

Improved Building Blocks for Secure Multi-Party Computation based on Secret Sharing with Honest Majority

Marina Blanton¹, Ahreum Kang², and Chen Yuan¹

¹ University at Buffalo (SUNY)

² Soonchunhyang University

Abstract. Secure multi-party computation permits evaluation of any desired functionality on private data without disclosing the data to the participants and is gaining its popularity due to increasing collection of user, customer, or patient data and the need to analyze data sets distributed across different organizations without disclosing them. Because adoption of secure computation techniques depends on their performance in practice, it is important to continue improving their performance. In this work, we focus on common non-trivial operations used by many types of programs, and any advances in their performance would impact the runtime of programs that rely on them. In particular, we treat the operation of reading or writing an element of an array at a private location and integer multiplication. The focus of this work is secret sharing setting with honest majority in the semi-honest security model. We demonstrate improvement of the proposed techniques over prior constructions via analytical and empirical evaluation.

1 Introduction

Secure multi-party computation refers to the ability of a number of participants to evaluate a function of their choice on private data without disclosing unintended information about the private data to the computation participants. It has been the subject of research for many years with its performance experiencing tremendous progress during the last decade. Such techniques are now suitable for data and computations of significant sizes. Furthermore, they can be increasingly applied to perform analysis of large private data sets distributed among a number of participants, as well as data analytics and decision making using private distributed data (including medical, financial, and other domains).

Of particular interest to the research community in recent years has been privacy-preserving machine learning, which uses non-trivial algorithms to analyze large volumes of data. Computation used in such analyses often requires access to data at private locations, be it due to the nature of data representation, e.g., in the form of sparse data sets or due to the nature of the algorithm itself. When such operations are executed as part of secure computation on private data, we must employ data-oblivious (i.e., data-independent) constructions for realizing the operations to eliminate leakage of private information. In the case of accessing memory at a private location, we could either access each location of the data set or array or employ more complex randomized techniques such as Oblivious RAM (ORAM) for secure computation. The latter has lower asymptotic complexities as a function of the memory size, but are more complex

to set up and invoke. As a result, the former approach known as linear scan outperforms known ORAM constructions for memory of small to medium sizes. While improved secure computation ORAM techniques in both two-party and multi-party settings are an active area of research, in this work we focus on improving performance of linear scan in the multi-party setting, which is also not entirely straightforward operation.

In addition, we revisit the multiplication operation with optimizations in the computational (as opposed to information-theoretic) setting. While information-theoretic security might be considered stronger than computational security, all known secure multi-party computation frameworks rely on secure channels implemented in the computational setting. Multiplication is a fundamental building blocks of many secure computation frameworks based on arithmetic circuits and is ubiquitously used for realizing more complex operations including linear scan.

Motivating example. Consider the problem of building an exact machine learning model such as a Bayesian network from a distributed data set located at different organizations (for instance, patient information located at different hospitals). A data set includes a number of attributes or features (e.g., age, gender, medical diagnosis, BMI, etc. in the case of medical data) with an instance of the data set corresponding to a user, customer, or patient. Constructing a model consists of determining correlation between different attributes based on instances located at different locations. This can be accomplished using the so-called variable assignment or parent assignment problem [22], which is the critical component of Bayesian network learning, Markov blankets identification and, more generally, feature selection [22, 18]. Correlation between different variables is commonly computed using the MDL score [27] which uses conditional entropy computation as part of the score calculation. This computation is heavy on the use of logarithms with private inputs, which is an expensive function when implemented within a secure multi-party computation framework. The computation takes the form of $\log(X)$, where X corresponds to the number of instances with a combination of specific values for some features (and thus is private), and the function is called extensively on different values of X .

Our observation is that, instead of evaluating the logarithm function directly, we could build an alternative solution of much higher speed. In particular, because X is integer and ranges between 0 and the (combined) data set size, we can pre-compute the values of $\log(i)$ for each i in that range and store the result in array A . Then executing $\log(X)$ would translate into retrieving the value of stored at the private location X of array A . When the size of the array (which is proportional to the number of data set instances) is not very large, a solution based on a linear scan would outperform other alternatives such as using ORAM or evaluating the logarithm function itself. In this work we thus revisit existing solutions to implementing read or write operations at private locations in the multi-party setting based on secret sharing (SS) and show that significant optimizations are possible.

Our contributions. In this work we develop new constructions for access to an array at a private location (read or write) that significantly outperform conventional implementations of this operations in the setting with honest majority based on secret sharing. We present a general construction which works for any number of computation participants n that uses conventional Shamir secret sharing. We also present a custom construction

for the common case of three parties, which outperforms the general construction. Because it uses 2-out-of-2 additive secret sharing, we show how to convert between that representation and conventional three-party Shamir SS.

We also develop optimizations for the multiplication operation based on Shamir SS in the computational setting. We provide two constructions: the first has communication complexity linear in the number of parties n (i.e., constant per party) and practical performance. The second construction has communication complexity quadratic in the number of parties, but offers lower communication complexity for the important case of $n = 3$ parties with a single field element transmitted by each party. This matches communication of the best known custom three-party multiplication protocols (designed for custom replicated secret sharing) from [3]. Our optimizations are tailored to the setting where the number of computational parties is small.

We implement our constructions on operations of varying sizes in the setting of the PICCO compiler [36] and show that they significantly outperform the previous implementations adopted by PICCO.

2 Related Work

Conventional implementations of performing an array access at a private location via a linear scan can be comparison-based or multiplexer-based as we further discuss in section 4. Optimizations to the simple solutions are available in both two-party setting based on garbled circuits (which does not directly apply to the content of this work) and in the multi-party setting. The closest to our work is the construction due to Laud [23] for array read, which is applicable to both Shamir SS and Sharemind framework. The goal of that work was to minimize the online work (which depends on the private inputs), while our goal is to minimize the overall work. As a result, the proposed solution from [23] has large round complexity. It also offers optimizations, the most effective of which is applicable only to the Sharemind framework. We draw a more detailed comparison to the construction from [23] and our solutions in section 4.

Laud [24] proposed efficient protocols for reading and writing elements of an array at private locations in parallel. The solution is based on sorting and for ℓ parallel read requests to an array of size m has complexity $O((m + \ell) \log(m + \ell))$. Because the solution is non-trivial and was implemented in the Sharemind framework, we are unable to empirically compare its runtime to our constructions designed for Shamir secret sharing. We, however, note that based on the best known oblivious sorting algorithms, this construction will require $O((m + \ell) \log(m + \ell))$ comparisons each of cost $O(k)$ for k -bit integers. Based on our detailed analysis of its possible implementation in our setting, we expect that it might outperform our proposed constructions only when both m and ℓ are large. That is, when the number of parallel invocations ℓ is small or when the array size m is not large (even with a very large number of parallel invocations), we expect our constructions to outperform the solution from [24].

Oblivious RAM [13, 25, 14] can also be used to realize array read or write at a private location, where a client outsources its private memory to a remote server without revealing any information including access patterns to the server. There are many results (see, e.g., [29, 31, 29, 32, 26, 11]) including publications in the multi-server setting

(e.g., [30, 15]), but they are still not applicable to secure multi-party computation because the client’s work is not distributed and the client has access to the data in the clear. In the context of ORAM for secure computation, SCORAM [34] was among the early constructions in the two-party setting, after which a number of improvements such as [33, 35, 9] followed. There are also multi-party or three-party constructions in different settings including [19, 20, 10, 17, 5]. These constructions become competitive with approaches based on a linear scan only for rather large array or dataset sizes. We compare performance of our constructions to the state-of-the-art ORAM in section 6.

Integer multiplication is a fundamental operation of any secure computation framework based on arithmetic circuits. Its original efficient implementation in Shamir SS is due to Gennaro et al. [12], while most recent version of multiplication for Sharemind can be found in [21]. The best known multiplication we are aware of is in a custom three-party replicated SS from [3], which we match with an n -party construction based on Shamir SS in this work. Additional information is provided in section 5.

3 Preliminaries

3.1 Secure Multi-Party Computation

We consider the conventional secure multi-party setting with n computational parties, out of which at most t can be corrupt. We work in the setting with honest majority, i.e., $t < n/2$ and focus on security against semi-honest participants, in which the participants are trusted to follow the prescribed computation, but might attempt to learn unauthorized information based on the information they possess. We use the standard simulation-based security definition that requires that the participants do not learn any information beyond their intended output. For completeness, it is provided in Appendix A. Due to space limitations, we provide only sketches of our security arguments.

As customary with techniques based on SS, the set of computational parties does not have to coincide with (and can be formed independently from) the set of parties supplying inputs in the computation (input providers) and the set of parties receiving output (output recipients). Then if a computational party learns no output, the computation should not reveal any information to that party. Consequently, if we wish to design a functionality that takes input in the secret-shared form and produces shares of the output, any computational party should learn nothing from protocol execution.

3.2 Secret Sharing

A secret sharing scheme allows one to produce shares of secret x such that access to a predefined number of shares reveals no information about x . In the case of (n, t) threshold secret sharing schemes, there are n participants, each of which receive their own shares. The security requirement is that possession of shares stored at any t or fewer parties reveals no information about x , while access to shares stored at $t + 1$ or more parties allows for efficient reconstruction of x . We refer to this type of secret sharing as t -sharing. Of particular importance to secure multi-party computation is linear secret sharing schemes, which have the property that a linear combination of secret shared values can be performed locally on the shares.

Shamir secret sharing. Shamir secret sharing [28] (SSS) is an (n, t) -linear SS scheme with $t < n/2$, where computation takes places over a finite field \mathbb{F} . Then each secret value $s \in \mathbb{F}$ is represented by a random polynomial of degree t with the free coefficient set to s . Each share of s corresponds to the evaluation of the polynomial on a unique non-zero point. Consequently, given $t + 1$ or more shares, the parties can reconstruct the polynomial and learn s using Lagrange interpolation. Possession of t or fewer shares, however, information-theoretically reveals no information about s . With this representation, computation of any linear combination of secret-shared values is performed locally by each party using its shares, while multiplication requires interaction.

In what follows, we use notation $[x]$ to denote that the value of x is secret shared among the participants using (Shamir) t -sharing. We also let $[x]_p$ denote the share of x stored at party $p \in [1, n]$. Because each secret-shared value is a field element, the size of the field \mathbb{F} needs to be large enough to be able to represent values in the desired range. For example, to be able to support computation on k -bit integers, we must have that $|\mathbb{F}| \geq 2^k$. Furthermore, because we rely on certain building blocks from [6], some of them may place additional constraints on the field size (to increase the size by a certain amount) as specified in [6].

As customary in the literature, we measure performance in the number of elementary interactive operations (such as multiplication, opening the shares of private value, etc.) and the number of sequential operations, i.e., rounds. Local operations are not included in the cost due to their speed.

Replicated secret sharing. Replicated secret sharing (RSS) [16] is another types of linear secret sharing, which can be used to realize (n, t) -threshold secret sharing (and can be defined for more general access structures Γ , but we limit our use to threshold structures only). RSS can be defined for any $n \geq 2$ and any $t < n$ and works over any finite ring. The RSS access structure uses the notion of qualified sets, which are all subsets of the participants who are permitted to reconstruct the secret, i.e. all subsets of $t + 1$ or more parties in our case, while all other subsets are called unqualified. To secret-share private $x \in \mathbb{F}$ using RSS, we additively split it into shares x_T such that $x = \sum_{T \in \mathcal{T}} x_T$ (in \mathbb{F}), where \mathcal{T} consists of all maximal unqualified sets (i.e., all sets of t parties in our case). Then each party $p \in [1, n]$ stores shares x_T for all $T \in \mathcal{T}$ subject to $p \notin T$. In the general case of (n, t) -threshold RSS, the total number of shares is $\binom{n}{t}$ with $\binom{n-1}{t}$ shares stored by each party, which can become large as n and t grow. However, for small n , the number of shares is small (e.g., with both $(3, 1)$ and $(3, 2)$ RSS, there are the total of 3 shares).

An important optimization on which we rely is non-interactive evaluation of a pseudo-random function (PRF) using RSS in the computational (as opposed to number-theoretic) setting as proposed in [7]. In particular, [7] provide a mechanism for non-interactive generation of Shamir secret shares from RSSs as follows: suppose that shares of secret key k have been distributed to the parties according to a (n, t) -threshold RSS and let $\text{PRF} : \{0, 1\}^\kappa \times \{0, 1\}^* \rightarrow \mathbb{F}$ denote a pseudo-random function that takes a sufficiently large κ -bit key (κ is a security parameter). On input a , the parties collectively compute shares of $x = \text{PRF}_k(a) = \sum_{T \in \mathcal{T}} \text{PRF}_{k_T}(a)$ (in \mathbb{F}), where each party $p \in [1, n]$ computes its (Shamir) share of x as $[x]_p = \sum_{T \in \mathcal{T}, p \notin T} \text{PRF}_{k_T}(a) \cdot f_T(p)$. Here, for each $T \in \mathcal{T}$, f_T refers to the unique polynomial of degree t such that

$f_T(0) = 1$ and $f_T(p) = 0$ for each $p \in T$. We note that each (replicated) share k_T needs to be sufficiently long, while the computed (Shamir) shares $[x]_p$ can be within a smaller field. We denote this operation by PRSS (pseudo-random secret sharing).

4 Array Access at a Private Location

We next proceed with describe our constructions, which are in the honest majority setting based on Shamir secret sharing. In this section we treat optimizations to array access at a private location, while the next section discusses integer multiplication.

Assume that we are given an array of m (private or public) elements a_0, \dots, a_{m-1} and would like to retrieve the element a_j at a private index j . Conventional implementations of this functionality via linear scan include (i) privately comparing j to every integer in the range $[0, m - 1]$ and computing the dot product of the resulting bits and the array elements and (ii) bit-decomposing the index and using a multiplexer to retrieve the desired element. The latter approach was implemented in the PICCO compiler (which uses conventional Shamir secret sharing arithmetic) [36], while the former was later shown to be slightly faster in this framework [4]. Array write is implemented similarly, where instead of computing the dot product (i.e., a sum of products), we update each element of the array based on the result of an individual product.

4.1 General Construction

Our starting point for the general construction was the first traditional approach above where we privately compare j to each position of the array and retrieve the element for which the result of the comparison was true. If we let EQ denote the operation of privately comparing two integers with at least one of them being private, this operation can be represented as follows:

```
[b] ← ArrayRead( $\langle [a_0], \dots, [a_{m-1}] \rangle, [j]$ )
1. for  $i = 0$  to  $m - 1$ , compute in parallel  $[c_i] \leftarrow \text{EQ}([j], i)$ ;
2.  $[b] \leftarrow \sum_{i=0}^{m-1} [c_i] \cdot [a_i]$ ;
3. return  $[b]$ ;
```

This computation is written to take an array of private elements as its input, but when the elements are public, the computation proceeds similarly. Also, to turn this operation into the write operation, we replace line 2 with $[b_i] \leftarrow [c_i] \cdot [w] + (1 - [c_i])[a_i] = [c_i]([w] - [a_i]) + [a_i]$ for $i = 0, \dots, m - 1$, where $[w]$ is the value to be written at location $[j]$, and return array $[b_0], \dots, [b_{m-1}]$.

To optimize performance of this solution, our first observation stems from the fact that j is compared to all indices and part of the computation might be redundant. Let us look at the details of the comparison operation. The most efficient constant-round equality testing in this setting is due to Catrina and de Hoogh [6]. It proceeds by comparing a single private integer a to 0, which is normally the difference between two values being compared. The bitlength k of the operand a is given as the second argument.

```
[b] ← EQZ( $[a], k$ )
1.  $([r'], [r], [r_{k-1}], \dots, [r_0]) \leftarrow \text{PRandM}(k, k)$ ;
```

2. $c \leftarrow \text{Open}([a] + 2^k[r'] + [r]);$
3. $(c_{k-1}, \dots, c_0) \leftarrow \text{Bits}(c, k);$
4. for $i = 0, \dots, k - 1$ do $[d_i] \leftarrow c_i + [r_i] - 2c_i[r_i];$
5. $[b] \leftarrow 1 - \text{KOr}([d_{k-1}], \dots, [d_0]);$
6. return $[b];$

Here, the operation $\text{PRandM}(k, x)$ assumes that we work with k -bit integers and generates a $k + \kappa$ -bit random integer for a statistical security parameter κ , the x least significant bits of which are available in the bit-decomposed form. The returned result is x random bits r_0, \dots, r_{x-1} , x -bit $r = \sum_{i=0}^{x-1} 2^i r_i$, and $k + \kappa - x$ -bit integer r' . The Open function reveals the value of its private argument. $\text{Bits}(c, x)$ simply returns the x least significant bits of its public argument c . Lastly, KOr computes the k -ary OR of its k private arguments.

This operation hides the value of a by adding large random $2^k \cdot r' + r$ to it and opening the sum. Because the bits of r are available (as r_0 through r_{k-1}), the remaining computation can efficiently compute the bits of a (in step 4) and consequently test whether at least one of them is 1 (in step 5) using k -ary OR of k bits. The cost of this operation is dominated by PRandM which contributes k (parallel) interactive operations, while KOr costs $4 \log(k)$ and Open costs 1 interactive operations, respectively. The overall number of rounds is 4.

When we compare private j to all possible indices i in the set, we know that the adjacent values of $j - i$ differ by 1. Thus, instead of generating independent random bits for each $j - i$ via a new call to PRandM , we could execute this function once, protect j using the random values as in step 2 above, and open this protected value as c . Given the protected value c of j , we can then form protected values of $j - i$ as $c - 0, c - 1, \dots, c - (m - 1)$ assuming that i ranges from 0 to $m - 1$. In other words, the computation for array read with a private index becomes:

- $$[b] \leftarrow \text{ArrayRead}(\langle [a_0], \dots, [a_{m-1}] \rangle, [j])$$
1. $([r'], [r], [r_{\log m-1}], \dots, [r_0]) \leftarrow \text{PRandM}(\log m, \log m);$
 2. $c \leftarrow \text{Open}([j] + 2^{\log m}[r'] + [r]);$
 3. for $i = 0$ to $m - 1$, do in parallel
 - (a) $v \leftarrow c - i;$
 - (b) $(v_{\log m-1}, \dots, v_0) \leftarrow \text{Bits}(v, \log m);$
 - (c) for $l = 0, \dots, \log m - 1$ do $[d_l] \leftarrow v_l + [r_l] - 2v_l[r_l];$
 - (d) $[b_i] \leftarrow 1 - \text{KOr}([d_{\log m-1}], \dots, [d_0]);$
 4. $[b] \leftarrow \sum_{i=0}^{m-1} [b_i] \cdot [a_i];$
 5. return $[b];$

This optimization reduces the cost of array read from $m(\log m + 4 \log \log m + 1) + 1$ interactive operations in 5 rounds to $4m \log \log m + \log m + 2$ in 5 rounds. Alternatively, we could use a simple tree-like implementation of KOr with $\log m - 1$ interactive operations in $\log \log m$ rounds, which makes the complexity of ArrayRead be $m(\log m - 1) + \log m + 1$ in $\log \log m + 2$ rounds.

This, however, still appears redundant because the bits of v , and consequently bits d provided as input into the k -ary OR in step 3(d), are often reused from one loop iteration i to another. For example, we know that c and $c - 1$ are going to differ in their least significant bits, but a number of most significant bits might be the same. Also,

because the bitlength of j is $\log m$, we know that most of (or all) possible combinations of bits assembled will be used in KOr across all i . In other words, for any given v , its i th bit will be either the i th bit of c or its complement, and most of all possible $2^{\log m}$ combinations of bits will be used across all i s to form vs . Thus, we design an efficient mechanism for computing OR of all possible combinations of bits and then incorporate it in the private lookup protocol.

Our algorithm for computing ORs of bits uses a divide-and-conquer approach, where we split the original size into two halves, recurse on each half, and then assemble the result. It is denoted as AllOr and given below. On input k bits d_i , it computes 2^k k -ary ORs of the form $\bigvee_{i=0}^{k-1} c_i$, where c_i is either d_i or its complement $\neg d_i$.

$\langle [b_0], \dots, [b_{2^k-1}] \rangle \leftarrow \text{AllOr}([d_{k-1}], \dots, [d_0])$

1. if $(k = 1)$ return $\langle [d_0], 1 - [d_0] \rangle$;
2. else
3. $\ell \leftarrow \lfloor k/2 \rfloor$;
4. $[u_0], \dots, [u_{2^\ell-1}] \leftarrow \text{AllOr}([d_{\ell-1}], \dots, [d_0])$;
5. $[v_0], \dots, [v_{2^{k-\ell}-1}] \leftarrow \text{AllOr}([d_{k-1}], \dots, [d_\ell])$;
6. for $i = 0, \dots, 2^{k-\ell} - 1$ and $j = 0, \dots, 2^\ell - 1$ do in parallel $[b_{2^\ell i + j}] \leftarrow [v_i] + [u_j] - [v_i] \cdot [u_j]$;
7. return $\langle [b_0], \dots, [b_{2^k-1}] \rangle$;

To integrate this solution into our array read protocol, we apply AllOr to the bits r_i s computed in step 1 of ArrayRead and, as before, reveal the value of j protected by r , with its $\log m$ least significant bits denoted by c' . The intuition is now that the computed k -ary ORs correspond to all possible k -ary ORs over all k -bit integers “shuffled” based on the value of r and the only OR that evaluates to 0 will be at position r . This means that if we would like to know whether, e.g., $j = 0$, we need to test whether $c' = r$ or if the c' th position in the array of k -ary ORs corresponds to 0. Similarly, for testing whether $j = i$, we test whether $c' = r + i$ (or $r = c' - i$) and retrieve the $(c' - i)$ th value in the returned array. Lastly, because we need a single OR evaluate to 1 with the remaining values being 0, we complement the result of the AllOr operation. (Note that the original implementation of EQZ from [6] computes $c \oplus r$ instead of $c - r$ prior to calling KOr using a more complex logic to show correctness of the algorithm, but the same approach does not work in our case.) We obtain the following solution:

$[b] \leftarrow \text{ArrayRead}(\langle [a_0], \dots, [a_{m-1}] \rangle, [j])$

1. $([r'], [r], [r_{\log m-1}], \dots, [r_0]) \leftarrow \text{PRandM}(\log m, \log m)$;
2. $\langle [b_0], \dots, [b_{2^{\log m}-1}] \rangle \leftarrow \text{AllOr}([r_{\log m-1}], \dots, [r_0])$;
3. for $i = 0, \dots, 2^{\log m} - 1$, $[b_i] = 1 - [b_i]$;
4. $c \leftarrow \text{Open}([j] + 2^{\log m} [r'] + [r])$;
5. $c' \leftarrow c \bmod 2^{\log m}$;
6. $[b] \leftarrow \sum_{i=0}^{m-1} [b_{c'-i \bmod 2^{\log m}}] \cdot [a_i]$;
7. return $[b]$;

To realize the write operation at private index $[j]$, we replace line 6 of ArrayRead above with the computation $[d_i] \leftarrow [b_{c'-i \bmod 2^{\log m}}]([w] - [a_i]) + [a_i]$ for $i = 0, \dots, m - 1$, where as before $[w]$ corresponds to the value being written, and return the updated array $[d_0], \dots, [d_{m-1}]$.

The cost of ArrayRead dominated by that of AllOr protocol. The recurrence in AllOr can be specified as $T(k) = 2T(k/2) + 2^k$ and thus the function has complexity $\Theta(2^k)$, or $\Theta(m)$ when $k = \log m$. Furthermore, the constant behind the asymptotic notation is low and the number of interactive operations per array element reduces as the array size increases. For example, with $m = 2^4$, AllOr executes 1.5 multiplication per array element (i.e., 24), with $m = 2^8$, it is < 1.19 multiplications per array element, and with $m = 2^{16}$, it is < 1.01 per array element. The remaining steps in ArrayRead contribute $\log m + 2$ interactive operations. The round complexity of AllOr with $\log m$ -bit argument is $\log \log m$, which means that the overall number of rounds of ArrayRead is $\log \log m + 3$. Furthermore, the first three steps can be precomputed, which makes the online number of rounds to be 2 and the online number of interactive operations is also 2. Implementing array write at a private location increases the total (and online) number of interactive operations by $m - 1$ without affecting round complexity.

An alternative solution for this operation developed by Laud in [23] uses $m + 3$ interactive operations in $m + 3$ rounds³ (using Shamir secret sharing), where most of the work can be carried offline with the online work being 3 interactive operations in 3 rounds. The linear round complexity is nevertheless prohibitive, especially in the WAN setting. The round complexity of the array read from [23] can be reduced to a constant at the cost of increasing the number of multiplications by several times, at which point our construction is attractive and uses only a fraction of that cost. Thus, we offer practical performance improvement over known results.

To demonstrate security, we note that all instructions are input-independent and follow a similar structure to that of EQZ from [6]. All steps operate on shares except step 4, in which the value of c is revealed. The value of c corresponds to private j protected by a random value at least κ bits longer than j . This means that the probability that any information is revealed about j is negligible in the security parameter κ and is therefore acceptable. This implies that we are able to simulate the adversarial view without access to the inputs.

4.2 Custom Three-Party Construction

We also provide a second construction which is designed to work only with $n = 3$ parties using custom computation, but offers superior performance compared to the general construction. The construction uses 2-out-of-2 additive secret sharing, which means that if we would like to use it together with a standard secret sharing framework such as Shamir secret sharing, we need to provide procedures for converting between the two representations. This is what we consequently do in this section as well.

In what follows, we use notation $\llbracket x \rrbracket$ to denote that the value of $x \in \mathbb{F}$ is secret shared using 2-out-of-2 additive secret sharing. We note that this solution works over any finite ring, which has performance benefits such as using native hardware implementations of arithmetic in \mathbb{Z}_{2^k} for some k . For the purposes of this work, we let computation to be over a finite field to be compatible with other constructions we propose.

Because in this representation the shares are held by two parties out of three, for concreteness of the presentation, we let the notation include the parties holding the

³ This information is not explicitly provided in [23], but rather is deduced by us.

shares. Thus, we use $\llbracket x \rrbracket_{p_1 p_2}$ to indicate that the value is split between parties $p_1, p_2 \in [1, 3]$ with $p_1 \neq p_2$. For example, we might use $\llbracket x \rrbracket_{12}$. Then notation $\llbracket x \rrbracket_{p_1}$ and $\llbracket x \rrbracket_{p_2}$ denotes the individual shares when x is secret shared as $\llbracket x \rrbracket_{p_1 p_2}$.

In our construction, the data set is originally additively shared between parties 1 and 2 (i.e., we have $\llbracket a_0 \rrbracket_{12}, \dots, \llbracket a_{m-1} \rrbracket_{12}$), while the private index j can be shared-shared using any linear secret sharing scheme and for simplicity we assume it is shared using Shamir secret sharing as $\llbracket j \rrbracket$. The intuition behind our solution is that the data set is rotated by a private number of positions and the value of j gets adjusted by that value. Then the parties who do not have information about the entire amount of rotation are revealed the modified value of j and read the element at that position. To implement this idea, we need to be careful to ensure is that reading the element is performed on the shares to prevent any single party from having access to the read element, while at the same time we must have that the parties with cleartext access to the modified j do not know by which value j modified from its original value.

To realize this, we instruct parties 1 and 2 to rotate their shares of the data set by random amount $r_1 \in \mathbb{Z}_m$ known only to the two of them. Next, party 1 re-shares its shares of the data set between parties 2 and 3, which makes the rotated data set to be shared between these two parties. Now parties 2 and 3 again rotate the shared data set by random amount r_2 known only to the two of them, after which party 2 re-shares its data set shares among parties 1 and 3. At this point we have the data set rotated by amount $r_1 + r_2$ is shared between parties 1 and 3 and neither of them knows the value of $r_1 + r_2$. Thus, we open $h = (j + r_1 + r_2) \bmod m$ to parties 1 and 3 who consequently retrieve the element at position h in their data sets and return their share as the output.

In our solution, we propose that the parties generate r_1 and r_2 non-interactively using a shared seed to a pseudo-random generator. That is, parties 1 and 2 share key k_{12} , while parties 2 and 3 share key k_{23} . Because generation of r_1 and r_2 is a one-time cost independent of the data set size, any other suitable mechanism for agreeing on these values will work as well (e.g., if one wants to maintain information-theoretic security of the protocol). The computation then proceeds as follows:

$\llbracket b \rrbracket \leftarrow \text{ArrayRead}(\langle \llbracket a_0 \rrbracket_{12}, \dots, \llbracket a_{m-1} \rrbracket_{12} \rangle, \llbracket j \rrbracket)$

1. Parties 1 and 2 agree on random $r_1 \in \mathbb{Z}_m$ and locally rotate their shares as $\langle \llbracket a_{r_1} \rrbracket_p, \dots, \llbracket a_{m-1} \rrbracket_p, \llbracket a_0 \rrbracket_p, \dots, \llbracket a_{r_1-1} \rrbracket_p \rangle \leftarrow \langle \llbracket a_0 \rrbracket_p, \dots, \llbracket a_{m-1} \rrbracket_p \rangle$, where $p \in [1, 2]$, and also let $\llbracket h \rrbracket \leftarrow \llbracket j \rrbracket + r_1$.
2. Party 1 randomly generates $s_i \in \mathbb{F}$ for $i \in [0, m-1]$ and sends $\langle s_0, \dots, s_{m-1} \rangle$ to party 2, who consequently sets $\llbracket a'_i \rrbracket_2 = \llbracket a_i \rrbracket_2 + s_i$ for $i \in [0, m-1]$.
3. Party 1 sets $\llbracket a'_i \rrbracket_3 = \llbracket a_i \rrbracket_1 - s_i$ for $i \in [0, m-1]$ and sends $\langle \llbracket a'_0 \rrbracket_3, \dots, \llbracket a'_{m-1} \rrbracket_3 \rangle$ to party 3.
4. Parties 2 and 3 agree on random $r_2 \in \mathbb{Z}_m$ and locally rotate their shares as $\langle \llbracket a'_{r_2} \rrbracket_p, \dots, \llbracket a'_{m-1} \rrbracket_p, \llbracket a'_0 \rrbracket_p, \dots, \llbracket a'_{r_2-1} \rrbracket_p \rangle \leftarrow \langle \llbracket a'_0 \rrbracket_p, \dots, \llbracket a'_{m-1} \rrbracket_p \rangle$ and let $\llbracket h \rrbracket \leftarrow \llbracket h \rrbracket + r_2$.
5. Party 2 randomly generates $s'_i \in \mathbb{F}$ for $i \in [0, m-1]$ and sends $\langle s'_0, \dots, s'_{m-1} \rangle$ to party 3, who consequently sets $\llbracket a''_i \rrbracket_3 = \llbracket a'_i \rrbracket_3 + s'_i$ for $i \in [0, m-1]$.
6. Party 2 sets $\llbracket a''_i \rrbracket_1 = \llbracket a'_i \rrbracket_2 - s'_i$ for $i \in [0, m-1]$ and sends $\langle \llbracket a''_0 \rrbracket_1, \dots, \llbracket a''_{m-1} \rrbracket_1 \rangle$ to party 1.
7. The value of $h \bmod m$ is opened to parties 1 and 3 who set $\llbracket b \rrbracket_p = \llbracket a''_h \rrbracket_p$ for $p \in [1, 3]$.

8. Return $\llbracket b \rrbracket_{13}$.

This computation is dominated by communicating $4m$ elements in two rounds, i.e., similar to that of executing m multiplications in parallel. There might also be communication with computing h or $h \bmod m$ depending on the underlying secret sharing scheme. In particular, if h is secret-shared using additive secret sharing in \mathbb{Z}_m , no additional communication is needed. In particular, with additive secret sharing, we would need to modify only one of the shares to perform addition of r_1 or r_2 , and using arithmetic in \mathbb{Z}_m the opened value will be in \mathbb{Z}_m as desired. When using a different type of secret sharing such as Shamir SS, the parties would need to update h and re-share its value across all parties with fresh randomness. Similarly, when computation is not in \mathbb{Z}_m , computing $h \bmod m$ is needed prior to opening the value. For example, with SSS, one might invoke efficient Mod protocol from [6] (integer division with public divisor). This is a one-time operation of cost at most $O(\log m)$ and does not have a significant impact on the performance of the overall protocol.

If the parties instead would like to evaluate the write operation and store value $\llbracket w \rrbracket$ at private index j , we modify the protocol above to have parties 1 and 3 update the element at position h with shares of w in step 7. This is sufficient for this operation. However, if the values are to be opened instead of being used in consecutive computation, they would need to be re-randomized.

To show security in the three-party setting with a single corrupt party, we argue that the data set remains information-theoretically protected from any participant. In particular, it is always secret-shared among two parties. Furthermore, the value of j is also information-theoretically protected from the parties if r_1 and r_2 are chosen randomly (and otherwise is computationally protected). Thus, it can be shown that the simulated view with no access to real data is indistinguishable from a real run of the protocol.

Lastly, to permit this construction to be used in conjunction with SSS techniques, we provide conversion procedures to and from 2-out-of-2 additive secret sharing and SSS over the same field \mathbb{F} . Due to space limitations, their description is deferred to Appendix B. The cost of converting a field element to or from additive SS is that of communicating two field elements. This means that for a read operation, the cost of converting the inputs and outputs of ArrayRead is communicating about $2m$ elements, while for a write operation it is $4m$. The cost of the conversion, however, can be amortized among multiple operations if these operations are repeatedly called on the same data set without other intermediate operations.

5 Multiplication

The conventional multiplication protocol from [12] results in communicating the total of $n(n - 1)$ field elements in the n -party setting, with each party sending $n - 1$ field elements. This means that in the 3-party setting, the total of 6 elements are transmitted. Sharemind's multiplication protocol from [21] also results in communicating 6 elements with 3 computational parties. Here we design and present two new multiplication protocols suitable for use with Shamir secret sharing that lower this communication cost. The first communicates at most $2(n - 1)$ field elements and thus has lower communication cost than the protocol from [12] for any n , and in particular communicates

Protocol	n -party			3-party		
	comm.	rounds	comp.	comm.	rounds	comp.
Mult1 (sec. 5.1)	$1 + \frac{2t-1}{n}$	2	$O(n^t)$	$1 \frac{1}{3}$	2	$O(1)$
Mult2 (sec. 5.2)	$n - t - 1$	1	$O(n)$	1	1	$O(1)$

Table 1. Summary of proposed multiplication protocols.

4 field elements with $n = 3$. The second, when instantiated with any n , has communication cost quadratic in it (specifically, it is nt), but for $n = 3$ communicates only 3 field elements. It also uses fewer local operations for larger n than our first construction. Our optimizations are tailored to the settings when the number of parties n is not large. Both of our multiplication protocols are secure in the computational setting (as opposed to information-theoretic in [12] in the presence of secure channels). We do not view this as a disadvantage because information-theoretically secure protocols rely on secure channels for communication, which are also built on computational assumptions.

A summary of our proposed multiplication protocols is given in Table 1. Communication refers to average number of field elements transmitted by a parties (i.e., all communication divided by the number of parties) and computation refers to the average work performed by a party including local and communication work. Performance is dominated by communication and round complexity unless local work is excessive.

5.1 Linear-Communication Multiplication

Our starting point was the multiplication protocol from [8] (figure 4 in section 3.3). The high-level structure of the computation is that each participant performs local multiplication of its shares (which makes the degree of the resulting polynomial to become $2t$), sends the result protected by a random element for reconstruction to a dedicated party (called the king). The king performs the reconstruction and announces the result to all other parties who use the opened value to adjust their respective shares. The protocol can be specified as given next and uses two different types of sharings of the same field element, namely, conventional t -sharing of x denoted by $[x]$ and $2t$ -sharing of x denoted by $\langle x \rangle$, i.e., shares are computed using a polynomial of degree $2t$ and at least $2t + 1$ different shares are required for reconstruction of x .

$[c] \leftarrow \text{Mult}([a], [b])$

1. $([r], \langle R \rangle) \leftarrow \text{DRand}()$;
2. Each $p \in [1, n]$ computes $\langle D \rangle_p = [a]_p \cdot [b]_p + \langle R \rangle_p$ and sends $\langle D \rangle_p$ to the king;
3. The king reconstructs $D \leftarrow \text{Open2}(\langle D \rangle)$ and sends D to each party;
4. $[c] = D - [r]$;
5. return $[c]$;

Operation DRand (double random) refers to generation of a random value under two different types of secret sharing: t -sharing and $2t$ -sharing. In other words, the execution of $([r], \langle R \rangle) \leftarrow \text{DRand}()$ produces two different sharings of the same value, i.e., $[r]$ and $\langle R \rangle$ reconstruct to the same field element, but each sharing uses its own randomness. Open2 is similar to Open that reconstructs a value from its shares, but Open2

takes its input represented using $2t$ -sharing and thus requires at least $2t + 1$ shares for reconstruction.

The conventional implementation of Open (or Open2) involves parties sending their shares to others, after which each party reconstructs the value locally using its own and received shares. This requires $O(n^2)$ communication for any $t = O(n)$. However, with the use of a dedicated king, the overall communication can be lowered to $O(n)$, where the value is reconstructed only by the king. To realize Open2 in this way, we need $2t$ participants to communicate their share to the king, who reconstructs the value and consequently communicates it to all other $n - 1$ participants. With $n = 2t + 1$, we obtain $2n - 2 = 4t$ transmitted field elements, which for the (3,1) setting corresponds to communicating the total of 4 elements.

Our main optimization consists of computing double randoms as in DRand non-interactively. While the goal of [8] was to design protocols secure in the stronger, malicious model, even their preliminary construction secure in the semi-honest (passive) security setting was not very cheap. Performing double random generation in a batch of size $\ell = n - t$ required $O(n\ell + n^2)$ communication measured in field elements. We can entirely eliminate this communication by utilizing replicated secret sharing and using computational security.

We start by saying that it is possible to generate pseudo-random $[r]$ non-interactively using RSS as described in Section 3.2. Then if the same key shares are used in a related setup with a threshold set to $2t$, we would be able to non-interactively generate $\langle R \rangle$, where $R = r$. This, however, leads to the use of correlated randomness in the generation of $[r]$ and $\langle R \rangle$, which is not sufficient to provide the necessary security guarantees for our use of these shares. Instead, our approach is as follows: we first generate $[r]$ non-interactively using RSS. To create a $2t$ -sharing of r using fresh randomness, we first raise the degree of r 's secret sharing representation to $2t$ by multiplying it by another degree- t polynomial corresponding to [1]. Lastly, we randomize the resulting shares by adding fresh $\langle 0 \rangle$ to the result. The last step is accomplished by calling the protocol for pseudo-random zero sharing from [7], denoted as PRZS. Luckily, that construction is already given for creating $\langle 0 \rangle$ where the representation uses a polynomial of degree $2t$. We obtain the following construction for DRand that assumes pre-distributed shares k_T and a fixed representation of [1]:

$([r], \langle R \rangle) \leftarrow \text{DRand}()$

1. $[r] \leftarrow \text{PRSS}()$;
2. $\langle 0 \rangle \leftarrow \text{PRZS}()$;
3. Each $p \in [1, n]$ computes $\langle u \rangle_p = [r]_p \cdot [1]_p$;
4. $\langle R \rangle = \langle u \rangle + \langle 0 \rangle$;
5. return $([r], \langle R \rangle)$;

Returning to the multiplication operation, for general n and $t (< n/2)$, the original multiplication uses $n(n - 1)$ messages, while this solution uses $2t + n - 1 \leq 2n - 2$ messages (note that when $n > 2t + 1$, only $2t$ parties and the king execute step 2). For (3, 1)-sharing, the reduction is $6/4 = 1.5$; for (5, 2)-sharing, it is 2.5, and the difference continues to grow with n .

To demonstrate security, we note that we only modified the DRand functionality from that of the multiplication protocol from [8]. Our DRand protocol, however, only

invokes secure building blocks (PRSS and PRZS) and only operates on shares for the remaining computation without disclosing any values. This means that we can easily create a simulator which will not be able to distinguish between the real and simulated (with no access to inputs) views.

5.2 Alternative Multiplication

As mentioned before, we present another multiplication protocol that outperforms the protocol above in terms of communication only when $n = 3$. It is, however, can still be useful for higher values of n because the total work is limited by $O(n)$ per party and does not require the use of replicated secret sharing.

The idea behind this solution is that the parties locally multiply their shares, which, as before, raises the polynomial degree to $2t$ and results in a $2t$ -sharing of the product. To convert the product to a t -sharing, each participant re-shares its value using t -sharing and uses interpolation to compute the result similar to [12]. The difference is that instead of choosing a new random polynomial to do re-sharing, each party uses t pseudo-random points to create the polynomial. These points, together with the party's secret, define the polynomial and allow for the evaluation of the polynomial on other points. Then the pseudo-random points serve the role of the shares of t out of n participants, while the remaining shares are computed by the owner of the secret and are communicated to the remaining parties. The idea is that a pseudo-random value can be generated by two participants without communication and this approach reduced overall communication from $n(n-1)$ to $n(n-t-1)$ field elements, which is a factor of 2 with $n = 2t + 1$. In particular, in the case of $(3, 1)$ secret sharing, we have each party transmitting 1 field element, for the total of 3 field elements and 25% bandwidth reduction compared to the previous multiplication protocol. This also matches best-known 3-party multiplication communication cost based on custom replicated secret sharing arithmetic from [3].

Before we proceed with the algorithm specification, we need to define additional notation. For a secret-shared $[x]$, we let $f_x()$ denote the underlying polynomial according to which the shares of x were computed (i.e., $[x]_p$ corresponds to $f_x(p)$ and $[x]_0 = f_x(0) = x$). We also denote the procedure for reconstructing the polynomial f_x from at least $t + 1$ shares of x by SSReconst_{t+1} . In addition, we let λ_p denote polynomial interpolation constants as defined in [12].

We define mapping γ , which for each participant p defines t other parties with whom p shares PRG seeds for the purpose of non-interactive share computation of its secret. Specifically, for each $\gamma(p, p') = 1$ we let $k_{p,p'}$ be the seed shared by parties p and p' and let $\text{PRG}(k_{p,p'})$.next() denote retrieval of the next field element from the PRG seeded by $k_{p,p'}$. Our multiplication protocol is given next.

$[c] \leftarrow \text{Mult}([a], [b])$

1. Each $p \in [1, n]$ computes $\langle c \rangle_p = [a]_p \cdot [b]_p$;
2. Each $p \in [1, n]$ sets t shares $[d_p]_{p'} \leftarrow \text{PRG}(k_{p,p'})$.next() for each $\{p' \mid \gamma(p, p') = 1\}$ and one more share $[d_p]_0 = \langle c \rangle_p$;
3. Each $p \in [1, n]$ executes $f_{\langle c \rangle_p} \leftarrow \text{SSReconst}_{t+1}([d_p])$;
4. Each $p \in [1, n]$ evaluates $[d_p]_{p'} = f_{\langle c \rangle_p}(p')$ for each $\{p' \mid \gamma(p, p') \neq 1\}$ and sends $[d_p]_{p'}$ to party p' (other than $p' = p$).

5. Each $p \in [1, n]$ computes $[c]_p = \sum_{p'=1}^n \lambda_{p'} [d_{p'}]_p$, where $[d_{p'}]_p$ was either received in step 4 or computed as $[d_{p'}]_p \leftarrow \text{PRG}(k_{p',p}).\text{next}()$ (for $\{p' \mid \gamma(p', p) = 1\}$);
6. return $[c]$;

As discussed before, this protocol communicates $n(n - t - 1)$ field elements across all parties in a single round, and the local work per party is $O(n)$.

Security follows from the fact all computation proceeds on secret-shared values and no intermediate values get revealed. Conceptually this construction follows the structure of the multiplication protocol from [12], where we replace a number of shares to be pseudo-random instead of chosen at random. Thus, while the construction of [12] is secure against unbounded adversaries (assuming secure channels), our security holds in the computational setting.

6 Performance Evaluation

We have implemented the proposed array read and multiplication operations in C. We used the GNU Multiple Precision Arithmetic Library (GMP) [2] for field arithmetic and executed SSS constructions within the PICCO compiler framework [36]. We also execute original array read with private index and multiplication operations as originally implemented in PICCO. All of our protocols are evaluated in the three-party setting with a single corrupt party. For comparison, we also include runtimes of two-party Floram CPRG [9] using their implementation from [1]. This is one of the best performing ORAM constructions (among two- and three-party implementations) and its performance tells us at which array sizes ORAM techniques outperform linear scan. Note that ORAM use might involve additional overhead beyond what we report (e.g., for initializing ORAM or converting between different data representations).

We provide experiments in the LAN and WAN configurations. Our LAN experiments were carried out on identical machines connected via 1Gbps Ethernet with one-way latency of 0.15ms. Each machine was running CentOS 6.9 on a 2.1GHz processor. Our WAN experiments used local and one remote machine running Ubuntu 16 with a 2.4GHz processor. The one-way latency between the remote and local machines was 23ms. All experiments except Floram used a single core. All experiments (except Floram) were executed over a 64-bit finite field and averaged over 100 executions.

Performance of array read is shown in Figures 1 in both LAN and WAN settings. We see that the custom three-party construction significantly outperforms other options. We also see that linear scan constructions outperform ORAM-based solutions for arrays of size up to 2^{16} in the LAN setting and up to 2^{21} in the WAN setting. The figure also shows the difference in the performance of our general array read protocol using the original multiplication protocol as implemented in PICCO (with 6 field elements communicated per multiplication) and the new multiplication protocol from section 5.2 (with 3 field elements per multiplication). The difference between the two multiplication protocols is further detailed in Table 2, which shows that improved multiplication protocol provides up to 25% performance improvement. We further note that a flatter curve in Figure 1 indicates that round complexity or another portion of the computation sub-linear in the array size (as is the case with Floram or linear scans for arrays of small




Fig. 1. Performances of array read with a private index of varying size on a LAN (left) and WAN (right).

Operation \ Batch size	Batch size						
	1	10	10^2	10^3	10^4	10^5	10^6
Original mult. [12]	0.139	0.169	0.521	2.24	23.5	246	2,520
New mult. (sect. 5.2)	0.121	0.144	0.482	1.59	15.5	170	1,720

Table 2. Performance of the original and new multiplication protocols in the (3,1) setting in batches of varying sizes in milliseconds.

sizes in the WAN setting) is the bottleneck. A steeper curve indicates that work linear in the array size (e.g., $O(m)$ communication in the case of linear scans) is the bottleneck.

We also provide measurement results for parallel execution of array read in Table 3. We compare the original PICCO multiplexer-based implementation with our new general array read with new multiplication from section 5.2. Substantial runtime reduction over single execution is observed for arrays of relatively small size. The largest difference between the original and our solution is by a factor of 16 with array size of 2^{16} and batch size of 10.

7 Conclusions

In this work we study performance improvements to certain common building blocks in secure multi-party computation based on secret sharing. We present optimized protocols for reading or writing an element of an array at a private index and for integer multiplication. Most of our constructions are based on Shamir secret sharing with the exception of one array access construction. The latter uses 2-out-of-2 additive secret sharing in the three-party setting with honest majority, but offer superior performance compared to general constructions. To be compatible with computation based on Shamir secret sharing, we also provide conversion procedures to convert between the two representations. We implement the presented constructions in the setting with three computational parties and show that they offer attractive performance in both LAN and WAN settings.

Batch size \ Array size	1	10	10 ²	10 ³
2 ⁴	0.00087 / 0.00217	0.00206 / 0.00582	0.00946 / 0.0253	0.0958 / 0.243
2 ⁷	0.00183 / 0.00846	0.00711 / 0.028	0.0441 / 0.257	0.463 / 2.33
2 ¹⁰	0.00491 / 0.0285	0.0281 / 0.28	0.289 / 2.88	2.98 / 27.2
2 ¹³	0.0218 / 0.269	0.219 / 2.77	2.21 / 28.8	22.5 / 276
2 ¹⁶	0.174 / 2.67	1.75 / 27.8	17.6 / 267	180 / 2,689

Table 3. Performance of the multiplexer-based (lower left) and our general (upper right) array read with a private index in the (3, 1) setting for batches of varying size in seconds.

References

1. The floram oblivious ram implementation for secure computation. <https://gitlab.com/neucrypt/floram/tree/floram-release>. [Online; accessed 20-Nov-2018].
2. The GNU multiple precision arithmetic library. <https://gmplib.org/>.
3. T. Araki, J. Furukawa, Y. Lindell, A. Nof, and K. Ohara. High-throughput semi-honest secure three-party computation with an honest majority. In *ACM Conference on Computer and Communications Security (CCS)*, pages 805–817, 2016.
4. F. Bayatbabolghani, M. Blanton, M. Aliasgari, and M. Goodrich. Secure fingerprint alignment and matching protocols. *arXiv Report 1702.03379*, 2017.
5. P. Bunn, J. Katz, E. Kushilevitz, and R. Ostrovsky. Efficient 3-party distributed ORAM. *Cryptology ePrint Archive Report 2018/706*, 2018.
6. O. Catrina and S. de Hoogh. Improved primitives for secure multiparty integer computation. In *Security and Cryptography for Networks*, 2010.
7. R. Cramer, I. Damgård, and Y. Ishai. Share conversion, pseudorandom secret-sharing and applications to secure computation. In *Theory of Cryptography Conference (TCC)*, pages 342–362, 2005.
8. I. Damgård and J. Nielsen. Scalable and unconditionally secure multiparty computation. In *Advances in Cryptology – CRYPTO*, pages 572–590, 2007.
9. J. Doerner and A. Shelat. Scaling ORAM for secure computation. In *ACM Conference on Computer and Communications Security (CCS)*, pages 523–535, 2017.
10. S. Faber, S. Jarecki, S. Kentros, and B. Wei. Three-party ORAM for secure computation. In *ASIACRYPT*, pages 360–385, 2015.
11. C. W. Fletcher, M. Naveed, L. Ren, E. Shi, and E. Stefanov. Bucket ORAM: Single online roundtrip, constant bandwidth oblivious RAM. *IACR Cryptology ePrint Archive*, 2015:1065, 2015.
12. R. Gennaro, M. Rabin, and T. Rabin. Simplified VSS and fast-track multiparty computations with applications to threshold cryptography. In *PODC*, pages 101–111, 1998.
13. O. Goldreich. Towards a theory of software protection and simulation by oblivious rams. In *ACM Symposium on Theory of Computing (STOC)*, pages 182–194, 1987.
14. O. Goldreich and R. Ostrovsky. Software protection and simulation on oblivious RAMs. *Journal of the ACM (JACM)*, 43(3):431–473, 1996.

15. T. Hoang, C. D. Ozkaptan, A. A. Yavuz, J. Guajardo, and T. Nguyen. S^3 ORAM: A computation-efficient and constant client bandwidth blowup ORAM with Shamir secret sharing. In *ACM Conference on Computer and Communications Security (CCS)*, pages 491–505, 2017.
16. M. Ito, A. Saito, and T. Nishizeki. Secret sharing schemes realizing general access structures. In *IEEE Global Telecommunication Conference (Globecom)*, pages 99–102, 1987.
17. S. Jarecki and B. Wei. 3PC ORAM with low latency, low bandwidth, and fast batch retrieval. In *Applied Cryptography and Network Security (ANCS)*, 2018.
18. S. Karan and J. Zola. Scalable exact parent sets identification in Bayesian networks learning with Apache Spark. In *IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC)*, pages 33–41, 2017.
19. M. Keller and P. Scholl. Efficient, oblivious data structures for MPC. In *ASIACRYPT*, pages 506–525, 2014.
20. M. Keller and A. Yanai. Efficient maliciously secure multiparty computation for RAM. In *EUROCRYPT*, pages 91–124, 2018.
21. L. Kerik, P. Laud, and J. Randmets. Optimizing MPC for robust and scalable integer and floating-point arithmetic. In *International Conference on Financial Cryptography and Data Security Workshops*, pages 271–287, 2016.
22. M. Koivisto. Parent assignment is hard for the MDL, AIC, and NML costs. In *International Conference on Computational Learning Theory*, pages 289–303, 2006.
23. P. Laud. A private lookup protocol with low online complexity for secure multiparty computation. In *International Conference on Information and Communications Security (ICICS)*, pages 143–157, 2014.
24. P. Laud. Parallel oblivious array access for secure multiparty computation and privacy-preserving minimum spanning trees. *PoPETs*, 2015(2):188–205, 2015.
25. R. Ostrovsky. Efficient computation on oblivious RAMs. In *ACM Symposium on Theory of Computing (STOC)*, pages 514–523, 1990.
26. L. Ren, C. W. Fletcher, A. Kwon, E. Stefanov, E. Shi, M. van Dijk, and S. Devadas. Ring ORAM: Closing the gap between small and large client storage oblivious ram. *IACR Cryptology ePrint Archive*, 2014:997, 2014.
27. G. Schwarz. *The Annals of Statistics*, 6:461–464, 1978.
28. A. Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979.
29. E. Shi, T-H. H. Chan, E. Stefanov, and M. Li. Oblivious RAM with $O((\log N)^3)$ worst-case cost. In *ASIACRYPT*, pages 197–214, 2011.
30. E. Stefanov and E. Shi. Multi-cloud oblivious storage. In *ACM Conference on Computer and Communications Security (CCS)*, pages 247–258, 2013.
31. E. Stefanov, E. Shi, and D. Song. Towards practical oblivious RAM. *arXiv preprint arXiv:1106.3652*, 2011.
32. E. Stefanov, M. Van Dijk, E. Shi, C. Fletcher, L. Ren, X. Yu, and S. Devadas. Path ORAM: An extremely simple oblivious RAM protocol. In *ACM Conference on Computer and Communications Security (CCS)*, pages 299–310, 2013.
33. X. Wang, H. Chan, and E. Shi. Circuit ORAM: On tightness of the Goldreich-Ostrovsky lower bound. In *ACM Conference on Computer and Communications Security (CCS)*, pages 850–861, 2015.
34. X. Wang, Y. Huang, TH H. Chan, A. Shelat, and E. Shi. SCORAM: oblivious ram for secure computation. In *ACM Conference on Computer and Communications Security (CCS)*, pages 191–202, 2014.
35. S. Zahur, X. Wang, M. Raykova, A. Gascón, J. Doerner, D. Evans, and J. Katz. Revisiting square root ORAM: Efficient random access in multi-party computation. In *IEEE Symposium on Security and Privacy*, pages 218–234, 2016.

36. Y. Zhang, A. Steele, and M. Blanton. PICCO: A general-purpose compiler for private distributed computation. In *CCS*, pages 813–826, 2013.

A Security Definitions

The security definition that we adopt in this work is as follows:

Definition 1. Let parties P_1, \dots, P_n engage in a protocol Π that computes function $f(\text{in}_1, \dots, \text{in}_n) = (\text{out}_1, \dots, \text{out}_n)$, where in_i and out_i denote the input and output of party P_i , respectively. Let $\text{VIEW}_\Pi(P_i)$ denote the view of participant P_i during the execution of protocol Π . More precisely, P_i 's view is formed by its input and internal random coin tosses r_i , as well as messages m_1, \dots, m_k passed between the parties during protocol execution: $\text{VIEW}_\Pi(P_i) = (\text{in}_i, r_i, m_1, \dots, m_k)$. Let $I = \{P_{i_1}, P_{i_2}, \dots, P_{i_t}\}$ denote a subset of the participants for $t < n$, $\text{VIEW}_\Pi(I)$ denote the combined view of participants in I during the execution of protocol Π (i.e., the union of the views of the participants in I), and $f_I(\text{in}_1, \dots, \text{in}_n)$ denote the projection of $f(\text{in}_1, \dots, \text{in}_n)$ on the coordinates in I (i.e., $f_I(\text{in}_1, \dots, \text{in}_n)$ consists of the i_1 th, \dots , i_t th element that $f(\text{in}_1, \dots, \text{in}_n)$ outputs). We say that protocol Π is t -private in the presence of semi-honest adversaries if for each coalition of size at most t there exists a probabilistic polynomial time (PPT) simulator S_I such that $\{S_I(\text{in}_I, f_I(\text{in}_1, \dots, \text{in}_n)), f(\text{in}_1, \dots, \text{in}_n)\} \equiv \{\text{VIEW}_\Pi(I), (\text{out}_1, \dots, \text{out}_n)\}$, where $\text{in}_I = \bigcup_{P_i \in I} \{\text{in}_i\}$ and \equiv denotes computational or statistical indistinguishability.

B Secret Sharing Conversion

In this section we provide conversion procedures between Shamir SS and 2-out-of-2 additive secret sharing in the three-party setting. It is assumed that both use shares over the same finite field \mathbb{F} .

We start from the SSS to additive SS conversion, which proceeds as follows:

- $\llbracket a \rrbracket_{12} \leftarrow \text{S2A}(\llbracket a \rrbracket)$
1. $\llbracket r \rrbracket \leftarrow \text{PRSS}()$;
 2. $\llbracket d \rrbracket \leftarrow \llbracket a \rrbracket - \llbracket r \rrbracket$;
 3. Open r to party 1 who sets $\llbracket a \rrbracket_1 = r$;
 4. Open d to party 2 who sets $\llbracket a \rrbracket_2 = d$;
 5. return $\llbracket a \rrbracket_{12}$;

If we assume that PRSS can be realized non-interactively as before, reconstructing r in step 3 and d in step 4 involves communicating one field element each. That is, party 2 or 3 sends its share of r to party 1 in step 3, from which party 1 recovers r . Thus, the cost is communicating 2 field elements in 1 round.

The conversion from additive SS to SSS is as follows:

- $\llbracket a \rrbracket \leftarrow \text{A2S}(\llbracket a \rrbracket_{12})$
1. Party 1 creates Shamir secret shares of $\llbracket a \rrbracket_1$ and distributes them among the parties;
 2. Party 2 creates Shamir secret shares of $\llbracket a \rrbracket_2$ and distributes them among the parties;
 3. $\llbracket a \rrbracket = \llbracket \llbracket a \rrbracket_1 \rrbracket + \llbracket \llbracket a \rrbracket_2 \rrbracket$;

4. return $[a]$;

Steps 1 and 2 can be accomplished by communicating a single field element each in the computational setting similar to the approach taken in section 5.2. That is, party 1 shares a secret key with another party for generating that party's share as a pseudo-random value. The remaining shares are computed to be consistent with the pseudo-random share and the value of $[[a]]_1$, which requires communication of a single share. Thus, the protocol's cost is communicating 2 field elements in 1 round.