₁, $\dots$, *x*_{*t*}) from all parties, wait for an input from the adversary. If the adversary inputs *Proceed*, store the value *x*_{*i*}, for all $i \in [t]$, and output *Proceed* to all parties, otherwise output *abort* and halt.

Abort: The adversary can at any time send *abort*, upon which send *abort* to all honest parties and halt.

Figure 4: The ideal preprocessing functionality for MPC over $\mathbb{F}_p$ and $\mathbb{F}_2$. We denote this functionality by $\mathcal{F}_{\text{Prep}}[\text{SPDZ}]$ when $\text{domain} = \mathbb{F}_p$.

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

This functionality runs with two parties P_S (the sender) and P_R (the receiver) and an ideal adversary $\mathcal{A}$.

Initialize: On input (*Init*, Δ), $\Delta \in \mathbb{K}$ from P_S and (*Init*) from P_R , store Δ .

Extend: On input (*Extend*, $b_1, \dots, b_m$) from P_R and (*Extend*) from P_S proceed as follows:

- (1) Sample $t_i \in \mathbb{K}$, for $i \in [m]$. If P_R is corrupted receive t_i from $\mathcal{A}$.
- (2) Set $q_i \leftarrow t_i + b_i \cdot \Delta \in \mathbb{K}$, for any $i \in [m]$. If P_S is corrupt, receive q_i from $\mathcal{A}$ and set $t_i \leftarrow q_i + b_i \cdot \Delta$ instead.
- (3) Output t_i to P_R and q_i to P_S , for $i \in [m]$

Figure 5: The ideal functionality for COT

subject to $\sum_i a_i = a$, $\sum_i b_i = b$, $\sum_i c_i = c$. It stores these values with the corresponding identifiers.

All the other commands, as $\mathcal{F}_{\text{n-TinyOT}}$.Open and $\mathcal{F}_{\text{n-TinyOT}}$.OpenTo, are similar to those in $\mathcal{F}_{\text{Prep}}[\text{AuthLSSS}]$.

Essentially, this functionality enables the parties to produce random authenticated bits and bit-triples to be used in the protocol $\Pi_{\text{Prep}}[\text{BMR}]$. To implement $\mathcal{F}_{\text{n-TinyOT}}$ in the $\mathcal{F}_{\text{COT}}$ -hybrid model (where the functionality $\mathcal{F}_{\text{COT}}$ is described in Figure 5), we use a

combination of protocols from [14, 16, 19, 28] as follows. We implement $\mathcal{F}_{\text{COT}}$ using the OT-extension protocols described in the two works [14, Full Version, Figure 19] and [19, Full Version, Figure 7]. These correlated OTs are easily converted into random sharings of authenticated bits (called aShares or aBits in previous papers and $\llbracket \cdot \rrbracket_2$ -sharings in this paper), by setting $M_i^j = t_i$ and $K_j^i = q_i$, where t_i and q_i are the output of $\mathcal{F}_{\text{COT}}$. For this step we use [16, Full Version, Figure 16]. Finally, these authenticated bits are converted into triples (aANDs in other works) using [28, Full Version, Figures 16, 8 and 18 in order].

3.2.2 Garbling. Using $\mathcal{F}_{n\text{-TinyOT}}$ we can now describe the protocol $\Pi_{\text{Prep}}[\text{BMR}]$. It is a (*minor*) modification of the n -party garbled circuit construction given in [16]. The main difference is that we are interested in a reactive functionality, i.e. we are going to evaluate multiple different functions F , with the inputs and outputs of this function being vectors of authenticated bits $\llbracket \mathbf{x} \rrbracket_2$. The function we want to compute will be given by

$$\llbracket \mathbf{y} \rrbracket_2 = F(\llbracket \mathbf{x} \rrbracket_2),$$

where $\llbracket \mathbf{x} \rrbracket_2$ and $\llbracket \mathbf{y} \rrbracket_2$ are vectors of authenticated bits. The number of bits in a vector is denoted by $|\mathbf{x}|$, and the j -th bit of $\mathbf{x}$ is denoted by $\mathbf{x}^j$. The common input of all parties in $\mathcal{P}$ is the function to be evaluated given as a binary *combinatorial* circuit consisting of three types of gates: XOR, AND and INV. The circuit representing F is assumed to be given in a topological order.

We describe at a high level the precise garbling being used, and why it works. First, the parties call the $\mathcal{F}_{n\text{-TinyOT}}$ functionality on command **Initialize** to obtain Δ_i , for $i \in [n]$. To each wire w in the circuit C_F we will associate a random wire mask $\llbracket \lambda_w \rrbracket_2$ that the parties obtain by calling $\mathcal{F}_{n\text{-TinyOT}}.\text{Get}$. It will be used to mask the actual value v_w of the wire w , so to define the *external value* as $\epsilon_w = v_w \oplus \lambda_w$. There is a special wire, called the *one wire*, which is used to represent the value one (needed in the INV gate which is given by $\text{INV}(a) = \text{XOR}(a, 1)$). This always has wire mask $\lambda_{\text{one}} = 0$ and, consequently, external value $\epsilon_{\text{one}} = 1$.

The wire masks are chosen independently for all wires in the circuit, *except* for the outputs of XOR gates, where if we have a gate $c = \text{XOR}(a, b)$ then we define $\lambda_c = \lambda_a \oplus \lambda_b$. Note, since an INV gate is a XOR gate this also applies to the outputs of INV gates where we write for $b = \text{INV}(a)$, $\lambda_b = \lambda_a \oplus 1$.

Associated to each wire w , we also have a *vector* of secret *base* keys, $\mathbf{k}_{w,0} = (\mathbf{k}_{w,0}^1, \dots, \mathbf{k}_{w,0}^n) \in \mathbb{K}^n$. These are the garbled circuit keys associated with the zero value on this wire. To enable free-XOR, the value for the one wire is given by $\mathbf{k}_{a,1} = \mathbf{k}_{a,0} + \underline{\Delta}$, where we $\underline{\Delta} = (\Delta_1, \dots, \Delta_n)$ is the *global difference* vector. The key observation [16] is that each party P_i uses the same Δ_i as in the $\llbracket \cdot \rrbracket_2$ -representation, i.e. the MAC keys used to authenticate in $\mathbb{F}_2$ and the global differences used in BMR are the same. The wire keys are often called wire labels in the literature. The *one wire* also has a set of base keys given by a vector $\mathbf{k}_{\text{one},0}$.

We will require a PRF, denoted by $F_{k_0, k_1}(m)$, which takes as input a key in $\mathbb{K}^2$, a one block message m and outputs an element in $\mathbb{K}$. The precise PRF we use, since we select $\mathbb{K} = \mathbb{F}_{2^{128}}$ is derived from Matyas–Meyer–Oseas hashing

$$M = m \oplus (k_0 \cdot X) \oplus (k_1 \cdot X^2),$$

$\Pi_{\text{Prep}}[\text{BMR}]$

COMMON INPUT: All the parties hold the circuit C_F to be evaluated in a topological order.

Initialize: Each party P_i calls $\mathcal{F}_{n\text{-TinyOT}}$ to obtain the global difference $\Delta_i \in \mathbb{K}$.

Garbling:

- For all wires w which are either an input wire to the circuit or the output wire of an AND gate compute the shared wire mask $\llbracket \lambda_w \rrbracket_2 \leftarrow \mathcal{F}_{n\text{-TinyOT}}.\text{Get}()$.
- For all wires w which are either an input wire or the output wire of an AND gate party P_i samples $\mathbf{k}_{w,0}^i \in \mathbb{K}$ at random and sets $\mathbf{k}_{w,1}^i \leftarrow \mathbf{k}_{w,0}^i + \Delta_i$.
- Each P_i selects $\mathbf{k}_{\text{one},0}^i \in \mathbb{K}$ at random and sets $\mathbf{k}_{\text{one},1}^i \leftarrow \mathbf{k}_{\text{one},0}^i \oplus \Delta_i$.
- The parties now progress through the circuit in the given topological ordering and process the gates as follows:
 - (1) $c = \text{XOR}(a, b)$ Gate:
 - Each party sets $\llbracket \lambda_c \rrbracket_2 \leftarrow \llbracket \lambda_a \rrbracket_2 \oplus \llbracket \lambda_b \rrbracket_2$, $\mathbf{k}_{c,0}^i \leftarrow \mathbf{k}_{a,0}^i \oplus \mathbf{k}_{b,0}^i$ and $\mathbf{k}_{c,1}^i \leftarrow \mathbf{k}_{a,0}^i \oplus \Delta_i$.
 - (2) $c = \text{INV}(a)$ Gate:
 - Each party sets $\llbracket \lambda_c \rrbracket_2 \leftarrow \llbracket \lambda_a \rrbracket_2 \oplus 1$, $\mathbf{k}_{c,0}^i \leftarrow \mathbf{k}_{a,0}^i \oplus \mathbf{k}_{\text{one},0}^i$ and $\mathbf{k}_{c,1}^i \leftarrow \mathbf{k}_{a,0}^i \oplus \Delta_i$.
- For each AND gate $c = \text{AND}(a, b)$ the parties now execute:
 - The parties runs the subprotocol Π_{Mult} to compute $\llbracket a \rrbracket_2 \cdot \llbracket b \rrbracket_2$ and obtain $\llbracket \lambda_{ab} \rrbracket_2 \leftarrow \Pi_{\text{Mult}}(\llbracket \lambda_a \rrbracket_2, \llbracket \lambda_b \rrbracket_2)$.
 - Each party locally obtains a share of $\langle \lambda_a \cdot \Delta_j \rangle_{\mathbb{K}}$, $\langle \lambda_b \cdot \Delta_j \rangle_{\mathbb{K}}$ and $\langle (\lambda_{ab} + \lambda_c) \cdot \Delta_j \rangle_{\mathbb{K}}$.
 - For each $\epsilon_a, \epsilon_b \in \{0, 1\}$ and $j \in [n]$ party i computes its share of $\langle \rho_{j, \epsilon_a, \epsilon_b} \rangle_{\mathbb{K}} = \langle (\lambda_{ab} + \lambda_c) \cdot \Delta_j \rangle_{\mathbb{K}} \oplus (\epsilon_a \cdot \langle \lambda_b \cdot \Delta_j \rangle_{\mathbb{K}}) \oplus (\epsilon_b \cdot \langle \lambda_a \cdot \Delta_j \rangle_{\mathbb{K}}) \oplus (\epsilon_a \cdot \epsilon_b \cdot \Delta_j)$.
 - Party i can now compute its share of the j column of the garbled gate via

$$\langle g_{j, \epsilon_a, \epsilon_b} \rangle_{\mathbb{K}}^j \leftarrow \langle \rho_{j, \epsilon_a, \epsilon_b} \rangle_{\mathbb{K}}^j \oplus F_{\mathbf{k}_{a, \epsilon_a}^j, \mathbf{k}_{b, \epsilon_b}^j}(g \| j) \oplus \mathbf{k}_{c,0}^j$$
 if $j = i$, and otherwise via

$$\langle g_{j, \epsilon_a, \epsilon_b} \rangle_{\mathbb{K}}^i \leftarrow \langle \rho_{j, \epsilon_a, \epsilon_b} \rangle_{\mathbb{K}}^i \oplus F_{\mathbf{k}_{a, \epsilon_a}^i, \mathbf{k}_{b, \epsilon_b}^i}(g \| j),$$
 where g is a unique identifier for this AND-gate.

Open Garbling: The parties open the garbled gates $\langle g_{j, \epsilon_a, \epsilon_b} \rangle_{\mathbb{K}}$, and sum the shares so as to obtain, for each $\epsilon_a, \epsilon_b \in \{0, 1\}$ and $j \in [n]$,

$$g_{j, \epsilon_a, \epsilon_b} \leftarrow \mathbf{k}_{c, \epsilon_c}^i \oplus \left(\bigoplus_{j=1}^n F_{\mathbf{k}_{a, \epsilon_a}^j, \mathbf{k}_{b, \epsilon_b}^j}(g \| j) \right).$$

The party i sends $\mathbf{k}_{\text{one},1}^i$ to all parties.

Figure 6: The BMR pre-processing protocol

$$F_{k_0, k_1}(m) = \sigma(M) \oplus \text{AES}_0(\sigma(M)),$$

where $k \cdot X$ denotes multiplication by X in the finite field $\mathbb{K}$ (for X the formal root of the defining polynomial of $\mathbb{K}$, i.e. we have $\mathbb{K} = \mathbb{F}_2/(f(X))$), and $\sigma(M)$ is the function on 128-bit datablocks defined in [15] by $\sigma(M_L \| M_R) = (M_R \oplus M_L) \| M_L$, where M_L and M_R are the left and right halves of the message M respectively.

A *garbled gate* is a means of translating the input wire keys (for actual wire values of 0 or 1) into the wire keys for the output of the gate. These are only needed for AND gates (and any other non-linear gate one wishes to support). The garbled gate, associated with $c = \text{AND}(a, b)$, is indexed by the *external values* ϵ_a and ϵ_b of the

input wires. Note, in this case we have $v_c = v_a \cdot v_b$ as we are evaluating an AND gate. The key trick in [16] is to enable the Tiny-OT authenticated shares $\llbracket \lambda_a \rrbracket_2$, $\llbracket \lambda_b \rrbracket_2$ and $\llbracket \lambda_c \rrbracket_2$ to be converted *without interaction* into a sharing $\langle \rho_{\epsilon_a, \epsilon_b, j} \rangle_{\mathbb{K}} = \langle E(\lambda_a, \lambda_b, \lambda_c, \epsilon_a, \epsilon_b) \cdot \Delta_j \rangle_{\mathbb{K}}$ for every value of $j \in [n]$. Thus each party P_i can derive a share $\langle \rho_{\epsilon_a, \epsilon_b, j} \rangle_{\mathbb{K}}^i$, for the values $\epsilon_a, \epsilon_b \in \{0, 1\}$ without knowing the values of λ_a, λ_b and λ_c . This is done via the following Lemma

LEMMA 3.1. *Let $\llbracket x \rrbracket_2$ be an authenticated bit with value/share/-MAC values (x_i, M_i^j, K_i^j) then*

$$y_i = M_i^j \text{ for } i \neq j,$$

$$y_j = (x_j \cdot \Delta_j) \oplus \left(\bigoplus_{k \neq j} K_k^j \right).$$

is a $\langle \cdot \rangle_{\mathbb{K}}$ -sharing of $x \cdot \Delta_j$.

To encrypt the gate each party P_i encrypts their *output* base wire key $k_{c,0}^i$ using the PRF applied to the external wire keys k_{a,ϵ_a}^i and k_{b,ϵ_b}^i . They also add on the sharing $\langle \rho_{\epsilon_a, \epsilon_b, j} \rangle_{\mathbb{K}}^i$. The other parties for this output key (column) do a similar operation but do not add on the base key. Each garbled gate consists of a table of $(4 \times n)$ elements in K held by party P_i . When one opens the gates at the end of garbling these tables are exchanged and added together.

The sub-protocol Π_{Mult}	
$\Pi_{\text{Mult}}(\llbracket x \rrbracket_2, \llbracket y \rrbracket_2)$:	This produces a sharing of $x \cdot y$.
(1)	Call <i>Triple</i> on $\mathcal{F}_{\text{n-TinyOT}}$ to obtain a triple $(\llbracket a \rrbracket_2, \llbracket b \rrbracket_2, \llbracket c \rrbracket_2)$.
(2)	The players compute $\llbracket d \rrbracket_2 \leftarrow \llbracket a \rrbracket_2 \oplus \llbracket b \rrbracket_2$ and $\llbracket e \rrbracket_2 \leftarrow \llbracket b \rrbracket_2 \oplus \llbracket y \rrbracket_2$.
(3)	The parties call $\mathcal{F}_{\text{n-TinyOT}}.\text{Open}(\llbracket d \rrbracket_2)$ and $\mathcal{F}_{\text{n-TinyOT}}.\text{Open}(\llbracket e \rrbracket_2)$ to obtain d and e .
(4)	The parties compute $\llbracket z \rrbracket_2 \leftarrow \llbracket c \rrbracket_2 \oplus d \cdot \llbracket e \rrbracket_2 \oplus e \cdot \llbracket a \rrbracket_2 \oplus d \cdot e$.
(5)	The parties output $\llbracket z \rrbracket_2$.

Figure 7: Subprotocols on authenticated bits

Since we are interested in reactive computations, in the opening garbling phase we do not reveal the wire mask λ_w associated to the circuit-output wires w , as happens in previous works. Our protocol $\Pi_{\text{Prep}}[\text{BMR}]$ securely implements an ideal functionality $\mathcal{F}_{\text{Prep}}[\text{BMR}]$ with active security in the $\mathcal{F}_{\text{n-TinyOT}}$ -hybrid model. We do not give here the details of $\mathcal{F}_{\text{Prep}}[\text{BMR}]$ since this functionality is essentially the same as the one given in [16]. Active security is obtained by using the actively secure $\mathcal{F}_{\text{n-TinyOT}}$ functionality, and the Open command of $\mathcal{F}_{\text{n-TinyOT}}$, which enables us to check the opened shares bits. The modifications of the proofs in [16] to the reactive case are minor, and thus we do not give them.

3.3 Complete Pre-processing Functionality

Other than the correlated randomness needed for the circuit evaluation, both in the LSSS and GC fashion, we need random doubly-authenticated bits (daBits) to perform the conversions in the on-line protocol. The command **daShare**, together with commands in $\mathcal{F}_{\text{Prep}}[\text{SPDZ}]$ and $\mathcal{F}_{\text{Prep}}[\text{BMR}]$, completes our ideal pre-processing functionality $\mathcal{F}_{\text{Prep}}$, given in Figure 8. Extending functionalities in the UC framework by combining and adding new commands is a well-known technique in the literature, e.g. in [20, 26].

In the next section we give a new procedure for generating daBits and show that it can be used to realize $\mathcal{F}_{\text{Prep}}$ UC securely, essentially by the same proof as in [26].

$\mathcal{F}_{\text{Prep}}$	
This functionality has all the commands of $\mathcal{F}_{\text{Prep}}[\text{SPDZ}]$ and all the commands of $\mathcal{F}_{\text{Prep}}[\text{BMR}]$ and additionally has the command daShare defined below.	
daShare:	On input $(\text{daShare}, \mathbb{F}_2, \text{varid}, \text{varid}_0, \dots, \text{varid}_{\ell-1}, b_1, \dots, b_{\ell})$, where $b_i \in \{0, 1\}$, $i \in [\ell]$, compute $r = \sum_{i=0}^{\ell} b_i \cdot 2^i \in \mathbb{F}_p$ and store $(\text{varid}, \mathbb{F}_p, r)$.

Figure 8: Ideal functionality for Preprocessing for Zaphod

4 DABITS

It was shown by Rotaru and Wood [26] that, if a subprotocol for generating daBits is given, in which each daBit is a sharing of the same uniformly-random bit in both fields, except with negligible probability, then there exists a protocol that UC-securely realizes the functionality $\mathcal{F}_{\text{Prep}}$. In this section, we give a new protocol for generating daBits to replace the procedure from [26]. The exposition focuses on the full-threshold case, since it is the most straightforward to explain. However, the protocol works for any Q2 access structure with some minor modifications, as described later.

The security of using the new protocol to realize our functionality $\mathcal{F}_{\text{Prep}}$ essentially follows immediately from the proof in [26]; however, there are two minor differences between our $\mathcal{F}_{\text{Prep}}$ and the corresponding functionality in [26], and so in the full version we also provide a sketch proof of the UC security. The differences are as follows. Firstly, we are only interested specifically in authenticated sharings of *field elements* in $\mathbb{F}_p$ whose bit-decomposition is also shared with authentication in $\mathbb{F}_2$ (not just bits). This is essentially a notational change since the conversion is entirely local, and so the functionalities are equivalent in the UC sense (i.e. there are UC reductions both ways). Secondly, Π_{daBits} refers explicitly to shares rather than abstract handles.

In Figure 9, we describe the protocol Π_{daShare} that produces double-authenticated shares $(\llbracket r \rrbracket_p, \llbracket r \rrbracket_2)$, where $\llbracket r \rrbracket_p \leftarrow \sum_{k=0}^{\ell-1} 2^k \cdot \llbracket b^{(k)} \rrbracket_p$ and $\llbracket r \rrbracket_2 = (\llbracket b^{(0)} \rrbracket_2, \dots, \llbracket b^{(\ell-1)} \rrbracket_2)$. This protocol makes use of our new Π_{daBits} protocol (Figure 10) to produce ℓ daBits $(\llbracket b^{(k)} \rrbracket_p, \llbracket b^{(k)} \rrbracket_2)_{k=0}^{\ell-1}$. The proof of the following theorem is deferred to the full version.

THEOREM 4.1. *The protocol Π_{daShare} UC-securely realizes $\mathcal{F}_{\text{Prep}}$ in the presence of a static, active adversary in the $\mathcal{F}_{\text{Prep}}[\text{AuthLSSS}], \mathcal{F}_{\text{Rand}}$ -hybrid model.*

Π_{daShare}	
daShare:	Parties do the following:
(1)	Execute Π_{daBits} to generate ℓ daBits, $(\llbracket b^{(k)} \rrbracket_p, \llbracket b^{(k)} \rrbracket_2)_{k=0}^{\ell-1}$.
(2)	Compute $\llbracket r \rrbracket_p \leftarrow \sum_{k=0}^{\ell-1} 2^k \cdot \llbracket b^{(k)} \rrbracket_p$ and set $\llbracket r \rrbracket_2 \leftarrow (\llbracket b^{(k)} \rrbracket_2)_{k=0}^{\ell-1}$.
(3)	Output $(\llbracket r \rrbracket_p, \llbracket r \rrbracket_2)$.

Figure 9: n-party daShare protocol

4.1 New daBit Protocol

The key observation leading to the new daBit protocol is that for a random bit shared in $\mathbb{F}_p$ between two parties, with probability $O(p^{-1})$ the XOR of the least significant bits (LSBs) of the two parties' shares is equal to the bit they share offset by $p \bmod 2$. Since typically $p = O(2^{\text{sec}})$, this probability is generally negligible in the security parameter sec . This is proved in Lemma 4.2 (Full version). In the protocol, this observation is used to turn $n/2$ random bits in $\mathbb{F}_p$, secret-shared among n parties, into one daBit shared in both $\mathbb{F}_p$ and in $\mathbb{F}_2$.

Before giving the full protocol for generating daBits amongst n parties, we first provide the intuition. We emphasize that the focus here is on the full-threshold case, with necessary adaptations in the case of a Q2 access structure summarized later.

First, in the **Input** step, the parties generate $n/2$ random bits in $\mathbb{F}_p$. In the **daBits Generation**, each of these is 'assigned' to one pair of adjacent parties as follows. For each pair, $(2 \cdot i - 1, 2 \cdot i)$, $i \in [n/2]$, the n -party sharing of one of the $n/2$ bits, $\llbracket b^{(i)} \rrbracket_p$, is turned into a 2-party sharing, which is achieved by each party other than $P_{2 \cdot i - 1}$ and $P_{2 \cdot i}$ sending their shares to one of these two parties. To decide which party sends to which in each pair, we define odd as the set of parties $\{P_1, P_3, \dots\}$ with odd index, and even as the set of parties $\{P_2, P_4, \dots\}$ with even index, and require parties always to open to the party with same parity as themselves. Thus for the adjacent pair $(P_{2 \cdot i - 1}, P_{2 \cdot i})$, they obtain a sharing $\widehat{b}_i^{(2 \cdot i - 1)} + \widehat{b}_i^{(2 \cdot i)} = b^{(i)} \bmod p$ of $b^{(i)}$. Now with high probability, it holds that $1 < \widehat{b}^{(2 \cdot i)}, \widehat{b}^{(2 \cdot i + 1)} < p$, so $(\widehat{b}_i^{(2 \cdot i)} \bmod 2) \oplus (\widehat{b}_i^{(2 \cdot i + 1)} \bmod 2) \oplus (p \bmod 2) = b_i \bmod 2$.

Now these two parties provide $\widehat{b}^{(2 \cdot i - 1)} \bmod 2$ and $\widehat{b}^{(2 \cdot i)} \bmod 2$ as inputs for $\mathcal{F}_{\text{Prep}}[\text{AuthLSSS}]$ with $\text{domain} = \mathbb{F}_2$, so that the parties obtain an n -party boolean sharing of these values. Then all the parties subtract (XOR) the public value $p \bmod 2$ from this sharing in $\mathbb{F}_2$. Finally, after doing this for each pair, the parties take the XOR of all $n/2$ bits in $\mathbb{F}_2$ and in $\mathbb{F}_p$, which requires $n/2 - 1$ triples in $\mathbb{F}_p$. In the protocol, instead of removing the offset $p \bmod 2$ once for each pair, they remove it after the XOR with the offset $(n/2 \cdot (p \bmod 2)) \bmod 2$ to obtain the final daBit.

To achieve active security, i.e. to make sure parties are honest when providing inputs to $\mathcal{F}_{\text{Prep}}[\text{AuthLSSS}]$, the parties perform the **daBits Check** step. Note that to produce m correct daBits securely except with probability $2^{-\text{sec}}$, the parties need to produce sec additional daBits that later will be discarded. This also means that the authentication of the 2-party bit is not necessarily correct when converting from an n -party sharing to a 2-party sharing since correctness is only guaranteed by the checking procedures.

As long as every party is involved in the generation of at least one daBit, any adversary corrupting up to $n - 1$ parties will not learn the value of the daBit, as shown in the full version of this paper. The protocol is given in Figure 10. It makes use of a subprotocol Π_{lsb} , of which we do not give a full description, that takes as input an authenticated secret-shared element of $\mathbb{F}_p$, $\llbracket x \rrbracket_p$ and returns the least significant bit of x in secret-shared form, $\llbracket b \rrbracket_p$. Protocols for this operation (with statistical security) can be found in the literature, e.g. [1, v1.5, §12.3.4]. It also uses a subprotocol Π_{xor} , which computes the (generalized) XOR function of a set of secret-shared bits in $\mathbb{F}_p$: given $\{\llbracket b^{(k)} \rrbracket_p\}_{k=1}^t$, it computes $\llbracket b \rrbracket_p \leftarrow \llbracket b^{(1)} \rrbracket_p$ and

then iteratively computes $\llbracket b \rrbracket_p \leftarrow \llbracket b \rrbracket_p + \llbracket b^{(i)} \rrbracket_p - 2 \cdot \llbracket b \rrbracket_p \cdot \llbracket b^{(i)} \rrbracket_p$ for all $i = 2$ to t , requiring $t - 1$ Beaver triples.

Figure 10: n -party daBit generation

In the proof of Lemma 4.2, which is left to the full version, we argue that the daBits protocol is correct except with negligible probability in the statistical security parameter.

LEMMA 4.2. *If the checks pass without parties aborting, then the protocol Π_{daBits} outputs correct daBits of uniformly random bits, except with probability $2^{-\text{sec}}$, provided $m > 2 \cdot \text{sec}$.*

General Q2 Protocol. In the case of an arbitrary Q2 access structure, when a bit is converted from an n -party sharing to a 2-party sharing, the pair must receive enough shares so that between them they hold an additive sharing of the secret, which typically requires less communication than in the full-threshold case. This is most easily done for a pair of parties (i, j) by taking a qualified set Q containing i and j and requiring all parties in Q to send their share(s) to one party or the other; then, using the recombination vector for Q , they can combine these shares into a two-party additive sharing.

For the secrecy of the final daBit (from the adversary, and the environment), it is not sufficient to generate $n/2$ bits and to turn them into 2-party sharings for pairs of *adjacent* parties as in the full-threshold case because it is possible that the adversary learns every bit during the conversion procedure. This happens (potentially) if in every pair, at least one party is corrupt and receives at least one share from an honest party.

Fortunately, there is an easy way to get around this problem by choosing the pairs so as to guarantee at least one pair exists for which both parties are honest. We give this algorithm in the full version, due to absence of space here.

5 ONLINE EVALUATION

Finally, we can give the online protocol Π_{MPC} (Figure 11) implementing the ideal functionality $\mathcal{F}_{\text{MPC}}$ given in Figure 2 and Figure 3.

Protocol Π_{MPC}
Initialize: Parties call $\mathcal{F}_{\text{Prep}}$ to activate the functionality.
Input, Evaluate, Output: There are two cases:
- If $\text{domain} = \mathbb{F}_p$, the parties run the sub-protocol $\Pi_{\text{MPC}}[\text{SPDZ}]$ on commands Input, Evaluate, Output
- Otherwise, the parties run the sub-protocol $\Pi_{\text{MPC}}[\text{BMR}]$ on commands Input, Evaluate, Output
Conversion: Parties call the sub-protocol $\Pi_{\text{ConvertToField}}$ and $\Pi_{\text{ConvertToBinary}}$

Figure 11: Zaphod online protocol

The sub-protocols $\Pi_{\text{MPC}}[\text{SPDZ}]$ (Figure 12) and $\Pi_{\text{MPC}}[\text{BMR}]$ (Figure 13) are the same as the SPDZ-online protocol [12] and BMR-online protocol [16], respectively. The only (minor) differences are:

- In $\Pi_{\text{MPC}}[\text{SPDZ}]$ we have the possibility to output to a single party, other than publicly to $\mathcal{P}$.
- As output of the BMR evaluation in $\Pi_{\text{MPC}}[\text{BMR}]$, the parties only obtain a sharing of the result of the computation. Note that this means that the final output of the online phase will always be a value in $\mathbb{F}_p$.

5.1 The Conversion Protocol

All that remains to complete the online protocol is to give the protocols for *ConvertToField* and *ConvertToBinary* and show that their implementation does not break either the security of the LSSS based MPC components or the Garbled circuit based MPC components. Given the functionality in Figure 8 our two conversion routines are trivial, see Figure 14 and Figure 15.

The most complex conversion is that in given in Figure 14 as it requires the evaluation of a garbled circuit. Here we utilize a simple circuit which takes as input the $\lceil \log_2 p \rceil$ bits of y , and the ℓ (secret shared) bits of r . The circuit computes $z = y + r \pmod{p}$, which can be done with a single addition and conditional subtraction. Then z is compared to $p/2$, if less than $p/2$ then the least 64 bits of z are extracted and assigned to x . Otherwise the circuit computes $w = p - z$, extracts the least 64 bits of w and then computes x the bitwise negation of these 64 bits plus one. As one can see this circuit involves only a small number of addition/subtraction operations, and is the main reason we restrict our selection of the prime p so that we do not have to perform more complex modular operations at this stage.

Sub-protocol $\Pi_{\text{MPC}}[\text{SPDZ}]$
Input: Given an input x entered by party P_i :
(1) The parties call the command (<i>Input</i> , P_i , r) of $\mathcal{F}_{\text{Prep}}[\text{SPDZ}]$, so that P_i obtains the <i>value</i> r , and the other parties obtain a handle on r , i.e. $\llbracket r \rrbracket_p$.
(2) Party P_i broadcasts the value $y \leftarrow x - r$.
(3) The parties compute $\llbracket x \rrbracket_p \leftarrow \llbracket r \rrbracket_p + y$.
Evaluate: - Add: On input $\llbracket x \rrbracket_p, \llbracket y \rrbracket_p$ the parties locally compute $\llbracket x + y \rrbracket_p \leftarrow \llbracket x \rrbracket_p + \llbracket y \rrbracket_p$.
- Multiply: On input $\llbracket x \rrbracket_p, \llbracket y \rrbracket_p$:
(1) The parties call <i>Triple</i> of $\mathcal{F}_{\text{Prep}}[\text{SPDZ}]$ to obtain a triple $(\llbracket a \rrbracket_p, \llbracket b \rrbracket_p, \llbracket c \rrbracket_p)$.
(2) The parties compute the handles $\llbracket e \rrbracket_p \leftarrow \llbracket x \rrbracket_p - \llbracket a \rrbracket_p$ and $\llbracket f \rrbracket_p \leftarrow \llbracket y \rrbracket_p - \llbracket b \rrbracket_p$.
(3) The parties call $\mathcal{F}_{\text{Prep}}[\text{SPDZ}].\text{Open}$ on $\llbracket e \rrbracket_p$ and $\llbracket f \rrbracket_p$.
(4) The parties evaluate $\llbracket z \rrbracket_p \leftarrow \llbracket c \rrbracket_p + e \cdot \llbracket y \rrbracket_p + f \cdot \llbracket a \rrbracket_p + e \cdot f$.
Output: To output a value $\llbracket y \rrbracket_p$ we distinguish two cases:
(1) (PublicOutput):
(a) Call $\mathcal{F}_{\text{Prep}}[\text{SPDZ}].\text{Check}$ on all opened handles/values since the last execution of $\mathcal{F}_{\text{Prep}}[\text{SPDZ}].\text{Check}$, if it aborts then abort.
(b) Call $\mathcal{F}_{\text{Prep}}[\text{SPDZ}].\text{Open}$ on the handle $\llbracket y \rrbracket_p$.
(c) Call $\mathcal{F}_{\text{Prep}}[\text{SPDZ}].\text{Check}$ on the single pair $(\llbracket y \rrbracket_p, y)$ if it aborts then abort, otherwise accept y as the valid output.
(2) (PrivateOutput):
(a) The parties execute (<i>Input</i> , P_i , r) so that P_i obtains the <i>value</i> r , and the other parties obtain a handle on r , i.e. $\llbracket r \rrbracket_p$.
(b) Call $\mathcal{F}_{\text{Prep}}[\text{SPDZ}].\text{Check}$ on all opened handles/values since the last execution of $\mathcal{F}_{\text{Prep}}[\text{SPDZ}].\text{Check}$, if it aborts then abort.
(c) The parties compute $\llbracket x \rrbracket_p \leftarrow \llbracket r \rrbracket_p + \llbracket y \rrbracket_p$.
(d) Call $\mathcal{F}_{\text{Prep}}[\text{SPDZ}].\text{Open}$ on the handle $\llbracket x \rrbracket_p$.
(e) Call $\mathcal{F}_{\text{Prep}}[\text{SPDZ}].\text{Check}$ on the single pair $(\llbracket x \rrbracket_p, x)$, if it aborts then abort. Otherwise, P_i computes and accepts $y \leftarrow x - r$.

Figure 12: The sub-protocol for realising evaluation over $\mathbb{F}_p$

Sub-protocol $\Pi_{\text{MPC}}[\text{BMR}]$
Input:
- The parties call $\mathcal{F}_{\text{Prep}}[\text{BMR}]$ on input C_f receiving all the wire masks $\llbracket \lambda_w \rrbracket_2$ and all the key $\mathbf{k}_{w,0}$.
- For each input wire w the parties have a shared value $\llbracket v_w \rrbracket_2$. The parties compute $\llbracket \epsilon_w \rrbracket_2 \leftarrow \llbracket v_w \rrbracket_2 \oplus \llbracket \lambda_w \rrbracket_2$. The value $\llbracket \epsilon_w \rrbracket_2$ is then opened to all parties.
Evaluate:
- For each input wire w P_i broadcasts the key $\mathbf{k}_{w,\epsilon_w}^i$.
- Each P_i now progresses through the circuit in the given topological ordering and locally processes the gates as follows:
(1) $c = \text{XOR}(a, b)$ Gate:
- $\epsilon_c \leftarrow \epsilon_a \oplus \epsilon_b$, and $\mathbf{k}_{c,\epsilon_c}^j \leftarrow \mathbf{k}_{a,\epsilon_a}^j \oplus \mathbf{k}_{b,\epsilon_b}^j$ for all $j \in [n]$.
(2) $c = \text{INV}(a)$ Gate:
- $\epsilon_c \leftarrow \epsilon_a \oplus 1$, and $\mathbf{k}_{c,\epsilon_c}^j \leftarrow \mathbf{k}_{a,\epsilon_a}^j \oplus \mathbf{k}_{\text{one},1}^j$ for all $j \in [n]$.
(3) $c = \text{AND}(a, b)$ Gate:
- Recover $\mathbf{k}_{c,\epsilon_c}^j$ for all $j \in [n]$ by computing
$\mathbf{k}_{c,\epsilon_c}^j \leftarrow g_{j,\epsilon_a,\epsilon_b} \oplus \left(\bigoplus_{j=1}^n F_{\mathbf{k}_{a,\epsilon_a}^j, \mathbf{k}_{b,\epsilon_b}^j} (g \parallel j) \right)$
- If $\mathbf{k}_{c,\epsilon_c}^i = \mathbf{k}_{c,0}^i$ set $\epsilon_c = 0$, if $\mathbf{k}_{c,\epsilon_c}^i = \mathbf{k}_{c,1}^i$ set $\epsilon_c = 1$, else <i>abort</i> .
Output: For every output wire w the parties compute $\llbracket v_w \rrbracket_2 \leftarrow \llbracket \lambda_w \rrbracket_2 \oplus \epsilon_w$, obtaining a (privately) share of $\llbracket v_w \rrbracket_2$.

Figure 13: The BMR-evaluation sub-protocol

Figure 14: Protocol to convert a $\llbracket \cdot \rrbracket_p$ -sharing into a vector of 64 $\llbracket \cdot \rrbracket_2$ -sharings

Figure 15: Protocol to convert a vector of $\llbracket \cdot \rrbracket_2$ -sharings into a $\llbracket \cdot \rrbracket_p$ -sharing

THEOREM 5.1. *The protocol Π_{MPC} securely implements the functionality $\mathcal{F}_{\text{MPC}}$ in the presence of a static active adversary corrupting up to $n - 1$ parties in the $\mathcal{F}_{\text{Prep}}$ -hybrid model.*

Correctness follows by inspection of the protocol. At a high level, the simulation of the *Input*, *Evaluation*, *Output* steps is almost identical to the respective simulations of $\Pi_{\text{MPC}}[\text{SPDZ}]$ and $\Pi_{\text{MPC}}[\text{BMR}]$ in [12] and [16], respectively. In particular, the input step is done by using random masks from the pre-processing phase, that can be used by the simulator to extract corrupt parties' inputs. Both protocols the conversion protocols only consists of calling to the $\mathcal{F}_{\text{Prep}}$ functionality and openings, so the simulation is straightforward. Indistinguishability follows from standard arguments and our choice of p . A more detailed proof is given in the full version.

6 ENGINEERING

In this section we outline a number of facets about how computing using both $\llbracket \cdot \rrbracket_p$ and $\llbracket \cdot \rrbracket_2$ sharings were integrated together into the SCALE-MAMBA system. This is a system which implemented LSSS-based MPC, i.e. MPC based on $\langle \cdot \rangle_p$ sharings. To perform this integration a number of design decisions had to be made.

6.1 Threads and Data Queues

The SCALE-MAMBA system allows a number of independent MPC programs to run in parallel, thus creating a number of independent online 'threads'. Since pre-processing for the $\langle \cdot \rangle_p$ -sharing based

MPC is more expensive than the online processing, each online thread has a number of 'helper' threads which provide the pre-processed data. Thus for each thread implementing $\Pi_{\text{MPC}}[\text{SPDZ}]$ there are a four threads providing the functionalities of $\mathcal{F}_{\text{Prep}}[\text{SPDZ}]$ (for example Triples and Random Shared Bits). Each of the data types produced by the offline threads are added to queues, for consumption by the associated online thread. Keeping one queue per data type, per online thread, means there is no need for complex signaling between threads as would be needed if there was a single global queue for each data type.

In addition, there are a variable number of global threads (in the case of full threshold access structures) which perform the necessary ZKPoKs associated to the FHE-based ciphertexts (essentially these threads implement the HighGear/TopGear ZKPoKs [2, 21]. In the default configuration there are two such ZKPoK threads. Thus if one is running o online threads then, in the default configuration for full threshold, one is actually running $4 \cdot o + 2$ actual threads.

This means that adding the garbled circuit functionality needs to be done in a way which does not degrade the performance of the existing $\langle \cdot \rangle_p$ -based MPC routines too much. We thus settled for the following architecture. We added a single additional global thread (independent of the number of online threads), which performs the base OTs and OT-extensions, enabling the implementation of the procedure Get from $\mathcal{F}_{\text{n-TinyOT}}$. This thread maintains a list of $o + 1$ queues of authenticated bits $\llbracket b \rrbracket_2$, for random bits b . The throughput of such a single thread turns out to be more than enough to support the online operations in up to four online threads, since OT extension is so fast. The first o of the $o + 1$ queues of authenticated bits are used to supply the online threads with authenticated bits, whereas the last of the authenticated bits queues is used to supply another thread whose task is to produce o queues of triples, one queue per online thread.

As a garbled circuit operation is encountered, the online thread polls the authenticated bit and triple thread for data; which is then supplied (if it is ready). When it is not yet ready (which happens if another GC operation has been performed recently, or if this is the first GC operation to be performed in a run) the online thread waits for the data to become available. When a garbled circuit needs to be evaluated, the garbling and evaluation are done on the fly by the online thread, resulting in authenticated bits and triples being consumed from the respective queues. We thus do not treat garbling as an "offline" operation.

The creation of daBits is performed within the online thread as and when needed. When a first request for daBits is encountered, the online thread produces a batch of daBits, consumes what it wants, and then stores any remaining for future use. Subsequent daBit requests either use the stored unused daBits, or if none are left more daBits are generated. When daBits are needed they are also consumed from the respective list.

Thus with this architecture we increase the number of actual threads from $4 \cdot o + 2$ to $4 \cdot o + 4$. By limiting the authenticated bit and triple producing threads to only produce a limited number of authenticated bits and triples, we can mitigate any slow-down of the online threads when no garbled circuit based operations are performed. However, we do get a small stall in performance when the first garbled circuit (resp. daBit) operation is met, and the associated queues of triples (resp. daBits) need to be filled.

6.2 Circuit Generation

By selecting to operate online on $\ell = 64$ -bit data types, it means that we restrict the user to a small set of well defined operations on the 64-bit signed integers. Thus the user is not able to call arbitrary circuits, but has access to a library of circuits. This design decision is because most users are unable to design their own circuits, with few having either VHDL experience or access to expensive state-of-the-art tool chains. Thus to operate on 64-bit data types we extended the MAMBA language to support secure signed 64-bit arithmetic, with each arithmetic operation corresponding to a pre-defined combinatorial circuit. The combinatorial circuits were generated using the Synopsis tool chain from a VHDL behavioural description of the desired circuit functionality.

However, a problem arises when using this methodology for the circuit needed in the $\mathbb{F}_2^\ell$ online evaluation. Users in SCALE-MAMBA can select their own prime p for the secret sharing scheme $\langle \cdot \rangle_p$, thus the circuit will depend not only on the bits of p , but also on the bit length. Thus it would seem that users need to be able to ‘compile’ their own circuit for such an operation, and thus would need access to tool chains such as that from Synopsis.

To avoid this issue we created a global circuit which performs the required operation for an arbitrary p (of bit length up to 512 bits). Then in a one-time setup procedure, this generic circuit, consisting of 6185 AND gates, 2893 XOR gates and 5865 INV gates, is specialized via a C-program to the specific prime selected by the user. This latter specialisation is then optimized further in C. The resulting circuit depends on the prime p being used, but for example in a 128-bit p variant the final optimized specialisation resulted in a circuit with 968 AND gates, 523 XOR gates and 629 INV gates.

We also extended SCALE-MAMBA to allow user-defined circuits to be executed. Thus, a user can define (say) a circuit for the Keccak round function and then evaluate it on 25 secure 64-bit integers (Keccak has a “block” size of 1600 bits). Combined with our daBit conversion routine this would allow the implementation of the distributed decryption algorithm in [23] completely within SCALE-MAMBA for an arbitrary number of parties. This algorithm requires switching between secret sharing modulo p , and then the resulting shared values need to be passed various garbled circuit constructions, and then passed back to secret sharing mod p .

6.3 Experiments

We benchmarked the daBit generation between computers with commodity hardware connected via a 10Gb/s LAN network with an average round-trip latency of 0.5ms. The results are presented in Table 1 with various statistical security parameters sec . According to the SCALE-MAMBA architecture, these numbers are produced using five FHE factories and eight online daBit generation threads.

The communication reported in [26] for statistical security $\text{sec} = 40$ in the two party case is 171kb which is about twice as low as the equivalent implementation in SCALE-MAMBA (384kb). The reason for this is that Rotaru and Wood estimated the pre-preprocessing per daBit using LowGear with every possible security parameter equal to 40 (zero knowledge proof soundness, distributed decryption, etc). The communication cost per triple used by them was 26.3kb. On the other hand we ran the daBit generation in parallel with TopGear which has a slightly higher communication cost per triple (78kb) but

also enjoys higher security parameters for the ZK proof soundness (2^{-128}) and distributed decryption (2^{-40}). SCALE-MAMBA already had the method from [26] implemented within it, thus making an apples-to-apples comparison with our method was easier.

For the two party case we see that communication cost are reduced by a quarter, with the throughput doubled, using the new method for daBit generation. In the three party case the reduction in communication is less pronounced, and the throughput is roughly the same. However for an odd number of parties n the new method has roughly the same complexity as the case for the $n + 1$ parties. Thus when we move to four parties the improvement in the new method in terms of communication again returns. At 80-bit security we found it impossible to run the method from [26] for three parties, as the limits on the queues of pre-processed data in SCALE-MAMBA were too small for what is required by the method.

We have also benchmarked the total conversion cost of a 64 bit integer with statistical security $\text{sec} = 40$ for the two party case. As expected, the conversion throughput from LSSS to GC is 104 times lower than generating a single daBit ending in 2.12MBytes per conversion (total communication) with a throughput of 20 conversions per second. The reason for a factor of 100 slowdown compared to the daBit generation is that a single conversion requires 104 daBits.

ACKNOWLEDGMENTS

This work has been supported in part by an ERC Advanced Grant ERC-2015-AdG-IMPACT, by the Defense Advanced Research Projects Agency (DARPA) and Space and Naval Warfare Systems Center, Pacific (SSC Pacific) under contract No. N66001-15-C-4070 and No. FA8750-19-C-0502, by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA) via Contract No. 2019-1902070006, and by the FWO under Odysseus project GOH9718N. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the any of the funders. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

REFERENCES

- [1] Abdelrahman Aly, Marcel Keller, Emmanuela Orsini, Dragos Rotaru, Peter Scholl, Nigel P. Smart, and Tim Wood. 2018. SCALE-MAMBA v1.2: Documentation. <https://homes.esat.kuleuven.be/~nsmart/SCALE/Documentation.pdf>
- [2] Carsten Baum, Daniele Cozzo, and Nigel P. Smart. 2019. Using TopGear in Overdrive: A more efficient ZKPoK for SPDZ. *IACR Cryptology ePrint Archive* 2019 (2019), 35. <https://eprint.iacr.org/2019/035>
- [3] Donald Beaver, Silvio Micali, and Phillip Rogaway. 1990. The Round Complexity of Secure Protocols (Extended Abstract). In *22nd Annual ACM Symposium on Theory of Computing*. ACM Press, Baltimore, MD, USA, 503–513. <https://doi.org/10.1145/100216.100287>
- [4] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. 1988. Completeness Theorems for Non-Cryptographic Fault-Tolerant Distributed Computation (Extended Abstract). In *20th Annual ACM Symposium on Theory of Computing*. ACM Press, Chicago, IL, USA, 1–10. <https://doi.org/10.1145/62212.62213>
- [5] Rikke Bendlin, Ivan Damgård, Claudio Orlandi, and Sarah Zakarias. 2011. Semi-homomorphic Encryption and Multiparty Computation. In *Advances in Cryptology – EUROCRYPT 2011 (Lecture Notes in Computer Science)*, Kenneth G. Paterson (Ed.), Vol. 6632. Springer, Heidelberg, Germany, Tallinn, Estonia, 169–188. https://doi.org/10.1007/978-3-642-20465-4_11
- [6] Justin Brickell, Donald E. Porter, Vitaly Shmatikov, and Emmett Witchel. 2007. Privacy-preserving remote diagnostics. In *ACM CCS 2007: 14th Conference on Computer and Communications Security*, Peng Ning, Sabrina De Capitani di Vimercati, and Paul F. Syverson (Eds.). ACM Press, Alexandria, Virginia, USA, 498–507. <https://doi.org/10.1145/1315245.1315307>

Protocol	n	sec	Comm. (kb)	Throughput (ops/s)
daBit [26]	2	40	384	1008
dabit (ours)	2	40	94	2150
daBit [26]	3	40	1640	560
dabit (ours)	3	40	1104	650
daBit [26]	4	40	4781	306
dabit (ours)	4	40	2173	552
daBit [26]	2	80	931	424
daBit (ours)	2	80	288	706
daBit [26]	3	80	-	-
daBit (ours)	3	80	2442	285

Table 1: 10Gb/s LAN experiments for n -party daBit generation

- [7] Niklas Büscher, Daniel Demmler, Stefan Katzenbeisser, David Kretzmer, and Thomas Schneider. 2018. HyCC: Compilation of Hybrid Protocols for Practical Secure Computation. In *ACM CCS 2018: 25th Conference on Computer and Communications Security*, David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang (Eds.). ACM Press, Toronto, ON, Canada, 847–861. <https://doi.org/10.1145/3243734.3243786>
- [8] Ran Canetti. 2000. Universally Composable Security: A New Paradigm for Cryptographic Protocols. Cryptology ePrint Archive, Report 2000/067. <http://eprint.iacr.org/2000/067>.
- [9] Nishanth Chandran, Divya Gupta, Aseem Rastogi, Rahul Sharma, and Shardul Tripathi. 2017. EzPC: Programmable, Efficient, and Scalable Secure Two-Party Computation for Machine Learning. Cryptology ePrint Archive, Report 2017/1109. <https://eprint.iacr.org/2017/1109>.
- [10] David Chaum, Claude Crépeau, and Ivan Damgård. 1988. Multiparty Unconditionally Secure Protocols (Extended Abstract). In *20th Annual ACM Symposium on Theory of Computing*. ACM Press, Chicago, IL, USA, 11–19. <https://doi.org/10.1145/62212.62214>
- [11] Koji Chida, Daniel Genkin, Koki Hamada, Dai Ikarashi, Ryo Kikuchi, Yehuda Lindell, and Ariel Nof. 2018. Fast Large-Scale Honest-Majority MPC for Malicious Adversaries. Cryptology ePrint Archive, Report 2018/570. <https://eprint.iacr.org/2018/570>.
- [12] Ivan Damgård, Valerio Pastro, Nigel P. Smart, and Sarah Zakarias. 2012. Multiparty Computation from Somewhat Homomorphic Encryption. In *Advances in Cryptology – CRYPTO 2012 (Lecture Notes in Computer Science)*, Reihaneh Safavi-Naini and Ran Canetti (Eds.), Vol. 7417. Springer, Heidelberg, Germany, Santa Barbara, CA, USA, 643–662. https://doi.org/10.1007/978-3-642-32009-5_38
- [13] Daniel Demmler, Thomas Schneider, and Michael Zohner. 2015. ABY - A Framework for Efficient Mixed-Protocol Secure Two-Party Computation. In *ISOC Network and Distributed System Security Symposium – NDSS 2015*. The Internet Society, San Diego, CA, USA.
- [14] Tore Kasper Frederiksen, Marcel Keller, Emmanuela Orsini, and Peter Scholl. 2015. A Unified Approach to MPC with Preprocessing Using OT. In *Advances in Cryptology – ASIACRYPT 2015, Part I (Lecture Notes in Computer Science)*, Tetsu Iwata and Jung Hee Cheon (Eds.), Vol. 9452. Springer, Heidelberg, Germany, Auckland, New Zealand, 711–735. https://doi.org/10.1007/978-3-662-48797-6_29
- [15] Chun Guo, Jonathan Katz, Xiao Wang, and Yu Yu. 2019. Efficient and Secure Multiparty Computation from Fixed-Key Block Ciphers. *IACR Cryptology ePrint Archive* 2019 (2019), 74. <https://eprint.iacr.org/2019/074>
- [16] Carmit Hazay, Peter Scholl, and Eduardo Soria-Vazquez. 2017. Low Cost Constant Round MPC Combining BMR and Oblivious Transfer. In *Advances in Cryptology – ASIACRYPT 2017, Part I (Lecture Notes in Computer Science)*, Tsuyoshi Takagi and Thomas Peyrin (Eds.), Vol. 10624. Springer, Heidelberg, Germany, Hong Kong, China, 598–628. https://doi.org/10.1007/978-3-319-70694-8_21
- [17] Wilko Henecka, Stefan Kögl, Ahmad-Reza Sadeghi, Thomas Schneider, and Immo Wehrenberg. 2010. TASTY: tool for automating secure two-party computations. In *ACM CCS 2010: 17th Conference on Computer and Communications Security*, Ehab Al-Shaer, Angelos D. Keromytis, and Vitaly Shmatikov (Eds.). ACM Press, Chicago, Illinois, USA, 451–462. <https://doi.org/10.1145/1866307.1866358>
- [18] Mauricio Karchmer and Avi Wigderson. 1993. On Span Programs. In *Proceedings of Structures in Complexity Theory*. 102–111.
- [19] Marcel Keller, Emmanuela Orsini, and Peter Scholl. 2015. Actively Secure OT Extension with Optimal Overhead. In *Advances in Cryptology – CRYPTO 2015, Part I (Lecture Notes in Computer Science)*, Rosario Gennaro and Matthew J. B. Robshaw (Eds.), Vol. 9215. Springer, Heidelberg, Germany, Santa Barbara, CA, USA, 724–741. https://doi.org/10.1007/978-3-662-47989-6_35
- [20] Marcel Keller, Emmanuela Orsini, and Peter Scholl. 2016. MAScot: Faster Malicious Arithmetic Secure Computation with Oblivious Transfer. In *ACM CCS 2016: 23rd Conference on Computer and Communications Security*, Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi (Eds.). ACM Press, Vienna, Austria, 830–842. <https://doi.org/10.1145/2976749.2978357>
- [21] Marcel Keller, Valerio Pastro, and Dragos Rotaru. 2018. Overdrive: Making SPDZ Great Again. In *Advances in Cryptology – EUROCRYPT 2018, Part III (Lecture Notes in Computer Science)*, Jesper Buus Nielsen and Vincent Rijmen (Eds.), Vol. 10822. Springer, Heidelberg, Germany, Tel Aviv, Israel, 158–189. https://doi.org/10.1007/978-3-319-78372-7_6
- [22] Florian Kerschbaum, Thomas Schneider, and Axel Schröpfer. 2014. Automatic Protocol Selection in Secure Two-Party Computations. In *ACNS 14: 12th International Conference on Applied Cryptography and Network Security (Lecture Notes in Computer Science)*, Ioana Boureanu, Philippe Owesarski, and Serge Vaudenay (Eds.), Vol. 8479. Springer, Heidelberg, Germany, Lausanne, Switzerland, 566–584. https://doi.org/10.1007/978-3-319-07536-5_33
- [23] Michael Kraitsberg, Yehuda Lindell, Valery Osheter, Nigel P. Smart, and Younes Talibi Alaoui. 2019. Adding Distributed Decryption and Key Generation to a Ring-LWE Based CCA Encryption Scheme. In *ACISP 19: 24th Australasian Conference on Information Security and Privacy (Lecture Notes in Computer Science)*, Julian Jang-Jaccard and Fuchun Guo (Eds.), Vol. 11547. Springer, Heidelberg, Germany, Christchurch, New Zealand, 192–210. https://doi.org/10.1007/978-3-030-21548-4_11
- [24] Payman Mohassel and Peter Rindal. 2018. ABY3: A Mixed Protocol Framework for Machine Learning. Cryptology ePrint Archive, Report 2018/403. <https://eprint.iacr.org/2018/403>.
- [25] Jesper Buus Nielsen, Peter Sebastian Nordholt, Claudio Orlandi, and Sai Sheshank Burra. 2012. A New Approach to Practical Active-Secure Two-Party Computation. In *Advances in Cryptology – CRYPTO 2012 (Lecture Notes in Computer Science)*, Reihaneh Safavi-Naini and Ran Canetti (Eds.), Vol. 7417. Springer, Heidelberg, Germany, Santa Barbara, CA, USA, 681–700. https://doi.org/10.1007/978-3-642-32009-5_40
- [26] Dragos Rotaru and Tim Wood. 2019. MArBled Circuits: Mixing Arithmetic and Boolean Circuits with Active Security. *IACR Cryptology ePrint Archive* 2019 (2019), 207. <https://eprint.iacr.org/2019/207>
- [27] Nigel P. Smart and Tim Wood. 2019. Error Detection in Monotone Span Programs with Application to Communication-Efficient Multi-party Computation. In *Topics in Cryptology – CT-RSA 2019 (Lecture Notes in Computer Science)*, Mitsuru Matsui (Ed.), Vol. 11405. Springer, Heidelberg, Germany, San Francisco, CA, USA, 210–229. https://doi.org/10.1007/978-3-030-12612-4_11
- [28] Xiao Wang, Samuel Ranellucci, and Jonathan Katz. 2017. Global-Scale Secure Multiparty Computation. In *ACM CCS 2017: 24th Conference on Computer and Communications Security*, Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu (Eds.). ACM Press, Dallas, TX, USA, 39–56. <https://doi.org/10.1145/3133956.3133979>
- [29] Andrew Chi-Chih Yao. 1982. Protocols for Secure Computations (Extended Abstract). In *23rd Annual Symposium on Foundations of Computer Science*. IEEE Computer Society Press, Chicago, Illinois, 160–164. <https://doi.org/10.1109/SFCS.1982.38>