Improved Primitives for MPC over Mixed Arithmetic-Binary Circuits

Daniel Escudero[1(✉)], Satrajit Ghosh[1(✉)], Marcel Keller[2(✉)],
Rahul Rachuri[1(✉)], and Peter Scholl[1(✉)]

[1] Aarhus University, Aarhus, Denmark
{escudero,satrajit,rachuri,peter.scholl}@cs.au.dk
[2] CSIRO's Data61, Sydney, Australia
mks.keller@gmail.com

Abstract. This work introduces novel techniques to improve the translation between arithmetic and binary data types in secure multi-party computation. We introduce a new approach to performing these conversions using what we call *extended doubly-authenticated bits* (edaBits), which correspond to shared integers in the arithmetic domain whose bit decomposition is shared in the binary domain. These can be used to considerably increase the efficiency of non-linear operations such as truncation, secure comparison and bit-decomposition.

Our edaBits are similar to the *daBits* technique introduced by Rotaru et al. (Indocrypt 2019). However, we show that edaBits can be directly produced much more efficiently than daBits, with active security, while enabling the same benefits in higher-level applications. Our method for generating edaBits involves a novel cut-and-choose technique that may be of independent interest, and improves efficiency by exploiting natural, tamper-resilient properties of binary circuits that occur in our construction. We also show how edaBits can be applied to efficiently implement various non-linear protocols of interest, and we thoroughly analyze their correctness for both signed and unsigned integers.

The results of this work can be applied to any corruption threshold, although they seem best suited to dishonest majority protocols such as SPDZ. We implement and benchmark our constructions, and experimentally verify that our technique yields a substantial increase in efficiency. EdaBits save in communication by a factor that lies between 2 and 60 for secure comparisons with respect to a purely arithmetic approach, and between 2 and 25 with respect to using daBits. Improvements in throughput per second slightly lower but still as high as a factor of 47. We also apply our novel machinery to the tasks of biometric matching and convolutional neural networks, obtaining a noticeable improvement as well.

1 Introduction

Secure multi-party computation, or MPC, allows a set of parties to compute some function f on private data, in such a way that the parties do not learn anything about the actual inputs to f, beyond what could be computed given

the result. MPC can be used in a wide range of applications, such as private statistical analysis, machine learning, secure auctions and more.

MPC protocols can vary widely depending on the adversary model that is considered. For example, protocols in the *honest majority* setting are only secure as long as fewer than half of the parties are corrupt and colluding, whilst protocols secure against a *dishonest majority* allow all-but-one of the parties to be corrupt. Another important distinction is whether the adversary is assumed to be *semi-honest*, that is, they will always follow the instructions of the protocol, or *malicious*, and can deviate arbitrarily.

The mathematical structure underpinning secure computation usually requires to fix what we call a computation domain. The most common examples of such domains are computation modulo a large number (prime or power of two) or binary circuits (computation modulo two). In terms of cost, the former is more favorable to integer computation such as addition and multiplication while the latter is preferable for highly non-linear functions such as comparisons.

Applications often feature both linear and non-linear functionality. For example, convolution layers in deep learning consist of dot products followed by a nonlinear activation function. It is therefore desirable to convert between an arithmetic computation domain and binary circuits. This has led to a line of works exploring this possibility, starting with the ABY framework [20] (Arithmetic-Boolean-Yao) in the two-party setting with semi-honest security. Other works have extended this to the setting of three parties with an honest majority [2,30], dishonest majority with malicious security [33], as well as creating compilers that automatically decide which parts of a program should done in the binary or arithmetic domain [9,12,28].

A particular technique that is relevant for us is so-called *daBits* [33] (doubly-authenticated bits), which are random secret bits that are generated simultaneously in both the arithmetic and binary domains. These can be used for binary/arithmetic conversions in MPC protocols with any corruption setting, and have in particular been used with the SPDZ protocol [19], which provides malicious security in the dishonest majority setting. Later works have given more efficient ways of generating daBits [1,3,32], both with SPDZ and in the honest majority setting.

Another recent work uses function secret sharing [6] for binary/arithmetic conversions and other operations such as comparison [7]. This approach leads to a fast online phase with just one round of interaction and optimal communication complexity. However, it requires either a trusted setup, or an expensive preprocessing phase which has not been shown to be practical for malicious adversaries.

Limitations of daBits. Using daBits, it is relatively straightforward to convert between two computation domains. However, we found that in application-oriented settings the benefit of daBits alone is relatively limited. More concretely, if daBits are used to compute a comparison between two numbers that are secret-shared in $\mathbb{Z}_M$, for large arithmetic modulus M, the improvement is a factor of three at best. The reason for this is that the cost of creating the required daBits comes quite close to computing the comparison entirely in $\mathbb{Z}_M$.

This limitation seems to be inherent with any approach based on daBits, since a daBit requires generating a random shared bit in $\mathbb{Z}_M$. The only known way of doing this with malicious security require first performing a multiplication (or squaring) in $\mathbb{Z}_M$ on a secret value [16,17]. However, secret multiplication is an expensive operation in MPC, and doing this for every daBit gets costly.

1.1   Our Contributions

In this paper, we present a new approach to converting between binary and arithmetic representations in MPC. Our method is general, and can be applied to a wide range of corruption settings, but seems particularly well-suited to the case of dishonest majority with malicious security such as SPDZ [18,19], over the arithmetic domain $\mathbb{Z}_p$ for large prime p, or the ring $\mathbb{Z}_{2^k}$ [13]. Unlike previous works, we do not generate daBits, but instead create what we call *extended daBits* (edaBits), which avoid the limitations above. These allow conversions between arithmetic and binary domains, but can also be used directly for certain non-linear functions such as truncations and comparisons. We found that, for two- and three-party computation, edaBits allow to reduce the communication cost by up to two orders of magnitude and the wall clock time by up to a factor of 50 while both the inputs as well as the output are secret-shared in an arithmetic domain.

Below we highlight some more details of our contribution.

Extended daBits. An edaBit consists of a set of m random bits $(r_{m-1}, \ldots, r_0)$, secret-shared in the binary domain, together with the value $r = \sum_{i=0}^{m-1} r_i 2^i$ shared in the arithmetic domain. We denote these sharings by $[r_{m-1}]_2, \ldots, [r_0]_2$ and $[r]_M$, for arithmetic modulus M. Note that a daBit is simply an edaBit of length $m = 1$, and m daBits can be easily converted into an edaBit with a linear combination of the arithmetic shares. We show that this is wasteful, however, and edaBits can in general be produced much more efficiently than m daBits, for values of m used in practice.

Efficient malicious generation of edaBits. Let us first consider a simple approach with semi-honest security. If there are n parties, we have each party locally sample a value $r^i \in \mathbb{Z}_M$, then secret-shares r^i in the arithmetic domain, and the bits of r^i in the binary domain. We refer to these sharings as a *private* edaBit known to P_i. The parties can combine these by computing $\sum_i r^i$ in the arithmetic domain, and executing $n - 1$ protocols for addition in the binary domain, with a cost $O(nm)$ AND gates. Compared with using daBits, which costs $O(m)$ secret multiplications in $\mathbb{Z}_M$, this is much cheaper if n is not too large, by the simple fact that AND is a cheaper operation than multiplication in MPC.

To extend this naive approach to the malicious setting, we need a way to somehow verify that a set of edaBits was generated correctly. Firstly, we extend the underlying secret-sharing scheme to one that enforces correct computations on the underlying shares. This can be done, for instance, using authenticated

secret-sharing with information-theoretic MACs as in SPDZ [19]. Secondly, we use a cut-and-choose procedure to check that a large batch of edaBits are correct. This method is inspired by previous techniques for checking multiplication triples in MPC [8,22,23]. However, the case of edaBits is much more challenging to do efficiently, due to the highly non-linear relation between sharings in different domains, compared with the simple multiplicative property of triples (shares of (a, b, c) where $c = ab$).

Cut-and-choose approach. Our cut-and-choose procedure begins as in the semi-honest case, with each party P_i sampling and inputting a large batch of private edaBits of the form $(r_{m-1}^i, \ldots, r_0^i), r^i$. We then run a verification step on P_i's private edaBits, which begins by randomly picking a small subset of the edaBits to be opened and checked for correctness. Then, the remaining edaBits are shuffled and put into buckets of fixed size B. The first edaBit in each bucket is paired off with every other edaBit in the bucket, and we run a checking procedure on each of these pairs. To check a pair of edaBits r, s, the parties can compute $r + s$ in both the arithmetic and binary domains, and check these open to the same value. If all checks pass, then the parties take the first private edaBit from every bucket, and add this to all the other parties' private edaBits, created in the same way, to obtain secret-shared edaBits. Note that to pass a single check, the adversary must have corrupted both r and s so that they cancel each other out; therefore, the only way to successfully cheat is if every bucket with a corrupted edaBit contains *only* corrupted edaBits. By carefully choosing parameters, we can ensure that it is very unlikely the adversary manages to do this. For example, with 40-bit statistical security, from the analysis of [23], we could use bucket size $B = 3$ when generating more than a million sets of edaBits.

While the above method works, it incurs considerable overhead compared with similar cut-and-choose techniques used for multiplication triples. This is because in every pairwise check within a bucket, the parties have to perform an addition of binary-shared values, which requires a circuit with $O(m)$ AND gates. Each of these AND gates consumes an authenticated multiplication triple over $\mathbb{Z}_2$, and generating these triples themselves requires additional layers of cut-and-choose and verification machinery, when using efficient protocols based on oblivious transfer [22,31,34].

To reduce this cost, our first optimization is as follows. Recall that the check procedure within each bucket is done on a pair of *private* values known to one party, and not secret-shares. This means that when evaluating the addition circuit, it suffices to use *private* multiplication triples, which are authenticated triples where the secret values are known to party P_i. These are much cheaper to generate than fully-fledged secret-shared triples, although still require a verification procedure based on cut-and-choose. To further reduce costs, we propose a second, more significant optimization.

Cut-and-choose with faulty check circuits. Instead of using private triples that have been checked separately, we propose to use *faulty private triples*, that is, authenticated triples that are not guaranteed to be correct. This immediately

raises the question, how can the checking procedure be useful, if the verification mechanism itself is faulty? The hope is that if we randomly shuffle the set of triples, it may still be hard for an adversary who corrupts them to ensure that any incorrect edaBits are canceled out in the right way by the faulty check circuit, whilst any correct edaBits still pass unscathed. Proving this, however, is challenging. In fact, it seems to inherently rely on the *structure* of the binary circuit that computes the check function. For instance, if a faulty circuit can cause a check between a good and a bad edaBit to pass, and the same circuit also causes a check between two good edaBits to pass, for some carefully chosen inputs, then this type of cheating can help the adversary.

To rule this out, we consider circuits with a property we call *weak additive tamper-resilience*, meaning that for any tampering that flips some subset of AND gate outputs, the tampered circuit is either incorrect for every possible input, or it is correct for all inputs. This notion essentially rules out input-dependent failures from faulty multiplication triples, which avoids the above attack and allows us to simplify the analysis.

Weak additive tamper-resilience is implied by previous notions of circuits secure against additive attacks [24], however, these constructions are not practical over $\mathbb{F}_2$. Fortunately, we show that the standard ripple-carry adder circuit satisfies our notion, and suffices for creating edaBits in $\mathbb{Z}_{2^k}$. However, the circuit for binary addition modulo a prime, which requires an extra conditional subtraction, does not satisfy this. Instead, we adapt the circuit over the integers to use in our protocol modulo p, which allows us to generate length-m edaBits for any $m < \log p$; this turns out to be sufficient for most applications.

With this property, we can show that introducing faulty triples does not help an adversary to pass the check, so we can choose the same cut-and-choose parameters as previous works on triple generation, while saving significantly in the cost of generating our triples used in verification. The bulk of our technical contribution is in analysing this cut-and-choose technique.

Silent OT-friendly. Another benefit of our approach is that we can take advantage of recent advances in oblivious transfer (OT) extension techniques, which allow to create a large number of random, or correlated, OTs, with very little interaction [5]. In practice, the communication cost when using this "silent OT" method can be more than 100x less than OT extension based on previous techniques [27], with a modest increase in computation [4]. In settings where bandwidth is expensive, this suits our protocol well, since we mainly use MPC operations in $\mathbb{F}_2$ to create edaBits, and these are best done with OT-based techniques. This reduces the communication of our edaBits protocol by an $O(\lambda)$ factor, in practice cutting communication by 50–100x, although we have not yet implemented this optimization.

Note that it does not seem possible to exploit silent OT with previous daBit generation methods such as by Aly et al. [1]. This is due to the limitation mentioned previously that these require a large number of random bits shared in $\mathbb{Z}_p$, which we do not know how to create efficiently using OT.

Applications: improved conversions and primitives. edaBits can be used in a natural way to convert between binary and arithmetic domains, where each conversion of an m-bit value uses one edaBit of length m, and a single m-bit addition circuit. (In the mod-p case, we also need one "classic" daBit per conversion, to handle a carry computation.) However, for many primitives such as secure comparison, equality test and truncation, a better approach is to exploit the edaBits to perform the operation without doing an explicit conversion. In the $\mathbb{Z}_{2^k}$ case, a similar approach was used previously when combining the SPDZ2k protocol with daBit-style conversions [16]. We adapt these techniques to work with edaBits, in both $\mathbb{Z}_{2^k}$ and $\mathbb{Z}_p$. As an additional contribution, more at the engineering level, we take great care in all our constructions to ensure they work for both signed and unsigned data types. This was not done by previous truncation protocols in $\mathbb{Z}_{2^k}$ based on SPDZ [15,16], which only perform a *logical shift*, as opposed to the *arithmetic shift* that is needed to ensure correctness on signed inputs.

Handling garbled circuits. Our conversion method can also be extended to convert binary shares to garbled circuits, putting the 'Y' into 'ABY' and allowing constant round binary computations. In this paper, we do not focus on this, since the technique is exactly the same as described in [1]; when using binary shares based on TinyOT MACs, conversions between binary and garbled circuit representation comes for free, based on the observation from Hazay et al. [26] that TinyOT sharings can be locally converted into shares of a multi-party garbled circuit.

Performance evaluation. We have implemented our protocol in all relevant security models and computation domains as provided by MP-SPDZ [29], and we found it reduces communication both in microbenchmarks and application benchmarks when comparing to a purely arithmetic or a daBit-based implementation. More concretely, the reduction in communication lies between a factor of 2 and 60 for comparisons from purely arithmetic to edaBits and between 2 and 25 from daBits to edaBits. Improvements in throughput per second are slightly lower but still as high as a factor of 47. Generally, the improvements are higher for dishonest-majority computation and semi-honest security.

We have also compared our implementation with the most established software for mixed circuits [9] and found that it still improves up to a factor of two for a basic benchmark in semi-honest two-party computation. However, they maintain an advantage if the parties are far apart (100 ms RTT) due to the usage of garbled circuits.

Finally, a comparison with a purely arithmetic implementation of deep-learning inference shows an improvement of up to a factor eight in terms of both communication and wall clock time.

1.2   Paper Outline

We begin in Sect. 2 with some preliminaries. In Sect. 3, we introduce edaBits and show how to instantiate them, given a source of private edaBits. We then present our protocol for creating private edaBits in Sect. 4, based on the new cut-and-choose procedure. Then, in Sect. 4.2–4.4 we describe abstract cut-and-choose games that model the protocol, and carry out a formal analysis. Then in Sect. 5 we show how to use edaBits for higher-level primitives like comparison and truncation. Finally, in Sect. 6, we analyze the efficiency of our constructions and present performance numbers from our implementation.

2   Preliminaries

In this work we consider three main algebraic structures: $\mathbb{Z}_M$ for $M = p$ where p is a large prime, $M = 2^k$ where k is a large integer, and $\mathbb{Z}_2$.

2.1   Arithmetic Black-Box

We model MPC via the arithmetic black box model (ABB), which is an ideal functionality in the universal composability framework [10]. This functionality allows a set of n parties $P_1, \ldots, P_n$ to input values, operate on them, and receive outputs after the operations have been performed. Typically (see for example Rotaru and Wood [33]), this functionality is parameterized by a positive integer M, and the values that can be processed by the functionality are in $\mathbb{Z}_M$, with the native operations being addition and multiplication modulo M.

   In this work, we build on the basic ABB to construct edaBits, which are used in our higher-level applications. We therefore consider an extended version of the arithmetic black box model that handles values in both binary and arithmetic domains. First, within one single instance of the functionality we can have both binary and arithmetic computations, where the latter can be either modulo p or modulo 2^k. Furthermore, the functionality allows the parties to convert a single binary share into an arithmetic share of the same bit (but not the other way round). We will use this limited conversion capability to bootstrap to our fully-fledged edaBits, which can convert larger ring elements in both directions, and with much greater efficiency. The details of the functionality are presented in the full version [21].

Notation. As shorthand, we write $[x]_2$ to refer to a secret bit x that has been stored by the functionality $\mathcal{F}_{\mathsf{ABB}}$, and similarly $[x]_M$ for a value $x \in \mathbb{Z}_M$ with $M \in \{p, 2^k\}$. We overload the operators $+$ and $\cdot$, writing for instance, $[y]_M = [x]_M \cdot [y]_M + c$ to denote that the secret values x and y are first multiplied using the Mult command, and then the public constant c is added using LinComb.

Functionality $\mathcal{F}_{\mathsf{edaBits}}$

The functionality is parametrized by $M \in \{2^k, p\}$ and $m \leq \log M$. It has the same features as $\mathcal{F}_{\mathsf{ABB}}$, together with the following command:

Create edaBits: On input $(\mathsf{edabit}, \mathsf{id}_M, \mathsf{id}_2)$ from all parties, sample $(r_0, \ldots, r_{m-1}) \in \mathbb{Z}_2^m$ uniformly at random and store $(\mathsf{binary}, \mathsf{id}_2, r_j)$ for $j = 0, \ldots, m-1$, together with $(\mathsf{arithmetic}, \mathsf{id}_M, r)$, where $r = \sum_{j=0}^{m-1} r_j 2^j$.

Fig. 1. Ideal functionality for extended daBits.

3   Extended daBits

The main primitive of our work is the concept of extended daBits, or *edaBits*. Unlike a daBit, which is a random bit b shared as $([b]_M, [b]_2)$, an edaBit is a collection of bits $(r_{m-1}, \ldots, r_0)$ such that (1) each bit is secret-shared as $[r_i]_2$ and (2) the integer $r = \sum_{i=0}^{m} r_i 2^i$ is secret-shared as $[r]_M$.

One edaBit of length m can be generated from m daBits, and in fact, this is typically the first step when applying daBits to several non-linear primitives like truncation. Instead of following this approach, we choose to generate the edaBits—which is what is needed for most applications where daBits are used—directly, which leads to a much more efficient method and ultimately leads to more efficient primitives for MPC protocols.

At a high level, our protocol for generating edaBits proceeds as follows. Let us think initially of the passively secure setting. Each party P_i samples m random bits $r_{i,0}^i, \ldots, r_{i,m-1}^i$, and secret-shares these bits towards the parties over $\mathbb{Z}_2$, as well as the integer $r_i = \sum_{j=0}^{m-1} r_{i,j} 2^j$ over $\mathbb{Z}_M$. Since each edaBit is known by one party, these edaBits must be combined to get edaBits where no party knows the underlying values. We refer to the former as *private* edaBits, and to the latter as *global* edaBits. The parties combine the private edaBits by adding them together: the arithmetic shares can be simply added locally as $[r]_M = \sum_{i=1}^{n} [r_i]_M$, and the binary shares can be added via an n-input binary adder. Some complications arise, coming from the fact that the r_i values may overflow mod p. Dealing with this is highly non-trivial, and we will discuss this in detail in the description of our protocol in Sect. 3.2. However, before we dive into our construction, we will first present the functionality we aim at instantiating. This functionality is presented in Fig. 1.

3.1   Functionality for Private Extended daBits

We also use a functionality $\mathcal{F}_{\mathsf{edaBitsPriv}}$, which models a *private* set of edaBits that is known to one party. This functionality is defined exactly as $\mathcal{F}_{\mathsf{edaBits}}$, except that the bits $r_0, \ldots, r_{m-1}$ are given as output to one party; additionally, if that party is corrupt, the adversary may instead choose these bits.

The heaviest part of our contribution lies on the instantiation of this functionality, which we postpone to Sect. 4.

3.2 From Private to Global Extended daBits

As we discussed already at the beginning of this section, one can instantiate $\mathcal{F}_{\mathsf{edaBits}}$ using $\mathcal{F}_{\mathsf{edaBitsPriv}}$, by combining the different private edaBits to ensure no individual party knows the underlying values. Small variations are required depending on whether $M = 2^k$ or $M = p$, for reasons that will become clear in a moment.

Now, to provide an intuition on our protocol, assume that the ABB is storing $([r_i]_M, [r_{i,0}]_2, \ldots, [r_{i,m-1}]_2)$ for $i = 1, \ldots, n$, where party P_i knows $(r_{i,0}, \ldots, r_{i,m-1})$ and $r_i = \sum_{j=1}^{m-1} r_{i,j} 2^j$. The parties can add their arithmetic shares to get shares of $r' = \sum_{i=1}^{n} r_i \mod M$, and they can also add their binary shares using a binary n-input adder, which results in shares of the bits of r', only without modular reduction.

Since we want to output a random m-bit integer, the parties need to remove the bits of r' beyond the m-th bit from the arithmetic shares. We have binary shares of these carry bits as part of the output from the binary adder, so using $\log(n)$ calls to ConvertB2A of $\mathcal{F}_{\mathsf{ABB}}$, each of which costs a (regular) daBit, we can convert these to the arithmetic world and perform the correction. Notice that for the case of $M = 2^k$, $m = k$, we can omit this conversion since the arithmetic shares are already reduced.

Even without the correction above, the least significant m bits of r' still correspond to $r_0, \ldots, r_{m-1}$. This turns out to be enough for some applications because it is easy to "delete" the most significant bit in $\mathbb{Z}_{2^k}$ by multiplying with two. We call such an edaBit loose as apposed to a strict one as defined in Fig. 1.

One must be careful with potential overflows modulo M. If $M = 2^k$, then any overflow bits beyond the k-th position can simply be discarded. On the other hand, if $M = p$, as long as $m < \log p$ then we can still subtract the $\log n$ converted carries from the arithmetic shares to correct for any overflow modulo p. The protocol is given in Fig. 2, and the security stated in Theorem 1 below, whose proof follows in a straightforward manner from the correctness of the additions in the protocol. In the protocol, nBitADD denotes an n-input binary adder on m-bit inputs. This can be implemented naively in a circuit with $<(m + \log n) \cdot (n - 1)$ AND gates.

Theorem 1. *Protocol* Π_{edaBits} *UC-realizes functionality* $\mathcal{F}_{\mathsf{edaBits}}$ *in the* $(\mathcal{F}_{\mathsf{edaBitsPriv}}, \mathcal{F}_{\mathsf{B2A}})$-*hybrid model.*

4 Instantiating Private Extended daBits

Our protocol for producing private edaBits is fairly intuitive. The protocol begins with each party inputting a set of edaBits to the ABB functionality. However, since a corrupt party may input inconsistent edaBits (that is, the binary part

---

Protocol Π_{edaBits}

Pre:

- Access to $\mathcal{F}_{\mathsf{edaBitsPriv}}$.
- If $M = p$, then $0 < m < \log(p)$.

Post: The parties get $([r]_M, [r_i]_2, \ldots, [r_i]_2)$ where $r = \sum_{j=1}^{m-1} r_i 2^j$ and the bits are uniform to the adversary.

1. The parties call the functionality $\mathcal{F}_{\mathsf{edaBitsPriv}}$ to get random shares $([r_i]_M, [r_{i,0}]_2, \ldots, [r_{i,m-1}]_2)$, for $i = 1, \ldots, n$. Party P_i additionally learns $r_{i,j}$ and $r_i = \sum_{j=1}^{m-1} r_{i,j} 2^j$.
2. The parties invoke $\mathcal{F}_{\mathsf{ABB}}$ to compute $[r']_M = \sum_{i=1}^{n} [r_i]_M$.
3. The parties invoke $\mathcal{F}_{\mathsf{ABB}}$ to compute $\mathsf{nBitADD}\left(([r_{1,j}]_2)_j, \ldots, ([r_{n,j}]_2)_j\right)$, obtaining $m + \log n$ bits $([b_0]_2, \ldots, [b_{m+\log(n)-1}]_2)$.
4. Call $\mathsf{ConvertB2A}$ from $\mathcal{F}_{\mathsf{ABB}}$ to convert $[b_j]_2 \mapsto [b_j]_M$ for $j = m, \ldots, m + \log(n) - 1$. If $M = 2^k$, values b_j for $j > k$ do not need to be converted, and for the sake of notation, we denote $[b_j]_{2^k} := 0$ for $j > k$.
5. Use $\mathcal{F}_{\mathsf{ABB}}$ to compute $[r]_M = [r']_M - 2^m \sum_{j=0}^{\log(n)-1} [b_{j+m}]_M 2^j$.
6. Output $([r]_M, [b_0]_2, \ldots, [b_{m-1}]_2)$.

Fig. 2. Protocol for generating global edaBits from private edaBits.

---

may not correspond to the bit representation of the arithmetic part), some extra checks must be set in place to ensure correctness. To this end, the parties engage in a consistency check, where each party must prove that their private edaBits were created correctly. We do this with a cut-and-choose procedure, where first a random subset of a certain size of edaBits is opened, their correctness is checked, and then the remaining edaBits are randomly placed into buckets. Within each bucket, all edaBits but the first one are checked against the first edaBit by adding the two in both the binary and arithmetic domains, and opening the result. With high probability, the first edaBit will be correct if all the checks pass.

This method is based on a standard cut-and-choose technique for verifying multiplication triples, used in several other works [22,23]. However, the main difference in our case is that the checking procedure for verifying two edaBits within a bucket is much more expensive: checking two multiplication triples consists of a simple linear combination and opening, whereas to check edaBits, we need to run a binary addition circuit on secret-shared values. This binary addition itself requires $O(m)$ multiplication triples to verify, and the protocol for producing these triples typically requires further cut-and-choose steps to ensure correctness and security.

In this work, we take a different approach to reduce this overhead. First, we allow some of the triples used to perform the check within each bucket to be incorrect, which saves in resources as a triple verification step can be omitted. Furthermore, we observe that these multiplication triples are intended to be

used on inputs that are known to the party proposing the edaBits, and thus it is acceptable if this party knows the bits of the underlying triples as well. As a result, we can simplify the triple generation by letting this party sample the triples together with the edaBits, which is much cheaper than letting the parties jointly sample (even incorrect) triples. Note that even though the triples may be incorrect, they must still be authenticated (in practice, with MACs) by the party who proposes them so that the errors cannot be changed after generating the triples.

To model this, we extend the arithmetic black box model with the following commands, for generating a private triple, and for faulty multiplication, which uses a previously stored triple to do a multiplication.

Input Triple. On input $(\mathsf{Triple}, \mathsf{id}, a, b, c)$ from P_i, where id is a fresh binary identifier and $a, b, c \in \{0, 1\}$, store $(\mathsf{Triple}, i, \mathsf{id}, a, b, c)$.

Faulty Multiplication. On input $(\mathsf{FaultyMult}, \mathsf{id}, \mathsf{id}_1, \mathsf{id}_2, \mathsf{id}_T, i)$ from all parties (where $\mathsf{id}_1, \mathsf{id}_2$ are present in memory), retrieve $(\mathsf{binary}, \mathsf{id}_1, x)$, $(\mathsf{binary}, \mathsf{id}_2, y)$, $(\mathsf{Triple}, i, \mathsf{id}_T, a, b, c)$, compute $z = x \cdot y \oplus (c \oplus a \cdot b)$, and store (id, z).

The triple command can be directly instantiated using Input from $\mathcal{F}_{\mathsf{ABB}}$, while FaultyMult uses Beaver's multiplication technique with one of these triples. Note that in Beaver-based binary multiplication, it is easy to see that any additive error in a triple leads to exactly the same error in the product.

Now we are ready to present our protocol to preprocess private edaBits, described in Fig. 3. The party P_i locally samples a batch of edaBits and multiplication triples, then inputs these into $\mathcal{F}_{\mathsf{ABB}}$. The parties then run the CutNChoose subprotocol, given in Fig. 4, to check that the edaBits provided by P_i are consistent. The protocol outputs a batch of N edaBits, and is parametrized by a bucket size B, and values C, C' which determine how many edaBits and triples are opened, respectively. BitADDCarry denotes a two-input binary addition circuit with a carry bit, which must satisfy the weakly additively tamper resilient property given in the next section. As we will see later, this can be computed with m AND gates and depth $m - 1$.[1]

The cut-and-choose protocol starts by using a standard coin-tossing functionality, $\mathcal{F}_{\mathsf{Rand}}$, to sample public random permutations used to shuffle the sets of edaBits and triples. The coin-tossing can be implemented, for example, with hash-based commitments in the random oracle model. Then the first C edaBits and $C'm$ triples are opened and tested for correctness; this is to ensure that not too large a fraction of the remaining edaBits and triples are incorrect. Then the edaBits are divided into buckets of size B, together with $B - 1$ sets of m triples in each bucket. Then, the top edaBit from each bucket is checked with every other edaBit in the bucket by evaluating a binary addition circuit using the

---

[1] This circuit is rather naive, and in fact there are logarithmic depth circuits with a greater number of AND gates. However, as we will see later in the section, it is important for our security proof to use specifically these naive circuits to obtain the tamper-resilient property. Furthermore, they are only used in the preprocessing phase, so the overhead in round complexity is insignificant in practice.

Protocol $\Pi_{\mathsf{edaBitsPriv}}$

Pre: $\mathcal{F}_{\mathsf{ABB}}$ with modulus M, length parameter $m \in \mathbb{Z}$ with $m \leq \log_2 M$
Post: Batch of N shared edaBits $\{([r_j]_M, [r_{j,0}]_2, \ldots, [r_{j,m-1}]_2)\}_{j=1}^{N}$, where party P_i knows the underlying bits.

1. P_i samples $r_{j,0}, \ldots, r_{j,m-1} \in \mathbb{Z}_2$, for $j = 1, \ldots, NB + C$, and inputs these to $\mathcal{F}_{\mathsf{ABB}}$ in $\mathbb{Z}_2$.
2. P_i computes $r_j = \sum_{i=0}^{m-1} r_{j,i} 2^i$ and inputs $r_j \in \mathbb{Z}_M$ to $\mathcal{F}_{\mathsf{ABB}}$.
3. P_i samples $(N(B-1) + C')m$ random bit triples and inputs these to $\mathcal{F}_{\mathsf{ABB}}$.
4. The parties run the CutNChoose procedure to check the consistency of these edaBits. If the check passes, then the parties obtain N edaBits. Otherwise, they abort.

Fig. 3. Protocol for producing private extended daBits.

triples, and comparing the result with the same addition done in the arithmetic domain. Each individual check in the CutNChoose procedure takes two edaBits of m bits each, and consumes m triples as well as a single regular daBit, needed to convert the carry bit from the addition into the arithmetic domain. Note that when working with modulus $M = 2^k$, if $m = k$ then this conversion step is not needed.

4.1   Weakly Tamper-Resilient Binary Addition Circuit

To implement the BitADDCarry circuit we use a ripple-carry adder, which computes the carry bit at every position with the following equation:

$$c_{i+1} = c_i \oplus ((x_i \oplus c_i) \wedge (y_i \oplus c_i)), \forall i \in \{0, m-1\} \tag{1}$$

where $c_0 = 0$, and x_i, y_i are the i-th bits of the two binary inputs. It then outputs $z_i = x_i \oplus y_i \oplus c_i$, for $i = 0, \ldots, m-1$, and the last carry bit c_m. Note that this requires m AND gates and has linear depth.

Below we define the tamper-resilient property of the circuit that we require. We consider an adversary who can additively tamper with a binary circuit by inducing bit-flips in the output wires of any AND gate.

Definition 1. *A binary circuit* $\mathcal{C} : \mathbb{F}_2^{2m} \to \mathbb{F}_2^{m+1}$ *is weakly additively tamper resilient, if given any tampered circuit* $\mathcal{C}^*$, *obtained by additively tampering* $\mathcal{C}$, *one of the following holds:*

1. $\forall (x, y) \in \mathbb{F}_2^m : \; \mathcal{C}(x, y) = \mathcal{C}^*(x, y)$.
2. $\forall (x, y) \in \mathbb{F}_2^m : \; \mathcal{C}(x, y) \neq \mathcal{C}^*(x, y)$.

Intuitively, this says that the tampered circuit is either incorrect on every possible input, or functionally equivalent to the original circuit. In our protocol, this property restricts the adversary from being able to pass the check with a

---

Procedure CutNChoose

Pre: A batch of $(NB+C)$ shared edaBits $\{([r]_M, [r_0]_2, \ldots, [r_{m-1}]_2)\}_{j=1}^{NB+C}$ and a batch of $(N \cdot (B-1) \cdot m + C' \cdot m)$ triples, all stored in $\mathcal{F}_{\mathsf{ABB}}$, where party P_i knows the underlying bits of the edaBits and the triples.
Post: N verified edaBits
The parties do the following:

1. Using $\mathcal{F}_{\mathsf{Rand}}$, sample two public random permutations and use these to shuffle the edaBits and the triples.
2. Open the first C of the shuffled edaBits in both worlds, and the first $C' \cdot m$ triples. Abort if any of the edaBits or the triples are inconsistent.
3. Place the remaining edaBits into buckets of size B and the triples into buckets of size $(B-1) \cdot m$.
4. For each bucket, select the first edaBit $([r]_M, [r_0]_2, \ldots, [r_{m-1}]_2)$, and for every other edaBit $([s]_M, [s_0]_2, \ldots, [s_{m-1}]_2)$ in the same bucket, perform the following check:
   (a) Let $[r+s]_M = [r]_M + [s]_M$.
   (b) Let $([c_0]_2, \ldots, [c_m]_2) = \mathsf{BitADDCarry}([r_0]_2, \ldots, [r_{m-1}]_2, [s_0]_2, \ldots, [s_{m-1}]_2)$, using the FaultyMult command to evaluate each AND gate.
   (c) Convert $[c_m]_2 \mapsto [c_m]_M$ with ConvertB2A.
   (d) Let $[c']_M = [r+s]_M - 2^m \cdot [c_m]_M$. Open c' and the corresponding bits $c_0, \ldots, c_{m-1}$ from the binary world, and check that $c' = \sum_{i=0}^{m-1} c_i 2^i$.
5. If all the checks pass, output the first edaBit from each of the N buckets.

Fig. 4. Cut-and-choose procedure to check correctness of input edaBits.

tampered circuit with bad edaBits as well as the same circuit with good edaBits. It ensures that if any multiplication triple is incorrect, then the check at that position would only pass with either a good edaBit, or a bad edaBit (but not both).

In the full version, we show that this property is satisfied by the ripple-carry adder circuit above, which we use.

Lemma 1. *The ripple carry adder circuit above is weakly additively tamper-resilient (Definition 1).*

In the case of generating edaBits over $\mathbb{Z}_p$, we still use the ripple-carry adder circuit, and our protocol works as long as the length of the edaBits satisfies $m < \log(p)$. If we wanted edaBits with $m = \lceil \log p \rceil$, for instance to be able to represent arbitrary elements of the field, it seems we would need to use an addition circuit modulo p. Unfortunately, the natural circuit consisting of a binary addition followed by a conditional subtraction is *not* weakly additively tamper resilient. One possible workaround is to use Algebraic Manipulation Detection (AMD) [24,25] circuits, which satisfy much stronger requirements than being weakly additively tamper resilient, however this gives a very large overhead in practice.

Table 1. Number of edaBits produced by CutNChoose for statistical security 2^{-s} and bucket size B, with $C = C' = B$.

s	B	# of edaBits
40	3	≥ 1048576
40	4	≥ 10322
40	5	≥ 1024
80	5	≥ 1048576

4.2 Overview of Cut-and-Choose Analysis

The remainder of this section is devoted to proving that the cut-and-choose method used in our protocol is sound, as stated in the following theorem.

Theorem 2. *Let $N \geq 2^{s/(B-1)}$ and $C = C' = B$, for some bucket size $B \in \{3, 4, 5\}$. Then the probability that the CutNChoose procedure in protocol $\Pi_{\mathsf{edaBitsPriv}}$ outputs at least one incorrect edaBit is no more than 2^{-s}.*

Assuming the theorem above, we can prove that our protocol instantiates the desired functionality, as stated in the following theorem. The only interesting aspect to note about security is that we need $m \leq \log M$ to ensure that the value c' computed in step 4d of CutNChoose does not overflow modulo p when $M = p$ is prime. This guarantees that the check values are computed the same way in the binary and arithmetic domains.

Theorem 3. *Protocol $\Pi_{\mathsf{edaBitsPriv}}$ securely instantiates the functionality $\mathcal{F}_{\mathsf{edaBitsPriv}}$ in the $\mathcal{F}_{\mathsf{ABB}}$-hybrid model.*

To give some idea of parameters, in Table 1 we give the required bucket sizes and number N of edaBits that must be produced to ensure 2^{-s} failure probability according to Theorem 2. Note that these are exactly the same bounds as the standard cut-and-choose procedure without any faulty verification steps from [23]. Our current proof relies on case-by-case analyses for each bucket size, which is why Theorem 2 is not fully general. We leave it as an open problem to obtain a general result for any bucket size.

Overview of Analysis. We analyse the protocol by looking at two abstract games, which model the cut-and-choose procedure. The first game, RealGame, models the protocol fairly closely, but is difficult to directly analyze. We then make some simplifying assumptions about the game to get SimpleGame, and show that any adversary who wins in the real protocol can be translated into an adversary in the SimpleGame. This is the final game we actually analyze.

4.3 Abstracting the Cut-and-Choose Game

We first look more closely at the cut-and-choose procedure by defining an abstract game, RealGame, shown in Fig. 5, that models this process. Note that

---

RealGame

1. $\mathcal{A}$ prepares $NB + C$ shared edaBits $\{([r_j]_M, [r_{j,0}]_2, \ldots, [r_{j,m-1}]_2)\}_{j=1}^{NB+C}$, and batch of $N(B-1) + C'$ potentially tampered circuits $\{\mathcal{C}^*_j\}_{j=1}^{N(B-1)}$ to send to the challenger.
2. The challenger shuffles the edaBits and the circuits using 2 permutations.
3. The challenger opens C edaBits in both worlds and C' circuits randomly. If any of the edaBits are inconsistent, or the circuits have been tampered, Abort.
4. Within each bucket, for every pair of edaBits $(r, (r_i)_i)$ and $(s, (s_i)_i)$, take the next circuit C^* and compute $(c_0, \ldots, c_m) = C^*(r_0, \ldots, r_{m-1}, s_0, \ldots, s_{m-1})$. Compute $c = \sum_{i=0}^{m-1} c_i 2^i$ and check that $r + s - 2^m \cdot c_m$ equals c.

The adversary wins if all the checks pass and there is at least one corrupted edaBit in the output.
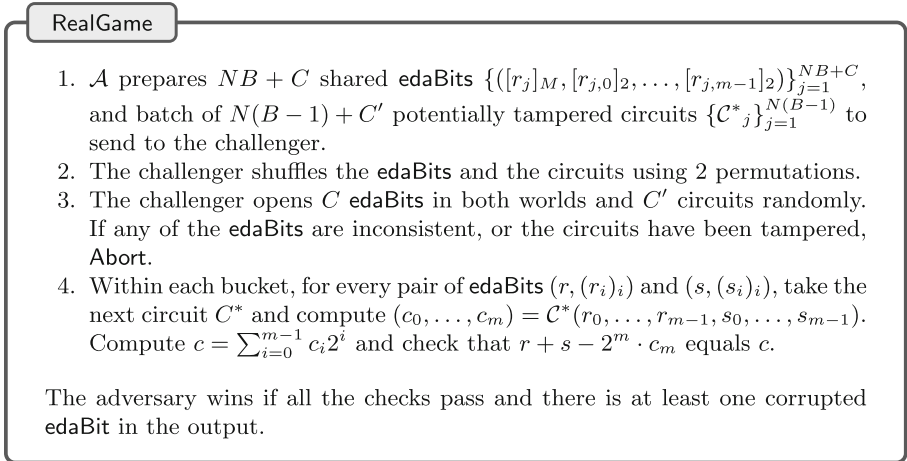

Fig. 5. Abstract game modelling the actual cut-and-choose procedure

in this game, the only difference compared with the original protocol is that the adversary directly chooses additively tampered binary circuits, instead of multiplication triples. The check procedure is carried out exactly as before, so it is clear that this faithfully models the original protocol.

Complexities of Analyzing the Game. In this game, the adversary can pass the check with a bad edaBit in two different ways. The first is to corrupt edaBits in multiples of the bucket size B, and hope that they all end up in the same bucket so that the errors cancel each other out. The second way is to corrupt a set of edaBits and guess the permutation in which they are most likely to end up. Once a permutation is guessed, the adversary will know how many triples it needs to corrupt in order to cancel out the errors, and must also hope that the triples end up in the right place.

To compute the exact probability of all these events, we will also have to consider the number of ways in which the bad edaBits can be corrupted. For edaBits which are m bits, there are up to $2^m - 1$ different ways in which they may be corrupted. On top of that, we have to consider the number of different ways in which these bad edaBits may be paired in the check. In order to avoid enumerating the cases and the complex calculation involved, we simplify the game in a few ways which can only give the adversary a better chance of winning. However, we show that these simplifications are sufficient for our purpose.

4.4 The SimpleGame

In this section we analyze a simplified game and bound the success probability of any adversary in that game by 2^{-s}. Before explaining the simple game, we will leave the complicated world of edaBits and triples. We define a TRIP to be a set of triples that is used to check two edaBits. In our simple world edaBits

SimpleGame

1. $\mathcal{A}$ prepares $NB + C$ balls, corrupts b of them and sends them to the challenger.
2. The challenger opens C of them randomly and checks whether all of them are good. If any one of them is not good, Abort.
3. The challenger permutes and throws NB balls into N buckets each of size B uniformly at random. Then sends the order of arrangement to $\mathcal{A}$.
4. $\mathcal{A}$ prepares $N(B-1) + C'$ triangles, corrupts t of them and sends them to the challenger.
5. The challenger opens C' of them randomly and checks whether all of them are good. If any one of them is not good, Abort.
6. The challenger permutes and throws $N(B-1)$ triangles into N buckets uniformly at random and runs the **Simple** BucketCheck subroutine.
7. If **Simple** BucketCheck returns 1, the challenger outputs first ball from each bucket. Else, Abort.

$\mathcal{A}$ wins if there is no Abort and at least one bad ball is in the output.

Fig. 6. Simplified CutNChoose game

transform into balls, $GOOD$ edaBits into white balls ($\bigcirc$) and BAD edaBits into gray balls ($\bullet$). An edaBit is BAD when at least one of the underlying bits are not correct. TRIPs transform themselves into triangles, $GOOD$ TRIPs into white triangles ($\triangle$) and BAD TRIPs into gray triangles ($\blacktriangle$). We define a TRIP to be BAD when it helps the adversary to win the game, in other words if it can alter the result of addition of two edaBits. Figure 6 illustrates the simple game.

In the SimpleGame $\mathcal{A}$ wins if there is no Abort (means $\mathcal{A}$ passes all the checks) and there is at least one bad ball in the final output. The **simple** BucketCheck checks all the buckets. Precisely, in each bucket two balls are being checked using one triangle. For example, let us consider the size of the buckets $B = 3$. Now one bucket contains three balls $[B1, B2, B3]$ and two triangles $[T1, T2]$. Then BucketCheck checks if the configurations $[B1, B2|T1]$ and $[B1, B3|T2]$ matches any one of these configurations $\{[\bigcirc, \bigcirc|\blacktriangle], [\bigcirc, \bullet|\triangle], [\bullet, \bigcirc|\triangle]\}$. If that is the case then BucketCheck Aborts. When there are two bad balls and one triangle the abort condition depends on the type of bad balls. That means we are considering all bad balls to be distinct, say with different color shades. As a result, in some cases challenger aborts if the checking configuration matches $[\bullet, \bullet|\blacktriangle]$ and in other cases it aborts due to $[\bullet, \bullet|\triangle]$ configuration.

In the simple world everyone has access to a public function f, which takes two bad balls and a triangle as input and outputs 0 or 1. If the output is zero, that means it is a bad configuration, otherwise it is good. This function is isomorphic to the check from step 4 of RealGame, which takes 2 edaBits and a circuit as inputs and outputs the result of the check. The BucketCheck procedure uses f to check all the buckets. Figure 7 illustrates the check in detail. $\mathcal{A}$ passes BucketCheck if

---

Simple BucketCheck

Input: N buckets and a function f. Each bucket contains B balls $\{x_1, \ldots, x_B\}$ and $(B-1)$ triangles $\{y_1, \ldots, y_{B-1}\}$.
Output: 0 or 1.
Runs this check in each bucket:

1. Check the configuration of $[x_1, x_i | y_{i-1}] \; \forall i \in [2, B]$.
   - If $[x_1, x_i | y_{i-1}] \in \{[\bigcirc, \bigcirc | \blacktriangle], [\bigcirc, \bullet | \triangle], [\bullet, \bigcirc | \triangle]\}$ return Reject.
   - If $[x_1, x_i | y_{i-1}] \in [\bullet, \bullet | \triangle]$ and $f(\bullet, \bullet, \triangle) = 0$ return Reject.
   - If $[x_1, x_i | y_{i-1}] \in [\bullet, \bullet | \blacktriangle]$ and $f(\bullet, \bullet, \blacktriangle) = 0$ return Reject.
2. Otherwise return Accept.

If check returns Accept for all the buckets, then output 1; Otherwise output 0.

Fig. 7. A simple bucket check procedure

all the check configurations are favorable to the adversary. These favorable check configurations are illustrated in Table 2.

After throwing triangles, in each bucket, if the check configuration of balls and triangles are from the first three entries of Table 2, then BucketCheck will not Abort. For the last entry BucketCheck will not Abort if the output of f is 1. Notice that if BucketCheck passes only due to the first configuration of Table 2 in all buckets, then the output from each bucket is going to be a good ball and $\mathcal{A}$ loses. So ideally we should take that into account while computing the winning probability of the adversary. However, for most of the cases it is sufficient to show that for large enough N the $\Pr[\mathcal{A}$ passes BucketCheck$]$ is negligible in the statistical security parameter s, as that will bound the winning probability of $\mathcal{A}$ in the simple game.

Before analyzing the SimpleGame, we show that security of RealGame follows directly from security of SimpleGame. Intuitively, that is indeed the case, as in the SimpleGame an adversary chooses number of bad triangles adaptively; Whereas in the RealGame it has to fix the tampered circuits before seeing the permuted edaBits. Thus, if an adversary cannot win the SimpleGame then it must be more difficult for it to succeed in the RealGame.

Lemma 2. *Security against all adversaries in* SimpleGame *implies security against all adversaries in* RealGame.

Table 2. Favorable combination of balls and triangles for the adversary.

Balls		Triangles
○	○	△
○	●	▲
●	○	▲
●	●	△/▲

Throughout the analysis, we use b to denote the number of bad balls and t to denote the number of bad triangles. Now in order to win the SimpleGame the adversary has to pass all the three checks, so let us try to bound the success probability of $\mathcal{A}$ for each of them. Throughout the analysis we consider $N \geq 2^{\frac{s}{B-1}}$, that is for $B \geq 3$, $N(B-1) \geq 2^{\frac{s}{B-1}+1}$ and we are opening $B(\geq 3)$ balls and B triangles in the first two checks.

Opening C balls: In the first check the challenger opens C balls and check whether they are good. So,

$$\Pr[C \text{ balls are good}] = \frac{\binom{NB+C-b}{C}}{\binom{NB+C}{C}} \approx (1 - b/(NB + C))^C.$$

Now for $b = (NB + C)\alpha$, where $1/(NB + C) \leq \alpha \leq 1$, the probability can be written as $(1 - \alpha)^C$. In order to bound the success probability of the adversary with the statistical security parameter s, let us consider the case when $\alpha \geq \frac{2^{s/B}-1}{2^{s/B}}$ and $C = B$. Thus,

$$\Pr[C \text{ balls are good}] \approx (1 - \alpha)^C = (2^{-s/B})^B = 2^{-s}.$$

So if the challenger opens B balls to check then in order to pass the first check $\mathcal{A}$ must corrupt less than α fraction of the balls, where $\alpha = \frac{2^{s/B}-1}{2^{s/B}}$. Lemma 3 follows from the above analysis.

Lemma 3. *The probability of $\mathcal{A}$ passing the first check in* SimpleGame *is less than 2^{-s}, if the adversary corrupts more than α fraction of balls for $\alpha = \frac{2^{s/B}-1}{2^{s/B}}$ and the challenger opens B balls.*

Opening C' triangles: In this case we'll consider the probability of $\mathcal{A}$ passing the second check. This is similar to the previous check, the only difference is that here the challenger opens C' triangles and checks whether they are good. Consequently,

$$\Pr[C' \text{ triangles are good}] = \frac{\binom{N(B-1)+C'-t}{C'}}{\binom{N(B-1)+C'}{C'}} \approx (1 - t/(N(B - 1) + C'))^{C'}.$$

As in the previous case, if t is more than β fraction of the total number of triangles for $\beta = \frac{2^{s/B}-1}{2^{s/B}}$, we can upper bound the success probability of $\mathcal{A}$ by $(2^{-s/B})^{C'}$. Thus for $C' = B$ the success probability of $\mathcal{A}$ in the second check can be bounded by 2^{-s}. Lemma 4 follows from the above analysis.

Lemma 4. *The probability of $\mathcal{A}$ passing the second check in* SimpleGame *is less than 2^{-s}, if the adversary corrupts more than β fraction of triangles for $\beta = \frac{2^{s/B}-1}{2^{s/B}}$ and the challenger opens B triangles.*

Lemmas 3–4 show that it suffices to only look at the first two checks to prove security when the fraction of bad balls or bad triangles is sufficiently large. However, when one of these is small, we also need to analyze the checks within each bucket in the game.

BucketCheck procedure: In this case we consider that the adversary passes first two checks and reaches the last level of the game. However, in order to win the game the adversary has to pass the BucketCheck. Note that now we are dealing with NB balls and the challenger already fixes the arrangement of NB balls in N buckets. Once the ball permutation is fixed that imposes a restriction on the number of favorable (for $\mathcal{A}$) triangle permutations. For example, let us consider that the challenger throws 12 balls into 4 buckets of size 3 and fixes this permutation:

$$\{[\bullet, \circ, \circ][\circ, \circ, \bullet][\bullet, \bullet, \circ][\circ, \circ, \circ]\}$$

Then there are only two possible favorable permutations of triangles:

$$\{[\blacktriangle, \blacktriangle][\triangle, \blacktriangle][\triangle, \blacktriangle][\triangle, \triangle]\}$$
$$\{[\blacktriangle, \blacktriangle][\triangle, \blacktriangle][\blacktriangle, \blacktriangle][\triangle, \triangle]\}$$

Two favorable permutations come from the fact that the third bucket contains two bad balls. From Table 2 we can see that whenever there are two bad balls in a bucket the adversary can pass the check in that bucket either with a good triangle or with a bad triangle. That means both configurations $[\bullet, \bullet|\triangle]$ and $[\bullet, \bullet|\blacktriangle]$ might be favorable to the adversary. Now $\mathcal{A}$ can use the public function f to determine the value of $f(\bullet, \bullet, \triangle)$ and $f(\bullet, \bullet, \blacktriangle)$. In this example, let us consider the value of $f(\bullet, \bullet, \triangle)$ to be 1; Then the first permutation of triangles is favorable to the adversary. As a result the probability of passing the BucketCheck essentially depends on the probability of hitting that specific permutation of triangles among all possible arrangements of triangles. Then the probability of the adversary passing the last check given a specific arrangement of balls L_i is given by:

$$\Pr[\mathcal{A} \text{ passes BucketCheck}|L_i] = 1/\binom{N(B-1)}{t}$$

where $t = N(B-1)\beta$. Thus,

$$\Pr[\mathcal{A} \text{ passes BucketCheck}|L_i] = \frac{(N(B-1)\beta)!(N(B-1)(1-\beta))!}{N(B-1)!}$$

In order to upper bound $\Pr[\mathcal{A}$ passes BucketCheck$]$ we will upper bound the probability for different ranges of α and β. Note that the total probability is given by:

$$\Pr[\mathcal{A} \text{ passes BucketCheck}] = \sum_i \Pr[\mathcal{A} \text{ passes BucketCheck}|L_i] \cdot \Pr[L_i]$$

If we can argue that for all possible $(2^{s/B}-1)/2^{s/B} \geq \alpha \geq 1/NB$, the maximum probability for $\Pr[\mathcal{A}$ passes BucketCheck$|L_i]$, for some configuration L_i, can be bounded by 2^{-s}, then:

$$\Pr[\mathcal{A} \text{ passes BucketCheck}] \leq \sum_i 2^{-s} \cdot \Pr[L_i]$$

Note that the maximum possible value of α is 1, however as the challenger opens C balls and C' triangles, the adversary cannot set α to be 1. To pass the first check $\mathcal{A}$ must set α to be less than $(2^{s/B} - 1)/2^{s/B}$ if the challenger opens B balls and B triangles.

Now let us try to bound $\Pr[\mathcal{A}$ passes BucketCheck$|L_i]$. The value of $\binom{N(B-1)}{t}$ maximizes at $t \approx N(B-1)/2$. Starting from the case when there is no bad triangle, the probability monotonically decreases from 1 to its minimum at $\beta \approx 1/2$, and then it monotonically increases to 1 when all triangles are bad. We analyze the success probability of $\mathcal{A}$ in three cases. These will be discussed in the full version. We summarize the analysis as follows.

Lemma 5. *The probability of $\mathcal{A}$ passing the* BucketCheck *in* SimpleGame *is less than* 2^{-s}, *if* $N \geq 2^{s/(B-1)}$ *and the challenger opens* $C = B$ *balls and* $C' = B$ *triangles during first two checks of* SimpleGame *for* $B \in \{3, 4, 5\}$ *given* $\frac{s}{B-1} > B$.

Combining Lemma 2 and Lemma 5, this completes the proof of Theorem 2.

Remark 1. As we already mentioned the bound we obtain is not general. However, from Lemma 5 it is evident that one can produce more than 1024 edaBits efficiently with 40-bit statistical security using different bucket sizes with our CutNChoose technique, which is sufficient for the applications we are considering in this work. It also shows that if we want to achieve 80-bit statistical security for $N \geq 2^{20}$, then increasing the bucket size from 3 to 5 would be sufficient. Table 1 shows the number of edaBits we can produce with different size of buckets.

5   Primitives

This section describes the high-level protocols we build using our edaBits, both over $\mathbb{Z}_{2^k}$ and $\mathbb{Z}_p$. We focus on secure truncation (Sect. 5.1) and secure integer comparison (Sect. 5.2), although our techniques apply to a much wider set of non-linear primitives that require binary circuits for intermediate computations. For example, our techniques also allow us to compute binary-to-arithmetic and arithmetic-to-binary conversions of shared integers, by plugging in our edaBits into the conversion protocols from [11] and [16] for the field and ring cases, respectively.

Throughout this section our datatypes are signed integers in the interval $[-2^{\ell-1}, 2^{\ell-1})$. On the other hand, our MPC protocols operate over a modulus $M \geq 2^\ell$ which is either 2^k or a prime p. Given an integer $\alpha \in [-2^{\ell-1}, 2^{\ell-1})$, we can associate to it the corresponding ring element in $\mathbb{Z}_M$ by computing $\alpha \bmod M \in \mathbb{Z}_M$ (modular reduction returns integers in $[0, M)$). We denote this map by $\mathsf{Rep}_M(\alpha)$, and we may drop the sub-index M when it is clear from context. Finally, in the protocols below LT denotes a binary less-than circuit.

5.1   Truncation

Recall that our datatypes are signed integers in the interval $[-2^{\ell-1}, 2^{\ell-1})$, represented by integers in $\mathbb{Z}_M$ where $M \geq 2^\ell$ via $\mathsf{Rep}_M(\alpha) = \alpha \bmod M$. The goal

of a truncation protocol is to obtain $[y]$ from $[a]$, where $y = \mathsf{Rep}\left(\left\lfloor\frac{\alpha}{2^m}\right\rfloor\right)$ and where $a = \mathsf{Rep}(\alpha)$. This is a crucial operation when dealing with fixed-point arithmetic, and therefore an efficient solution for it has a substantial impact in the efficiency of MPC protocols for a wide range of applications. An important observation is that, as integers, $\left\lfloor\frac{\alpha}{2^m}\right\rfloor = \frac{\alpha - (\alpha \bmod 2^m)}{2^m}$. If M is an odd prime p, this corresponds in $\mathbb{Z}_p$ to $y = (\mathsf{Rep}(\alpha) - \mathsf{Rep}(\alpha \bmod 2^m)) \cdot \mathsf{Rep}(2^m)^{-1}$. Furthermore, $\mathsf{Rep}(\alpha \bmod 2^m) = \alpha \bmod 2^m = a \bmod 2^m$ and $\mathsf{Rep}(2^m) = 2^m$, so $y = \frac{a - (a \bmod 2^m)}{(2^m)^{-1}}$.

We focus below in truncation over $\mathbb{Z}_{2^k}$ as it is the less studied case. For the case of truncation over $\mathbb{F}_p$ we refer the reader to the full version [21].

Truncation over $\mathbb{Z}_{2^k}$. Truncation protocols over fields typically exploit the fact that one can divide by powers of 2 modulo p. This is not possible when working modulo 2^k. Instead, we take a different approach. Let $[a]_{2^k}$ be the initial shares, where $a = \mathsf{Rep}(\alpha)$ with $\alpha \in [-2^{\ell-1}, 2^{\ell-1})$ (notice that it may be the case that $\ell < k$). First, we provide a method, $\mathsf{LogShift}$, for computing the *logical* right shift of a by m positions, assuming that $a \in [0, 2^\ell)$. That is, if a is

$$(\underbrace{0, \ldots, 0}_{k-\ell}, \underbrace{a_{\ell-1}, \ldots, a_0}_{\ell}),$$

this procedure will yield shares of

$$(\underbrace{0, \ldots, 0}_{k-\ell+m}, \underbrace{a_{\ell-1}, \ldots, a_m}_{\ell-m}).$$

Then, to compute the arithmetic shift, we use the fact that[2]

$$\left\lfloor\frac{\alpha}{2^m}\right\rfloor \equiv \mathsf{LogShift}_m(a + 2^{\ell-1}) - 2^{\ell-m-1} \bmod 2^k.$$

Now, to compute the logical shift, our protocol begins just like in the field case by computing shares of $a \bmod 2^m$ and subtracting them from a, which produces shares of $(a_{k-1}, \ldots, a_m, 0, \ldots, 0)$. The parties then open a masked version of $a - (a \bmod 2^m)$ which does not reveal the upper $k - \ell$ bits, and then shift to the right by m positions in the clear, and undo the truncated mask. One has to account for the overflow that may occur during this masking, but this can be calculated using a binary LT circuit.

The details of our logical shift protocol are provided in Fig. 8, and we analyze its correctness next. First, it is easy to see that $c = 2^{k-m}((a + r) \bmod 2^m)$, so $c/2^{k-m} = (a \bmod 2^m) + r - 2^m v$, where v is set if and only if $c/2^{k-m} < r$. From this we can see that the first part of the protocol $[a \bmod 2^m]_{2^k}$ is correctly computed. Privacy of this first part follows from the fact that $r \bmod 2^m$ completely masks $a \bmod 2^m$ when c is opened.

---

[2] Notice that we can use the $\mathsf{LogShift}$ method on $a + 2^{\ell-1}$ since, $\alpha + 2^{\ell-1} \in [0, 2^\ell)$, which implies that $(a + 2^{\ell-1}) \bmod 2^k = \alpha + 2^{\ell-1}$ and therefore $(a + 2^{\ell-1}) \bmod 2^k$ is ℓ-bits long, as required.

Logical right shift over $\mathbb{Z}_{2^k}$

Pre:
- $\mathcal{F}_{\mathsf{ABB}}$
- Input $[a]_{2^k}$ where $a \in [0, 2^\ell)$.
- Number of bits to shift m
- edaBit $([r]_{2^k}, [r]_2)$ of length m
- edaBit $([r']_{2^k}, [r']_2)$ of length $\ell - m$

Post: $[y]_{2^k}$, where $y = \mathsf{LogShift}_m(a)$.

1. The parties compute shares of $a \bmod 2^m$ as follows:
   (a) Call $c = \mathsf{open}\left(2^{k-m} \cdot ([a]_{2^k} + [r]_{2^k})\right)$
   (b) Compute $[v]_2 = \mathsf{LT}((c_i)_{i=k-m+1}^k, ([r_i]_2)_{i=0}^{m-1})$
   (c) Convert $[v]_2 \mapsto [v]_{2^k}$
   (d) Let $[a \bmod 2^m]_{2^k} = 2^m [v]_{2^k} - [r]_{2^k} + c/2^{k-m}$.
2. The parties compute the truncation:
   (a) Compute $[b]_{2^k} = [a]_{2^k} - ([a]_{2^k} \bmod 2^m)$.
   (b) Call $d = \mathsf{open}(2^{k-\ell} \cdot ([b]_{2^k} + 2^m [r']_{2^k}))$.
   (c) Compute $[u]_2 = \mathsf{LT}((d_i)_{i=k-\ell+m}^{k-1}, ([r'_i]_2)_{i=0}^{\ell-m-1})$
   (d) Convert $[u]_2 \mapsto [u]_{2^k}$.[a]
   (e) Output $[y]_{2^k} = 2^{\ell-m} [u]_{2^k} + d/2^{k-\ell+m} - [r']_{2^k}$

---

[a] One can optimize this by noticing that we only need shares of u modulo $2^{k-\ell+m}$.

Fig. 8. Protocol for performing logical right-shift

For the second part, let us write $b = 2^m a'$, then $d = 2^{k-\ell+m}((a' + r') \bmod 2^{\ell-m})$, so $d/2^{k-\ell+m} = a' + r' - 2^{\ell-m} u$, where u is set if and only if $d/2^{k-\ell+m} < r'$, as calculated by the protocol. We get then that $a' = d/2^{k-\ell+m} - r' + 2^{\ell-m} u$, and since a' is precisely $\mathsf{LogShift}_m(a)$, we conclude the correctness analysis.

Probabilistic Truncation. Recall that in the field case one can obtain probabilistic truncation avoiding a binary circuit, which results in a constant number of rounds. Over rings this is a much more challenging task. For example, probabilistic truncation with a constant number of rounds is achieved in ABY3 [30], but requires, like in the field case, a 2^s gap between the secret values and the actual modulus, which in turn implies that only small non-negative values can be truncated.

In Fig. 9, we take a different approach. Intuitively, we follow the same approach as in ABY3, which consists of masking the value to be truncated with a shared random value for which its corresponding truncation is also known, opening this value, truncating it and removing the truncated mask. In ABY3 a large gap is required to ensure that the overflow that may happen by the masking process does not occur with high probability. Instead, we allow this overflow bit to be non-zero and remove it from the final expression. Doing this naively would

> **Probabilistic truncation over $\mathbb{Z}_{2^k}$**
>
> **Pre:**
> - $\mathcal{F}_{\mathsf{ABB}}$
> - Input $[a]_{2^k}$ where $a \in [0, 2^\ell)$.
> - $\ell < k$
> - Number of bits to truncate m
> - edaBit $([r]_{2^k}, [r]_2)$ of length $(\ell - m)$
> - edaBit $([r']_{2^k}, [r']_2)$ of length m
> - Random bit $[b]_{2^k}$
>
> **Post:** $[y]_{2^k}$ where $y = \lfloor a/2^m \rfloor + u$ with $u = 1$ with probability $(a \bmod 2^m)/2^m$.
>
> 1. Call $c = \mathsf{open}(2^{k-\ell-1} \cdot ([a]_{2^k} + 2^\ell [b]_{2^k} + 2^m [r]_{2^k} + [r']_{2^k}))$. Write $c = 2^{k-\ell-1}c'$.
> 2. Compute $[v]_{2^k} = [b \oplus c'_\ell]_{2^k} = [b]_{2^k} + c'_\ell - 2c'_\ell [b]_{2^k}$
> 3. Output $[y]_{2^k} = (c' \bmod 2^\ell)/2^m - [r]_{2^k} + 2^{\ell-m} [v]_{2^k}$

Fig. 9. Probabilistic truncation in domain modulo power of two using edaBits

require us to compute a LT circuit, but we avoid doing this by using the fact that, because the input is positive, the overflow bit can be obtained from the opened value by making the mask value also positive. This leaks the overflow bit, which is not secure, and to avoid this we mask this single bit with another random bit. This protocol can be seen as an extension of the probabilistic truncation protocol by Dalskov et al. [15]. Below, we provide an analysis for our extension that also applies to said protocol.

Now we analyze the protocol. First we notice that $c = 2^{k-\ell-1}c'$ where $c' = (2^m r + r') + a + 2^\ell b - 2^{\ell+1}vb$, where v is set if and only if $(2^m r + r') + a$ overflows modulo 2^ℓ. It is easy to see that this implies that $c'_\ell = v \oplus b$, so we see that $v = c'_\ell \oplus b$, as calculated in the protocol.

On the other hand, we have that $(c' \bmod 2^\ell) = (2^m r + r') + a - 2^\ell v$, so $a \bmod 2^m = (c' \bmod 2^m) - r' + 2^m u$, where u is set if $(c' \bmod 2^m) < r'$. From this it can be obtained that $\lfloor (c' \bmod 2^\ell)/2^m \rfloor - r + 2^{\ell-m} = \lfloor a/2^m \rfloor + u$.

Remark 2. The protocol we discussed above only works if $a \in [0, 2^\ell)$, that is, if the value α represented $\alpha \in [0, 2^{\ell-1})$. We can extend it to $\alpha \in [-2^{\ell-1}, 2^{\ell-1})$ by using the same trick as in the deterministic truncation: The truncation is called with $a + 2^{\ell-1}$ as input, and $2^{\ell-m-1}$ is subtracted from the output.

5.2   Integer Comparison

Another important primitive that appears in many applications is integer comparison. In this case, two secret integers $[a]_M$ and $[b]_M$ are provided as input, and the goal is to compute shares of $\alpha \overset{?}{<} \beta$, where $a = \mathsf{Rep}(\alpha)$ and $b = \mathsf{Rep}(\beta)$.

As noticed by previous works (e.g. [11,16]), this computation reduces to extracting the MSB from a shared integer as follows: If $\alpha, \beta \in [-2^{k-2}, 2^{k-2})$,

Table 3. Amortized costs for generating 1 Private, and 1 Global edaBit. Costs for Global edaBits do not include the cost of the n additional sets of Private edaBits that are needed.

	Private edaBits		Global edaBits	
	$\mathbb{Z}_{2^k}$	$\mathbb{F}_p$	$\mathbb{Z}_{2^k}$	$\mathbb{F}_p$
Faulty edaBits	B	B	0	0 $(l-m+s,\,m)$
Faulty Triples	$(B-1)m$	$(B-1)m$	0	0
Secure Triples	0	0	$(\log n)(n-1)$	$(\log n)(n-1)$
daBits	0	$(B-1)$	0	$\log n$
Openings $(\mathbb{Z}_2)$	$(3m+1)(B-1)$	$(3m+1)(B-1)$	$(2m+2\log n)(n-1)$	$(2m+3\log n)(n-1)$
Openings $(\mathbb{Z}_M)$	$(B-1)$	$(B-1)$	0	0

then $\alpha - \beta = [-2^{k-1}, 2^{k-1})$, so $a - b = \mathsf{Rep}(\alpha - \beta)$ corresponds to the sign of $\alpha - \beta$, which is minus (i.e. the bit is 1) if and only if α is smaller than β.

To extract the MSB, we simply notice that $\mathsf{MSB}(\alpha) = -\left\lfloor \frac{\alpha}{2^{k-1}} \right\rfloor \bmod 2^k$, so this can be extracted with the protocols we have seen in the previous sections.

6 Applications and Benchmarks

6.1 Theoretical Cost

We present the theoretical costs of the different protocols in the paper, starting with the cost for producing Private and Global edaBits in terms of the different parameters.

Table 3 shows the main amortized costs for generating a Private and Global edaBit of length m. For Global edaBits, we assume have the required correct Private edaBits to start with, which is why number of Faulty edaBits needed is 0. B is the bucket size for the cut-and-choose procedure and n is the number of parties.

Table 4 shows the cost for two of our primitives from Sect. 5, namely comparison of m-bit numbers and truncation of an ℓ-bit number by m binary digits. For computation modulo a prime, there is also a statistical security parameter s.

Comparison in $\mathbb{Z}_{2^k}$ is our only application where it suffices to use loose edaBits (where the relation between the sets of shares only holds modulo 2^m, c.f. Sect. 3.2). This is because the arithmetic part of an edaBit is only used in the first step (the masking) but not at the end. Recall that the truncation protocols always use the arithmetic part of an edaBit twice, once before opening and once to compute an intermediate or the final result. Using a loose edaBit would clearly distort the result. With comparison on the other hand, an edaBit is only used to facilitate the conversion to binary computation, after which the result is converted back to arithmetic computation using a classic daBit.

6.2 Implementation Results

We have implemented our approach in a range of domains and security models, and we have run the generation of a million edaBits of length 64 on AWS

Table 4. Cost of our primitives. Numbers in brackets indicate edaBit length.

	Comparison		Truncation	
	$\mathbb{Z}_{2^k}$	$\mathbb{F}_p$	$\mathbb{Z}_{2^k}$	$\mathbb{F}_p$
Strict edaBits	0	2 $(m+1, s+1)$	2 $(l-m, m)$	2$(l-m+s, m)$
Loose edaBits	1 $(m+1)$	0	0	0
Classic daBits	1	1	2	1
Online ANDs	$\sim 2m$	$\sim 2m$	$\sim 2m$	$\sim 2k$

Table 5. Number of edaBits generated (in $1000\,$s) per second in various settings

		Domain	Strict edaBits	Loose edaBits
Dishonest maj.	Malicious	2^k (OT)	4.6	7.3
		p (OT)	3.6	4.2
		p (HE)	2.7	3.4
	Semi-hon.	2^k (OT)	456.7	922.5
		p (OT)	228.0	892.6
		p (HE)	470.5	905.6
Honest maj.	Malicious	2^k	191.5	205.8
		p	156.6	162.1
	Semi-hon.	2^k	2032.1	7180.0
		p	1367.7	4934.3

c5.9xlarge with the minimal number of parties required by the security model (two for dishonest majority and three for honest majority). Table 5 shows the throughput for various security models and computation domains, and Table 6 does so for communication. For computation modulo a prime with dishonest majority, we present figures for arithmetic computation both using oblivious transfer (OT) and LWE-based semi-homomorphic encryption (HE). Note that the binary computation is always based on oblivious transfer for dishonest majority and that all our results include all consumable preprocessing such as multiplication triples but not one-off costs such as key generation. The source code of our implementation has been added to MP-SPDZ [14].

We have also implemented 63-bit[3] comparison using edaBits, only daBits, and neither, and we have run one million comparisons in parallel again on AWS c5.9xlarge. Table 7 shows the throughput for our various security models and computation domains, and Table 8 does so for communication. Note that the arithmetic baseline uses either the protocol of Catrina and de Hoogh [11] ($\mathbb{F}_p$) or the variant by Dalskov et al. [15] ($\mathbb{Z}_{2^k}$).

---

[3] Comparison in secure computation is generally implemented by extracting the most significant bit of difference. This means that 63-bit is the highest accuracy achievable in computation modulo 2^{64}, which the natural modulus on current 64-bit platforms.

Table 6. Communication per edaBit (in kbit) in various settings

		Domain	Strict edaBits	Loose edaBits
Dishonest maj.	Malicious	2^k (OT)	1335.5	480.2
		p (OT)	1936.9	1473.2
		p (HE)	940.8	779.7
	Semi-hon.	2^k (OT)	22.5	9.6
		p (OT)	43.9	9.6
		p (HE)	11.8	9.6
Honest maj.	Malicious	2^k	5.6	3.7
		p	7.6	6.4
	Semi-hon.	2^k	0.3	0.2
		p	0.5	0.2

Table 7. Number of comparisons (in 1000 s) per second in various settings

		Domain	Arithm.	daBits	edaBits
Dishonest maj.	Malicious	2^k (OT)	0.5	1.2	4.4
		p (OT)	0.3	0.3	1.6
		p (HE)	0.6	0.7	2.0
	Semi-hon.	2^k (OT)	5.2	14.4	275.6
		p (OT)	1.6	3.3	79.7
		p (HE)	5.9	12.8	170.6
Honest maj.	Malicious	2^k	76.4	119.2	170.4
		p	66.9	78.3	80.1
	Semi-hon.	2^k	500.6	1007.7	1607.6
		p	157.8	277.1	457.6

Our results highlight the advantage of our approach over using only daBits. The biggest improvement comes in the dishonest majority with semi-honest security model. For the dishonest majority aspect, this is most likely because there is a great gap in the cost between multiplications and inputs (the latter is used extensively to generate edaBits). For the semi-honest security aspect, note that our approach for malicious security involves a cascade of sacrificing because the edaBit sacrifice involves binary computation, which in turn involves further sacrifice of AND triples. Finally, the improvement in communication is generally larger than the improvement in wall clock time. We estimate that this is due to the fact that switching to binary computation clearly reduces communication but increases the computational complexity.

Table 8. Communication per comparison (in kbit) in various settings

		Domain	Arithm.	daBits	edaBits
Dishonest maj.	Malicious	2^k (OT)	21737.7	9058.6	1310.5
		p (OT)	40108.5	34019.1	4783.3
		p (HE)	3020.5	3210.9	1584.8
	Semi-hon.	2^k (OT)	2283.0	830.2	39.0
		p (OT)	7353.1	3503.0	134.9
		p (HE)	411.6	219.1	38.7
Honest maj.	Malicious	2^k	63.4	27.8	5.4
		p	94.3	85.0	19.9
	Semi-hon.	2^k	14.5	7.1	0.4
		p	37.4	23.1	1.4

Table 9. Overall time and communication for biometric matching

		LAN (s)	WAN (s)	Comm. (MB)
$n = 1000$	ABY/HyCC (A+Y)	0.22	2.5	9.5
	ABY/HyCC (A+B)	0.22	6.1	10.6
	Ours	0.12	8.3	7.4
$n = 4096$	ABY/HyCC (A+Y)	0.63	6.6	40.4
	ABY/HyCC (A+B)	0.72	13.6	43.6
	Ours	0.48	12.6	29.1
$n = 13684$	ABY/HyCC (A+Y)	3.66	17.5	138.0
	ABY/HyCC (A+B)	5.4	26.2	190.8
	Ours	2.00	22.9	111.8

6.3 Comparison to Previous Works

Dishonest majority. The authors of HyCC [9] report figures for biometric matching with semi-honest two-party computation in ABY [20] and HyCC. The algorithm essentially computes the minimum over a list of small-dimensional Euclidean distances. The aforementioned authors report figures in LAN (1 Gbps) and artificial WAN settings of two machines with four-core i7 processors. For a fair comparison, we have run our implementation using one thread limiting the bandwidth and latency accordingly. Table 9 shows that our results improves on the time in the LAN setting and on communication generally as well as on the in the WAN setting for larger instances compared to their A+B setting (without garbled circuits). The WAN setting is less favorable to our solution because it is purely based on secret sharing and we have not particularly optimized the number of rounds.

Honest majority. Our approach is not directly comparable to the one by Mohassel and Rindal [30] because they use the specifics of replicated secret sharing for the conversion. We do note however that their approach of restricting binary circuits to the binary domain is comparable to our solution, and that they use the same secret sharing schemes as us in the 2^k domain. The full version [21] shows a comparison of their results with our approach applied to logistic regression.

daBits. Aly et al. [1] report figures for daBit generation with dishonest majority and malicious security in eight threads over a 10 Gbps network. For two-party computation using homomorphic-encryption, they achieve 2150 daBits per second at a communication cost of 94 kbit per daBit. In a comparable setting, we found that our protocol produces 12292 daBits per second requiring a communication cost of 32 kbit. Note however that Aly et al. use somewhat homomorphic encryption while our implementation is based on cheaper semi-homomorphic encryption.

Convolutional Neural Networks. We also apply our techniques to the convolutional neural networks considered be Dalskov et al. [15]. See the full version for details.

References

1. Aly, A., Orsini, E., Rotaru, D., Smart, N.P., Wood, T.: Zaphod: efficiently combining LSSS and garbled circuits in SCALE. In: WAHC 2019: Proceedings of the 7th ACM Workshop on Encrypted Computing & Applied Homomorphic Cryptography. ACM (2019). https://eprint.iacr.org/2019/974
2. Araki, T., et al.: Generalizing the SPDZ compiler for other protocols. In: Lie, D., Mannan, M., Backes, M., Wang, X. (eds.) ACM CCS 2018, pp. 880–895. ACM Press, October 2018
3. Bonte, C., Smart, N.P., Tanguy, T.: Thresholdizing HashEdDSA: MPC to the rescue. Cryptology ePrint Archive, Report 2020/214 (2020). https://eprint.iacr.org/2020/214
4. Boyle, E., et al.: Efficient two-round OT extension and silent non-interactive secure computation. In: Cavallaro, L., Kinder, J., Wang, X., Katz, J. (eds.) ACM CCS 2019, pp. 291–308. ACM Press, November 2019
5. Boyle, E., Couteau, G., Gilboa, N., Ishai, Y., Kohl, L., Scholl, P.: Efficient pseudorandom correlation generators: silent OT extension and more. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part III. LNCS, vol. 11694, pp. 489–518. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-26954-8_16

6. Boyle, E., Gilboa, N., Ishai, Y.: Function secret sharing. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015, Part II. LNCS, vol. 9057, pp. 337–367. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46803-6_12

7. Boyle, E., Gilboa, N., Ishai, Y.: Secure computation with preprocessing via function secret sharing. In: Hofheinz, D., Rosen, A. (eds.) TCC 2019, Part I. LNCS, vol. 11891, pp. 341–371. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-36030-6_14

8. Burra, S.S., et al.: High performance multi-party computation for binary circuits based on oblivious transfer. Cryptology ePrint Archive, Report 2015/472 (2015). http://eprint.iacr.org/2015/472

9. Büscher, N., Demmler, D., Katzenbeisser, S., Kretzmer, D., Schneider, T.: HyCC: compilation of hybrid protocols for practical secure computation. In: Lie, D., Mannan, M., Backes, M., Wang, X. (eds.) ACM CCS 2018, pp. 847–861. ACM Press, October 2018

10. Canetti, R.: Universally composable security: a new paradigm for cryptographic protocols. In: 42nd FOCS, pp. 136–145. IEEE Computer Society Press, October 2001

11. Catrina, O., de Hoogh, S.: Improved primitives for secure multiparty integer computation. In: Garay, J.A., De Prisco, R. (eds.) SCN 2010. LNCS, vol. 6280, pp. 182–199. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-15317-4_13

12. Chandran, N., Gupta, D., Rastogi, A., Sharma, R., Tripathi, S.: EzPC: programmable and efficient secure two-party computation for machine learning. In: 2019 IEEE European Symposium on Security and Privacy (EuroS&P), pp. 496–511. IEEE (2019)

13. Cramer, R., Damgård, I., Escudero, D., Scholl, P., Xing, C.: SPD $\mathbb{Z}_{2^k}$: efficient MPC mod 2^k for dishonest majority. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018. LNCS, vol. 10992, pp. 769–798. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-96881-0_26

14. CSIRO's Data61. MP-SPDZ (2020). https://github.com/data61/MP-SPDZ

15. Dalskov, A., Escudero, D., Keller, M.: Secure evaluation of quantized neural networks. Cryptology ePrint Archive, Report 2019/131 (2019). https://eprint.iacr.org/2019/131

16. Damgård, I., Escudero, D., Frederiksen, T.K., Keller, M., Scholl, P., Volgushev, N.: New primitives for actively-secure MPC over rings with applications to private machine learning. In: 2019 IEEE Symposium on Security and Privacy, pp. 1102–1120. IEEE Computer Society Press, May 2019

17. Damgård, I., Fitzi, M., Kiltz, E., Nielsen, J.B., Toft, T.: Unconditionally secure constant-rounds multi-party computation for equality, comparison, bits and exponentiation. In: Halevi, S., Rabin, T. (eds.) TCC 2006. LNCS, vol. 3876, pp. 285–304. Springer, Heidelberg (2006). https://doi.org/10.1007/11681878_15

18. Damgård, I., Keller, M., Larraia, E., Pastro, V., Scholl, P., Smart, N.P.: Practical covertly secure MPC for dishonest majority – or: breaking the SPDZ limits. In: Crampton, J., Jajodia, S., Mayes, K. (eds.) ESORICS 2013. LNCS, vol. 8134, pp. 1–18. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40203-6_1

19. Damgård, I., Pastro, V., Smart, N., Zakarias, S.: Multiparty computation from somewhat homomorphic encryption. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 643–662. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32009-5_38

20. Demmler, D., Schneider, T., Zohner, M.: ABY - a framework for efficient mixed-protocol secure two-party computation. In: NDSS 2015. The Internet Society, February 2015

21. Escudero, D., Ghosh, S., Keller, M., Rachuri, R., Scholl, P.: Improved primitives for MPC over mixed arithmetic-binary circuits. Cryptology ePrint Archive, Report 2020/338 (2020). https://eprint.iacr.org/2020/338

22. Frederiksen, T.K., Keller, M., Orsini, E., Scholl, P.: A unified approach to MPC with preprocessing using OT. In: Iwata, T., Cheon, J.H. (eds.) ASIACRYPT 2015, Part I. LNCS, vol. 9452, pp. 711–735. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48797-6_29

23. Furukawa, J., Lindell, Y., Nof, A., Weinstein, O.: High-throughput secure three-party computation for malicious adversaries and an honest majority. In: Coron, J.-S., Nielsen, J.B. (eds.) EUROCRYPT 2017, Part II. LNCS, vol. 10211, pp. 225–255. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-56614-6_8

24. Genkin, D., Ishai, Y., Prabhakaran, M., Sahai, A., Tromer, E.: Circuits resilient to additive attacks with applications to secure computation. In: Shmoys, D.B. (ed.) 46th ACM STOC, pp. 495–504. ACM Press, May/June (2014)

25. Genkin, D., Ishai, Y., Weiss, M.: Binary AMD circuits from secure multiparty computation. In: Hirt, M., Smith, A. (eds.) TCC 2016, Part I. LNCS, vol. 9985, pp. 336–366. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53641-4_14

26. Hazay, C., Scholl, P., Soria-Vazquez, E.: Low cost constant round MPC combining BMR and oblivious transfer. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017, Part I. LNCS, vol. 10624, pp. 598–628. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-70694-8_21

27. Ishai, Y., Kilian, J., Nissim, K., Petrank, E.: Extending oblivious transfers efficiently. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 145–161. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-45146-4_9

28. Ishaq, M., Milanova, A.L., Zikas, V.: Efficient MPC via program analysis: a framework for efficient optimal mixing. In: Cavallaro, L., Kinder, J., Wang, X., Katz, J. (eds.) ACM CCS 2019, pp. 1539–1556. ACM Press, November 2019

29. Keller, M.: MP-SPDZ: a versatile framework for multi-party computation. Cryptology ePrint Archive, Report 2020/521 (2020). https://eprint.iacr.org/2020/521

30. Mohassel, P., Rindal, P.: ABY^3: a mixed protocol framework for machine learning. In: Lie, D., Mannan, M., Backes, M., Wang, X. (eds.) ACM CCS 2018, pp. 35–52. ACM Press, October 2018

31. Nielsen, J.B., Nordholt, P.S., Orlandi, C., Burra, S.S.: A new approach to practical active-secure two-party computation. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 681–700. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32009-5_40

32. Rotaru, D., Smart, N.P., Tanguy, T., Vercauteren, F., Wood, T.: Actively secure setup for SPDZ. Cryptology ePrint Archive, Report 2019/1300 (2019). https://eprint.iacr.org/2019/1300

33. Rotaru, D., Wood, T.: MArBled circuits: mixing arithmetic and Boolean circuits with active security. In: Hao, F., Ruj, S., Sen Gupta, S. (eds.) INDOCRYPT 2019. LNCS, vol. 11898, pp. 227–249. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-35423-7_12

34. Wang, X., Ranellucci, S., Katz, J.: Global-scale secure multiparty computation. In: Thuraisingham, B.M., Evans, D., Malkin, T., Xu, D. (eds.) ACM CCS 2017, pp. 39–56. ACM Press, October 2017