



# MP-SPDZ: A Versatile Framework for Multi-Party Computation

Marcel Keller

CSIRO's Data61

Australia

marcel.keller@data61.csiro.au

mks.keller@gmail.com

## ABSTRACT

Multi-Protocol SPDZ (MP-SPDZ) is a fork of SPDZ-2 (Keller et al., CCS '13), an implementation of the multi-party computation (MPC) protocol called SPDZ (Damgård et al., Crypto '12). MP-SPDZ extends SPDZ-2 to 30 MPC protocol variants, all of which can be used with the same high-level programming interface based on Python. This considerably simplifies comparing the cost of different protocols and security models.

The protocols cover all commonly used security models (honest/dishonest majority and semi-honest/malicious corruption) as well as computation of binary and arithmetic circuits (the latter modulo primes and powers of two). The underlying primitives employed include secret sharing, oblivious transfer, homomorphic encryption, and garbled circuits.

The breadth of implemented protocols coupled with an accessible high-level interface makes it suitable to benchmark the cost of computation in various security models for researchers both with and without a background in secure computation

This paper aims to outline the variety of protocols implemented and the design choices made in the development of MP-SPDZ as well as the capabilities of the programming interface.

## CCS CONCEPTS

• Security and privacy → Cryptography; Privacy-preserving protocols.

## KEYWORDS

Multi-party computation; implementation

### ACM Reference Format:

Marcel Keller. 2020. MP-SPDZ: A Versatile Framework for Multi-Party Computation. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security (CCS '20)*, November 9–13, 2020, Virtual Event, USA. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3372297.3417872>

## 1 INTRODUCTION

Multi-party computation allows a set of parties to compute collaboratively on their private inputs without revealing anything other than the result. A range of applications have been suggested such

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

CCS '20, November 9–13, 2020, Virtual Event, USA

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-7089-9/20/11...\$15.00

<https://doi.org/10.1145/3372297.3417872>

as truthful auctions [15], avoiding satellite collisions [44], computing the gender pay gap [57], or privacy-preserving machine learning [62]. After the development of the basic theory in the 1980s [10, 21], the first implementation in software was created by Pinkas et al. [66] in 2009. Since then, a number of frameworks have been created for a range of security models and computation domains. By security model we mean the assumptions on the behavior of the parties, in particular how many are assumed to behave honestly and whether dishonest parties follow the protocol or try to gain information or compromise the result by deviating. On the other hand, the computation domain denotes a mathematical structure that is used to represent secret information in the computation as required by the underlying schemes. In applications, this domain usually takes the form of a ring, that is, a set with the two operations akin to addition and multiplication. Often said ring is simply defined by integer operations with modulus, but there are other examples such as the Galois fields of characteristic two (e.g., the one defining the arithmetic structure of AES).

Almost all openly available frameworks for multi-party computation are restricted to a particular security model and computation domain, which makes it hard to assess the cost of one security model compared to another because one has to implement the same computation several times. Similarly, it is hard to compare protocols in the same security model. MP-SPDZ<sup>1</sup> aims to change this by providing an implementation of 30 protocol variants with the same virtual machine, and a compiler that compiles high-level code into bytecode to be executed by said virtual machine. This allows one to implement a computation only once before benchmarking it in a variety of settings.

At the core of the approach of combining many protocols in one framework lies the intuition that, despite the differences, all commonly used protocols for secure computation can largely be reduced to a few operations that are similar in all of them, and these are input, output, locally computable linear operations, one much more involved operation like AND or multiplication, and generating correlated randomness such as random bits. It therefore seems natural to build a general framework facilitating this blueprint in order to reuse components and optimizations as much as possible. MP-SPDZ does so while still allowing specific protocols to be added if necessary.

MP-SPDZ is based on SPDZ-2 [51, 53], an implementation of the SPDZ protocol [28, 31]. The infrastructure of SPDZ-2 has been used and extended in a line works [27, 47–50, 52, 55]. Araki et al. [4] have undertaken steps to integrate protocols other than SPDZ in SPDZ-2. However, their approach is marred by overly relying on library calls for basic operations such as share additions, and thus

<sup>1</sup><https://github.com/data61/mp-spdz>

```

a = sint.Array(n)
b = sint.Array(n)
a.input_from(0)
b.input_from(1)
res = sint.dot_product(a, b)
print_ln('%s', res.reveal())

```

**Figure 1: Inner product computation in MP-SPDZ.**

it is relatively inefficient. This can be seen by how well SPDZ fares in comparison to simpler honest-majority protocols (Figures 6-8 by Araki et al.) because the former is natively implemented in SPDZ-2. MP-SPDZ is not based on the work by Araki et al. but directly forked from SPDZ-2 instead.

Hastings et al. [41] have assessed eleven multi-party computation frameworks available in the first half of 2018. They estimate that creating and running three simple example programs in all of the frameworks together took them 750 person-hours. MP-SPDZ aims to lower the entry barrier to secure computation by providing standalone Linux binaries for every release and extensive documentation for the high-level language.<sup>2</sup> In addition, it uses a continuous integration facility, and issues filed on GitHub are usually attended within a few days.

*Paper organization.* After introducing similar frameworks (with comparison benchmarks) and basic concepts, we will proceed layer by layer, from the underlying protocols to the high-level library.

As an anchor throughout the paper, we will use the example of computing an inner product of two vector supplied to by two different parties. Hastings et al. have used this as one of their examples. The code for MP-SPDZ is shown in Figure 1.

## 1.1 Comparison to Other Frameworks

In the following we will consider all frameworks described by Hastings et al. and the ones subsequently added to their repository containing example programs and build environments [42] at the time of writing. Where available we also indicate the approximate size of the inner product example. Table 1 highlights how MP-SPDZ compares to other frameworks. SH stands for semi-honest security. See Section 1.2 for details.

*ABY.* This framework only implements two-party computation with semi-honest security [32]. Unlike MP-SPDZ it provides conversion between secret-sharing computation and garbled circuits, however. The inner product example takes about 60 lines of code without comments or empty lines.

*ABY<sup>3</sup>.* This framework only implements three-party computation with honest majority and semi-honest security [61]. The inner product example takes about 40 lines of code without comments or empty lines.

*CBMC-GC.* This is a compiler that compiles C code to a binary circuit description [18] to be executed by ABY. Hastings et al. were unable to make it work with their examples.

*EMP-toolkit.* This framework only implements garbled circuits in various security models [74]. If not using Yao’s garbled circuits with semi-honest security, loops are unrolled at compile time. The inner product example takes about 60 lines of code without comments or empty lines.

*FRESCO.* This framework [1] only implements dishonest majority computation, with malicious security for arithmetic circuits (SPDZ and  $\text{SPDZ}_{2k}$ ) and with semi-honest security for binary circuits [29]. The inner product example takes about 30 lines of code without comments or empty lines.

*Frigate.* This is a compiler that compiles C-like code to a binary circuit description [64]. Unlike with MP-SPDZ, loops are always unrolled at compile-time. The inner product example takes about 20 lines of code without comments or empty lines.

*JIFF.* This JavaScript library only implements honest-majority computation with semi-honest security as a whole [73]. Unlike MP-SPDZ it allows changing the security model between offline and online phase. The inner product example takes about 50 lines of code without comments or empty lines.

*MPyC.* This Python framework [69] only implements computation with semi-honest security based on Shamir’s secret sharing [70]. The inner product example takes about 20 lines of code without comments or empty lines.

*Obliv-C.* This framework compiles an extension of C to machine code via standard C [76]. It only supports Yao’s garbled circuits with semi-honest security. The inner product example takes about 20 lines of code without comments or empty lines.

*OblivVM.* This framework compiles an extension of Java to Java bytecode [60]. It only supports Yao’s garbled circuits with semi-honest security. The inner product example takes about 20 lines of code without comments or empty lines.

*PICCO.* This framework compiles an extension of C to native binaries via standard C [78]. It only implements semi-honest computation with an honest majority based on Shamir’s secret sharing. The inner product example takes about 10 lines of code without comments or empty lines.

*SCALE-MAMBA.* This framework [22] is another fork of SPDZ-2 [51, 53]. Despite the common roots, the two forks have diverged considerably since 2018. SCALE-MAMBA only implements arithmetic computation modulo a prime (not modulo a power of two), garbled circuits according to Hazay et al. [43], and binary computation based on secret sharing [36, 75]. All computations are implemented with malicious security only, and dishonest majority computation modulo a prime is only implemented using homomorphic encryption. On the other hand, SCALE-MAMBA implements honest-majority computation for any access structure possible in theory.

The frontend is similar to the one in MP-SPDZ but without later additions such as the dynamic loop optimization (Section 6.3), the repetitive code optimization (Section 6.4), and machine learning functionality (Section 7.3). Furthermore, the authors have started to move away from the Python compiler in favor of a new one based on Rust.

<sup>2</sup><https://mp-spdz.readthedocs.org>

**Table 1: Time in seconds to compute an inner product of two 100,000-element vectors of 64-bit integers.**

	SH 2-party (OT)	SH replicated 3-party	Shamir	Malicious Shamir	SPDZ (HighGear)	Yao’s GC (32-bit)
ABY	1.8	⊥	⊥	⊥	⊥	⊥
ABY <sup>3</sup>	⊥	0.02	⊥	⊥	⊥	⊥
EMP-toolkit	⊥	⊥	⊥	⊥	⊥	10
MPyC	⊥	⊥	13.21	⊥	⊥	⊥
Obliv-C	⊥	⊥	⊥	⊥	⊥	29
ObliVM	⊥	⊥	⊥	⊥	⊥	700
PICCO	⊥	⊥	0.10	⊥	⊥	⊥
SCALE-MAMBA	⊥	⊥	⊥	8.3	314	⊥
MP-SPDZ	0.9	0.03	0.08	0.4	67	9

The inner product example takes less than 10 lines of code without comments or empty lines.

*Sharemind MPC.* This framework implements a frontend for a variety of backends but its own backend only uses three-party semi-honest computation with an honest majority [14]. It also allows the use of ABY and FRESCO as backend while the proprietary backend is not freely available. The inner product example takes less than 10 lines of code without comments or empty lines.

*TinyGarble.* This framework only implements Yao’s garbled circuits with semi-honest security [71]. Hastings et al. were unable to make it work with their examples.

*Wysteria.* This framework implements a domain-specific language with only binary computation in the semi-honest setting with dishonest majority [67]. Hastings et al. were unable to run all their examples in this framework.

## 1.2 Benchmarks

We have benchmarked MP-SPDZ against other frameworks listed in the previous section.<sup>3</sup> Table 1 shows how many seconds it takes to compute an inner product of length 100,000 on one machine with a 7-th generation i7 processor (baseline frequency 2.8 GHz). We have opted for a local computation because not all frameworks support execution on different hosts.

For benchmarking, we have replaced `get_input_from()` with `get_raw_input_from()` in the MP-SPDZ code. This is because the former leads to inputs being read using the `istream` functionality of the C++ standard library, which considerably increases the time in some benchmarks. The examples in other frameworks either use faster functions like `atoi()` or the inputs are hard-coded.

Furthermore, the timings for computation modulo a prime (Shamir and SPDZ) have been done using a 128-bit prime because that is a common choice given that many algorithms require a  $(k + s)$ -bit prime for  $k$ -bit computation and security parameters  $s$ . We have reduced all statistical security parameters in SCALE-MAMBA to 40 for a fairer comparison. Similarly, the HighGear figure excludes key generation because the two frameworks handle this differently.

The table shows that the more sophisticated approach using a virtual machine does not degrade performance considerably with arithmetic computation.

<sup>3</sup>The code can be found at <https://github.com/mkskeller/mpc-benchmarks>.

In addition, we have benchmarked frameworks using Yao’s garbled circuits but with 32-bit integers instead of 64 because some examples only support the former.

Note that ⊥ means that a particular framework does not implement a particular protocol, except for the case of ABY and Yao where there is no example code available.

We have identified a number of reasons for the frameworks performing significantly worse. MPyC and OblivM are not implemented with C/C++. For SCALE-MAMBA, we have found that it utilizes frequent allocation/free calls due to a regular usage of variable-length integers, and it uses a high number of relatively expensive function calls to MPIR for fixed-length arithmetic. Furthermore, it also generates more preprocessing data than necessary.

Finally, Tables 5 and 6 in the appendix show our figures for the framework in a simulated WAN setting (100 Mbit/s and 100 ms RTT) as well as the total communication. The latter was obtained using `ip -s link lo` in Docker for all other frameworks.

*Missing frameworks.* We do not provide benchmarks for various frameworks mentioned in Section 1.1 for the following reasons:

- No working inner product example: CBMC-GC, TinyGarble, Wysteria
- No backend executing the actual computation: Frigate
- Inner product example incomplete (without preprocessing): FRESCO
- Using JavaScript: JIFF
- Software not available: Sharemind MPC

## 2 PRELIMINARIES

In this section we explain various basic concepts in secure multi-party computation.

*Security models.* A core property of multi-party computation protocols is what assumptions are made about the parties, which happens along two axes. The first question is how many parties are assumed to behave honestly and how many are assumed to be “corrupted” in some way. The protocols in MP-SPDZ all assume threshold corruption, that is a constant number of parties are thought to be corrupted. There is a considerable difference between an honest majority and a dishonest majority, that is, whether this threshold is strictly below half the number of parties or not. The protocols for a dishonest majority are generally more intricate and more expensive.

The second question is how corrupted parties behave. The two main categories here are whether they still follow the protocol but collude on extracting information (called semi-honest or passive corruption) or whether they also deviate from the protocol in order to gain information or distort the result of the computation (called malicious or active corruption). The latter setting raises further questions as to whether the computation can still continue in case of deviation (guaranteed output) and whether the corrupted parties can learn the result without honest parties doing so (fairness). All protocols in MP-SPDZ with malicious security work in the malicious-with-abort model, which means that deviation is detected but the protocol cannot recover after that. The reason is that fairness is relatively expensive, and guaranteed output is outright impossible with dishonest majority. Consider the case of a two-party protocol: One party cannot continue when the other stops communicating.

*Computation domains.* There is a variety of mathematical structures used in multi-party computation. Often this takes the form of integers modulo some number  $M$ . If  $M = 2$  the computation reduces to binary circuits because addition then corresponds to XOR and multiplication corresponds to AND. For larger  $M$ , the literature has established the term arithmetic circuits because the basic operations provided remain addition and multiplication. If  $M$  is a prime, the domain satisfies the definition of a field, that is, all numbers except zero have a multiplicative inverse. This is required for a number of schemes including Shamir's secret sharing [70] and SPDZ [31]. For the latter, these requirements were later adapted for computation modulo a power  $2^k$ , resulting in the SPDZ $_{2^k}$  protocol. However, integer computation modulo a prime is not the only example of a finite field. MP-SPDZ also implements computation in  $\mathbb{F}_{2^k}$ . The latter domain particularly has applications for symmetric cryptosystems such as AES, which is based on arithmetic in  $\mathbb{F}_{2^8}$ .

*Secret sharing.* This concept of distributing information is at the core of many multi-party computation protocols. Information is distributed in a way that allows some sets (called qualified) of parties to reconstruct it while some smaller sets (called unqualified) cannot deduce anything from the shares they are given. The exact definition of which subsets can and cannot recover is called an access structure. Clearly, a superset of a qualified set has to be qualified as well, and a subset of an unqualified set has to be unqualified for the access structure to make sense. Furthermore, the common definition of secret sharing requires that all subsets of parties are either qualified or unqualified.

As an example consider additive secret sharing. A number  $x$  in a range  $[0, M - 1]$  is shared among parties  $P_1, \dots, P_n$  by sending a random number  $x_i \in [0, M - 1]$  to  $P_i$  such that  $\sum_i x_i = x \pmod{M}$ . It is easily seen that the all parties together can reconstruct  $x$  while the view of any strict subset is indistinguishable to a random set of numbers. This secret sharing scheme is used in all protocols with dishonest majority except schemes based on garbling.

Other secret sharing schemes used in MP-SPDZ include replicated secret sharing [11] where one starts with additive secret sharing but sends more than one share to every party and Shamir's secret sharing [70] where shares are determined using a random

polynomial of fixed degree equal to the maximum number of corrupted parties.

All of these secret sharing schemes are linear, that is, any (affine) linear combination of secret-shared values can be computed locally because the reconstruction is linear as well.

*Beaver's multiplication.* This technique [7] reduces multiplication of secret numbers to multiplication of secret random numbers, which is useful both in terms of security and practicality. First, it allows executing the multiplication of secret random numbers optimistically before checking whether the parties followed the protocol. If the check fails the protocol can abort without consequences because the secret randomness is independent of the actual secret data. Second, Beaver's multiplication enables batch-wise preprocessing even if the actual computation is sequential, which is a particular advantage when using lattice-based cryptosystems that (only) allow to encrypt many values at once efficiently.

The multiplication works as follows: Let  $[x]$  denote a secret sharing of  $x$ , and assume that a triple  $([a], [b], [ab])$  for independent random  $a$  and  $b$  is available. The multiplication of  $[x]$  and  $[y]$  can then be computed after revealing  $x + a$  and  $y + b$ . Revealing these does not reveal  $x$  or  $y$  because  $a$  and  $b$  are secret and randomly chosen. After revealing the sums, the parties compute

$$\begin{aligned} [xy] &= [(x + a - a) \cdot (y + b - b)] \\ &= (x + a) \cdot (y + b) - (x + a) \cdot [b] - (y + b) \cdot [a] + [ab]. \end{aligned}$$

Note that this computation is affine linear with respect to the secret values and thus it is possible to compute it with any linear secret sharing scheme.

*Oblivious transfer.* This is a basic two-party cryptographic protocol that was the first to be used for dishonest-majority computation. Essentially, one party (called sender) inputs two strings  $s_0, s_1$ , and the other party (called receiver) inputs a bit  $b$  and learns  $s_b$ . Crucially, the sender does not learn  $b$ , and the receiver does not learn  $s_{1-b}$ . It is relatively straightforward to construct a protocol from oblivious transfer that allows two parties to compute a secret sharing of the product of private numbers known to each without revealing the inputs. The construction is leveraged by protocols such as MASCOT [49]. Oblivious transfer is only known to be constructed from public-key cryptography. However, OT extension [45] can be used keep the use of such more expensive schemes at a low constant while using more symmetric cryptography instead.

*Homomorphic encryption.* This refers to encryption schemes that enable operating on ciphertexts efficiently in a way that implies some operation without revealing the cleartext. While encryption schemes with limited homomorphism have been known for a while (textbook RSA is homomorphic in some sense), only the emergence of lattice-based cryptography led to schemes that are homomorphic with respect to two operations. MP-SPDZ uses the leveled encryption scheme by Brakerski et al. [16], which requires that a vector of numbers are encrypted at once. The minimum length depends on the plaintext modulus but generally ranges from several thousand to several ten thousand. The concrete implementation is based on the one by Gentry et al. [39], which is defined by the use of Montgomery representation [63] and fast Fourier transform. The first facilitates multiplication modulo a prime without expensive

modular reduction, and the second is an efficient method for converting a list of numbers to a polynomial in the plaintext space such that the multiplication of two such polynomials corresponds to an element-wise product of the numbers. This approach imposes a restriction on the prime modulus that can be used, but many applications requiring only general integer-like computation are indifferent to the actual prime.

### 3 THE PROTOCOLS

Table 2 shows which protocol variants in each combination of security model and computation domain are implemented in MP-SPDZ. In the following sections we will briefly describe them.

#### 3.1 Dishonest Majority

All protocols for dishonest majority use techniques related to public-key cryptography, either oblivious transfer of homomorphic encryption. Research on practical protocols in this line was started by Nielsen et al. [65] and Damgård et al. [31]. The former proposed a two-party protocol computing binary circuits based on oblivious transfer (later dubbed TinyOT), and the latter proposed a multi-party protocol for computation in fields (modulo a prime or  $\mathbb{F}_{2^k}$ ) using homomorphic encryption (named SPDZ after the authors). Both use Beaver’s technique of dividing the computation in a data-independent and a data-dependent phase [7]. The former (sometimes called “offline”) outputs correlated information that is used by latter together with the actual private data.

*MASCOT.* This protocol denotes a way of computing Beaver triples used in SPDZ using oblivious transfer (OT) with malicious security [49], which was the first offline phase released in SPDZ-2. A core technology here is the use of so-called OT extension with malicious security [48], which considerably increases the throughput because public-key cryptography is only used briefly at the beginning of the computation.

*SPDZ<sub>2<sup>k</sup></sub>.* MASCOT has been adapted to computation modulo a power of two [23, 27]. The main challenge here is that not every element in  $\mathbb{Z}_{2^k}$  has an inverse, which is a crucial tool in proving the security of MASCOT (and SPDZ). SPDZ<sub>2<sup>k</sup></sub> mitigates this by moving to  $\mathbb{Z}_{2^{k'}}$  for a larger  $k'$  in order to counter the effect of zero divisors. MP-SPDZ fully implements SPDZ<sub>2<sup>k</sup></sub> using its own implementation of  $\mathbb{Z}_{2^k}$  optimized with compile-time  $k$ .

*SPDZ and Overdrive.* LowGear and HighGear [50] are the names of two protocols computing Beaver triples for SPDZ using semi-homomorphic and somewhat homomorphic encryption, respectively. They were part of SPDZ-2. Semi-homomorphic means that it is possible to add ciphertexts and multiply ciphertexts with clear-texts in order to obtain an encryption of the sum and the product, respectively. By somewhat homomorphic we mean that in addition it is also possible to multiply two ciphertexts but the result cannot be further multiplied by a ciphertext. All known homomorphic cryptosystems incur a considerable performance penalty for every additional level of ciphertext multiplication. The core idea of SPDZ, LowGear, and HighGear is to minimize the number of such multiplication levels using multi-party computation techniques. LowGear and HighGear represent a trade-off in that LowGear runs a sub-protocol between all pairs of parties (similar to Bendlin et

al. [12]) and therefore does not scale as well with the number of parties as HighGear.

Since MP-SPDZ does not implement the key generation needed for LowGear and HighGear with malicious security, variants of them using covertly secure key generation account for the two covertly secure protocols in Table 2. Optionally, the TopGear zero-knowledge proof [6] can be used instead of the original variants.

MP-SPDZ also retains the offline phase of earlier variants of the SPDZ protocol, with malicious [31] and covert security [28]. Since they are superseded by HighGear, they are not integrated with the online phase and therefore do not appear in Table 2, however.

*Binary secret sharing.* MP-SPDZ supports two ways of computing binary circuits with secret sharing, dishonest majority, and malicious security. The first is adapted from SPDZ<sub>2<sup>k</sup></sub> with  $k = 1$ . However, this variant suffers from the fact that communication is quadratic in the security parameters. Even though that also holds for SPDZ<sub>2<sup>k</sup></sub> but has less of impact if  $k$  is larger than the security parameter as is the case for SPDZ<sub>2<sup>k</sup></sub> because the communication of protocols based on oblivious transfer is in  $O(k^2)$ .

The second variant is based on a multi-party generalization of TinyOT by Frederiksen et al. [36] and has communication linear in the security parameter.

*BMR.* Beaver et al. [8] have presented a way to construct garbled circuits from any multi-party computation scheme inheriting the security properties. Their approach was later refined by Lindell et al. [59] by using SPDZ as the underlying protocol. MP-SPDZ implements BMR using SPDZ/MASCOT as well as protocols in other security models. This functionality was never part of SPDZ-2 but has been partially published before the first version of MP-SPDZ because it was used by Keller and Yanai to implement oblivious RAM [55].

*Yao’s Garbled Circuits.* Bellare et al. [9] have presented a variant of Yao’s garbled circuits that works particularly well with AES-NI, the native implementation of AES on modern processors. An implementation was added after the last version of SPDZ-2, and it was recently extended to include the half-gate technique [77].

*Semi-honest security.* It is relatively straight-forward to convert a protocol with malicious security to one with semi-honest security by removing all aspects that solely contribute to malicious security. All protocols in the line of SPDZ and TinyOT feature two main elements for this: an information-theoretic tag (called “MAC” in most of the literature) and a procedure called sacrificing that checks the correctness of correlated randomness using more correlated randomness. The latter has to be discarded thereafter because the procedure reveals correlation between the two parts, hence the name sacrificing. Stripping both leaves a protocol that comes close to canonical semi-honest protocol using the underlying techniques (oblivious transfer or homomorphic encryption). The main difference is that in the semi-honest setting one could use OT or HE directly with secret data instead of produced correlated randomness. Given the cost of OT or HE, the overhead is relatively small compared to the additional cost of Beaver’s multiplication, however. Furthermore, using correlated randomness with homomorphic encryption can easily work with the SIMD structure of LWE-based schemes, that is the fact that such schemes only can

**Table 2: Protocols implemented in MP-SPDZ. ‘-’ indicates removing parts needed only for malicious security, ‘\*\*’ denotes other modifications, and numbers in brackets stand for the number of protocol variants. All protocols in the first column support both types of fields.**

Security model	Modulo prime / Galois field	Modulo power of two	Binary sharing	Garbling
Malicious, dishonest majority	[KOS16]	[DEF <sup>+</sup> 19]	[DEF <sup>+</sup> 19]* [FKOS15]	[LPSY15]
Covert, dishonest majority	[KPR18]* (2)			
Semi-honest, dishonest majority	[KOS16]- [KPR18]- (2)	[KOS16]-	[FKOS15]-	[LPSY15] [ZRE15]
Malicious, honest majority	[LN17] (3)	[EKO <sup>+</sup> 20] (4)	[FLNW17] [LN17]	[LPSY15]
Semi-honest, honest majority	[AFL <sup>+</sup> 16] [CDM00]	[AFL <sup>+</sup> 16]	[AFL <sup>+</sup> 16] [CCD88]	[LPSY15]

encrypt thousands of field elements at once for reasonable security parameters. This enables using a single encryption even for computing sequential multiplications.

As a result, MP-SPDZ implements an OT-based protocol in all computation domains as well protocols for fields both with semi-homomorphic and somewhat homomorphic encryption.

### 3.2 Honest Majority

The honest-majority setting allows computing securely entirely using secret sharing, without oblivious transfer or homomorphic encryption. MP-SPDZ uses two secret sharing schemes to this end, replicated secret sharing [11] and Shamir’s secret sharing [70]. Both are multiplicative, that is, it is possible to locally compute a sharing of the multiplication of two sharings without communication, albeit resulting in a different secret sharing scheme. A resharing protocol involving communication can then convert back to the original sharing scheme, thus facilitating further multiplication.

The resharing protocol is linear in that resharing of a sum is equivalent to the sum of resharings, which enables the parties to compute an inner product at the communication cost of a single multiplication. The framework reflects this by providing a particular interface for inner products as seen later.

*Semi-honest computation based on replicated secret sharing.* Araki et al. [5] have observed that resharing can be done by every party sending exactly one element of the computation domain by using pseudo-random zero sharing [24]. The latter refers to the generation of a fresh random sharing of zero without communication using pseudo-random number generation with a set of keys shared between every pair of parties. A similar technique facilitates reducing the communication for private inputs as noted by Eerikson et al. [33].

*Semi-honest computation based on Shamir’s secret sharing.* This goes back to Ben-Or et al. [10] for computation modulo a prime and Chaum et al. [21] for computation in extension fields of characteristic two (and thus binary circuits). MP-SPDZ implements resharing as described by Cramer et al. [25].

*Malicious computation.* Lindell and Nof [58] have adapted the SPDZ sacrifice to the setting of multiplicative secret sharing modulo a prime (including replicated and Shamir’s secret sharing), and Araki et al. [37] have done so for TinyOT. Both works are based on the observation that the optimistic triple production using homomorphic encryption or oblivious transfer can be replaced by producing a triple using multiplicative secret sharing instead. A

pair of such triples can then be checked as in the dishonest-majority setting, resulting in the ratio of triples needed to be produced per triple available for the online phase. This is two for large fields (SPDZ) and three for bit-based protocols (TinyOT).

For computation modulo a power of two, Eerikson et al. [33] have presented several variants, one more akin to TinyOT, another inspired by the SPDZ<sub>2<sup>k</sup></sub> sacrifice, a third based on a general compiler [30], and a fourth called post-sacrifice. The last was introduced by Lindell and Nof [58] for computation modulo a prime, and it works by executing any multiplication with only semi-honest security but storing the inputs and output for a later check with a random triple similar to a SPDZ sacrifice. This deviates from Beaver’s technique and demonstrates that MP-SPDZ is not restricted to it.

### 3.3 Higher-Level Protocols

All protocols described above are generally concerned with an implementation of an arithmetic black-box, that is private input, addition, multiplication, and public output. However, many computations such as comparison require further correlated randomness, most notably secret random bits in the larger domains. A simple way of obtaining such bits is to have sufficiently many parties (depending on the security model) input a random bit and then computing the XOR of all of them. However, this does not scale optimally with the number of parties because computing XOR in a larger domain reduces to a multiplication:  $a \oplus b = a + b - 2ab$ . Furthermore, for malicious security, it has to be checked whether the output actually is a bit.

For computation modulo a prime, Damgård et al. [28] have shown that a secret random bit can be produced at the cost of only one multiplication and one opening by using the fact that the square of a random number does not reveal which of two possible square roots it corresponds to (if not zero). Damgård et al. [27] later extended this to computation modulo a power of two.

*Arithmetic-binary conversion.* MP-SPDZ implements several ways of converting between arithmetic (modulo a larger number) and binary (modulo two) computation. The general way uses correlated randomness in the two domains, so-called doubly-authenticated bits (daBits) introduced by Rotaru and Wood [68], and extended daBits introduced by Escudero et al. [34]. In semi-honest computation with a low number of parties, there is also the possibility of a more direct conversion [4, 61].

### 3.4 Benchmarks

In order to compare protocols for arithmetic circuits, we have benchmarked the computation of one epoch of logistic regression for 1280 examples of 10 features and batch size 128. The resulting circuit features more than one million multiplications in more than 1000 rounds. Tables 3 and 7 show our results for one and two corrupted players, respectively, on co-located AWS c5.9xlarge instances. Note that we compute the sigmoid function accurately unlike for example ABY<sup>3</sup> [61].

For protocols computing binary circuits, we have implemented the computation of an inner product of vectors input by two different parties with 32-bit numbers. Tables 4 and 8 show our results for one and two corrupted players, respectively, on co-located AWS c5.9xlarge instances. Note that [CCD88] stands for computing a binary circuit by embedding bits in into fields of characteristic two to satisfy the conditions of Shamir’s secret sharing. Furthermore, the protocol after [LPSY15] names the underlying protocol for the BMR construction.

## 4 INTERNAL INTERFACES

In multi-party computation based on secret sharing it is essential that parallel communication is bundled up to some extent. This raises the question of how to translate that requirement into an interface that is as easy to use as possible. Consider the following simple operator overloading:

```
a = b * c
e = d * f
```

The two multiplications can be run in parallel. However, the communication implied by the first line has to be deferred at least until the execution of the second line, which means that that a cannot represent an actual share but a deferred value instead, which in turn requires a considerable machinery in the background to handle the the communication and the availability of deferred values. Such an approach has been taken by VIFF [38], its successor MPyC [69], and FRESCO [1]. In contrast, MP-SPDZ aims for more efficiency at this stage and defers usability to a higher level. Therefore, the internal interface requires the programmer indicate parallel computation. A strictly vector-based interface is one way of doing this, but MP-SPDZ offers a slightly more amenable variant. Assume that a and b are vectors of multiplicands of size n and c is supposed to hold the output.

```
protocol.init_mul();
for (int i = 0; i < n; i++)
    protocol.prepare_mul(a[i], b[i]);
protocol.exchange();
for (int i = 0; i < n; i++)
    c[i] = protocol.finalize_mul();
```

This interface allows a more dynamic use without the necessity of copying information into vectors first.

As every protocol offers the same interface, it is straightforward to implement the same computation for several protocols at once using C++ templating also because the input and output protocols use the same approach.

Coming back to our inner product example, Figure 2 shows its implementation using the internal interface of an arithmetic

```
for (int i = 0; i < n; i++)
    input.input_from_all();
input.exchange();
for (int i = 0; i < n; i++)
{
    a[i] = input.finalize(0);
    b[i] = input.finalize(1);
}
protocol.init_dotprod();
for (int i = 0; i < n; i++)
    protocol.prepare_dotprod(a[i], b[i]);
protocol.next_dotprod();
protocol.exchange();
for (int i = 0; i < n; i++)
    c = protocol.finalize_dotprod();
output.prepare_open(c);
output.exchange();
result = output.finalize_open();
```

Figure 2: Inner product computation with C++ in MP-SPDZ.

protocol. As the communication takes place in the exchange calls, it is parallelized as much as possible. Furthermore, the code uses the optimized dot product facility if available for the protocol. Note that the code is the same for all protocols. It is only the type of the input, protocol, and output objects that changes from protocol to protocol.

### 4.1 Templating

C++ templating is widely used in MP-SPDZ because it allows reusing code whenever suitable without performance penalty. If runtime polymorphism with virtual functions would be used instead, every additional protocol would incur extra cost at every point of the execution where protocols differ. Given that the protocols implemented in the framework differ in very small units (e.g., adding shares can involve as little as adding two 64-bit numbers but also adding several pairs of numbers modulo a large prime), we estimate that the cost of runtime polymorphism would be considerable due to the continuous branching.

As an example of the benefits of templating, the implementation of Beaver’s multiplication [7] requires about 100 lines of code while being used for over ten protocols spanning all computation domains and several security models.

The central class of any protocol represents a secret value in that protocol. Everything else like types for corresponding clear-text values and sub-protocols are derived from this class using typedef. Usually classes describing a particular protocol have one template variable for the secret value class, from which all derived information can be accessed.

### 4.2 Preprocessing

A considerable part of the backend is dedicated to the so-called preprocessing, which denotes the generation of correlated randomness. The most prominent example of such randomness is multiplication triples  $([a], [b], [ab])$  where  $a$  and  $b$  are uniformly random and all numbers are secret-shared. Such triples are central to the Beaver’s

**Table 3: Time and communication for logistic regression with various protocols tolerating one corrupted party.**

Security	Parties	Modulo	Protocol	Time (s)	Comm. (MB)
Semi-honest	3	$2^{64}$	[AFL <sup>+</sup> 16]	0.62	62
		128-bit prime	[AFL <sup>+</sup> 16]	4.41	245
		128-bit prime	[CDM00]	5.31	490
	2	$2^{64}$	[DEF <sup>+</sup> 19]-	27.52	8,173
		128-bit prime	[KOS16]-	267.24	41,424
		128-bit prime	[KPR18]- (semi HE)	134.44	2,367
128-bit prime		[KPR18]- (somewhat HE)	539.50	2,593	
Covert	2	128-bit prime	[KPR18]* (semi HE)	1173.91	15,578
			[KPR18]* (somewhat HE)	3708.61	27,130
Malicious	3	$2^{64}$	[EKO <sup>+</sup> 20] (post-sacrifice)	11.09	249
		128-bit prime	[LN17] (replicated)	11.09	566
		128-bit prime	[LN17] (Shamir)	18.25	1,888
	2	$2^{64}$	[DEF <sup>+</sup> 19]	973.17	109,189
		128-bit prime	[KOS16]	2336.57	225,670

**Table 4: Time and communication per party for inner product of length 100,000 with various binary circuit protocols tolerating one corrupted party. Numbers with \* are scaled up from inner products of length 1,000.**

Security	Parties	Type	Protocol	Time (s)	Comm. (GB)
Semi-honest	3	Secret sharing	[AFL <sup>+</sup> 16]	18	0.1
		Secret sharing	[CDM00][CCD88]	85	1.8
		Garbled circuit	[LPSY15][AFL <sup>+</sup> 16]	594	68.7
		Garbled circuit	[LPSY15][CDM00]	1,002	184.5
	2	Secret sharing	[FKOS15]-	59	5.4
		Garbled circuit	[ZRE15]	8	3.6
Garbled circuit		[LPSY15][KOS16]-	13,502*	4,178.4*	
Malicious	3	Secret sharing	[FLNW17]	27	1.9
		Secret sharing	[LN17][CCD88]	191	14.1
		Garbled circuit	[LPSY15][LN17]	1,777	163.6
		Garbled circuit	[LPSY15][LN17]	3,241	518.6
	2	Secret sharing	[FKOS15]	897	73.4
		Garbled circuit	[LPSY15][KOS16]	132,243*	24,486.9*

multiplication. Other examples include squares ( $a, a^2$ ), inverses ( $a, a^{-1}$ ), random bits, daBits, and edaBits [34].

MP-SPDZ provides implementations to generate preprocessing generically for semi-honest and malicious security as well as specific implementations for certain protocols. For example, it is more efficient to generate squares without using multiplication in SPDZ [28], or daBits can be generated from random bits in some protocols for computing modulo a power of two.

Since the marginal cost per preprocessed element is lower if they are generated in batches, the framework provides infrastructure to do so. Furthermore, it is possible to read such information from disk, which facilitates benchmarking only the phase of the computation that is dependent on secret data (sometimes called “online phase”).

Preprocessing is generally executed on demand to avoid unnecessary computation, which is different to SCALE-MAMBA [22] where even the lesser used squares are generated whenever a computation is run. This does not only slow down the computation, it can also cause the application to seemingly hang at the end because the preprocessing has not finished.

## 5 THE VIRTUAL MACHINE

Keller et al. [53] have presented a virtual machine designed for the specifics of multi-party computation. This virtual machine is at the core of SPDZ-2 [51], the predecessor of MP-SPDZ. The main characteristic of the virtual machine is that instructions involving communication allow an unrestricted number of arguments and thus minimizing the number of communication rounds. It is this

property that distinguishes it from many other virtual machines and processors. Consider for example the 64-bit x86 instruction set. While instructions vary in the number of clock cycles required for completion (as little as one for addition but more than 100 for sine or cosine [35]), the difference between instructions is easily quantified by the difference in the binary circuits required to compute them. In multi-party computation, there is not only a quantitative difference between an addition of shares and a multiplication but also a qualitative one because the former can be done locally while the latter involves communication. This qualitative difference has implications that vary from protocol to protocol, but the benefit of unrestricted parallelization of communication is straightforward due to network latencies.

In the context of garbled circuits, there is another benefit of the parallelization property. Bellare et al. [9] proposed implementing garbled circuits using AES-NI, the CPU-native implementation of AES. Their scheme makes use of so-called pipelining, which means that several AES operations can be run in parallel on the same CPU core if the instructions are executed directly one after the other. The virtual machine design allows one to make use of pipelining while relieving the user in this respect.

*High-level design.* SPDZ-2 implements the SPDZ online phase for computations in  $\mathbb{F}_p$  for a large prime  $p$  and  $\mathbb{F}_{2^n}$  with either  $n = 40$  or  $n = 128$ . MP-SPDZ adds a range of protocols, replacing  $\mathbb{F}_p$  with  $\mathbb{Z}_{2^k}$  for some protocols and adding computations of Boolean circuits (technically in  $\mathbb{F}_2$ ). The latter is rooted in the implementation used by Keller and Yanai [55], and it works with vectors of length 64 by default. This is necessary for efficiency with some protocols where a share only consists of a bit or a pair of bits, and it leads to a natural optimization by using the 64-bit machine words of contemporary processors, computing 64 bit operations at once. As a result, every instantiation of the virtual machine offers integer-like computation (in  $\mathbb{F}_p$  or  $\mathbb{Z}_{2^k}$ ), and computation for  $\mathbb{F}_{2^n}$  and Boolean circuits, all in the same security model.

*Basic data types.* The virtual machine allows handling data of public and secret values for every computation domain and 64-bit integers. The reason for having clear data types in every domain on top of another integer type is that the size of numbers in the computation domains vary and numbers in  $\mathbb{F}_p$  are stored using Montgomery representation, which does not lend itself for purposes such as loop counters or addressing memory. It is therefore cleaner to have public data types reflecting all computation domains. Furthermore, in the context of garbled circuits there is a conceptual difference between public numbers that are known prior to garbling such as loop counters and numbers that result from revealing a secret value. This difference matters because the garbling part of a computation depends on a revealed value has to happen after the evaluation resulting in said value whereas garbling only depending on a loop counter can be processed at any time.

*Registers.* The virtual machine provides an unlimited number of registers for every basic data type. While this is less sophisticated than a stack-based design, it allows for a simpler implementation. Register numbers are hard-coded into the bytecode, which enables

the virtual machine to allocate sufficient numbers for the computation. Registers are generally used to store inputs and outputs of instructions, and they are local to a thread.

*Memory.* For more complex data structures such as arrays, matrices, and higher-dimensional structures, the virtual machine provides another facility for every basic data type, called the memory. The memory arrays are global, and thus allows communicating information between threads. Unlike registers, the memory can be accessed using runtime values stored in integer registers. Memory has to be allocated at compile time.

*Instructions.* Most of the instructions supported by the virtual machine can be roughly categorized as follows:

**Copying** This includes initializing registers, copying between registers and memory, as well as conversion between registers of different public data types.

**Simple computation** Instructions resembling the common three-argument format is used for all computation that does not require communication, including linear operations on secret values.

**Secure protocols** All instructions of this category allow an infinite number of arguments to facilitate reducing communication rounds as described above. The protocols include essential ones such as private input, multiplication, and public output, but also more specialized ones such as inner product, matrix multiplication, 2D convolution, and the special truncation by Dalskov et al. [26].

**Preprocessed information** As mentioned in Section 4.2, preprocessing is executed batch-wise. This means that instructions of this type only incur communication if necessary. It is also possible to communicate the amount of preprocessed information required in a tape in order to optimize the preprocessing.

**Control flow** The virtual machine allows jumps as well as spawning and joining threads.

**Further input/output** Instructions in this category facilitate functionality outside secure protocols such as printing or client communication.

**Protocol information** It is possible to use the same program in various player configurations by accessing information such as the number of players and then for example loop over all available players when gathering inputs.

*Vectorization.* Most instructions are vectorized, that is, they imply the execution of the same operation for as many consecutive registers as requested. This considerably reduces the overhead of representing repetitive computation both during compilation as well as during execution. The virtual machine also facilitates structured loading of values into consecutive registers, for example, loading rows at any dimension in multi-dimensional arrays.

*Threads.* The virtual machine implements multi-threading as follows: A computation is run in the main thread as described by list of instructions called a “tape”. The latter can start further tapes in other threads and wait for their completion. However, the maximum number of threads has to be known at compile time. Despite this limitation, the available functionality is powerful enough for

many applications that benefit from multi-threading such as matrix multiplication or convolutions.

*Inner product example.* Figure 3 shows the bytecode representation of our inner product example with vectors of length three. Note that both `inputmixed`, `dotprods`, `asm_open` give the number of parameters as first argument. The first instruction lets private inputs from parties 0 and 1 to be stored in registers `s1`, `s3`, `s5` and `s2`, `s4`, `s6`, respectively, and the result of the inner product is then stored in `s0`. Eventually, the inner product is revealed and printed followed by a new line character.

## 6 THE COMPILER

Similar to SPDZ-2, the compiler runs high-level code written in Python, and outputs bytecode to be run by the virtual machine. Some aspects have been changed as pointed out throughout this chapter. The following core aspects remain the same, however.

*Type system.* MP-SPDZ follows the dynamic typing paradigm of Python. This makes programming more intuitive in that, for example, any operation involving a secret and a public value results in secret value. Consider that in a stronger type system, the assignment of a secret type to a public type would involve automatic revealing, which can be unintended, or the compiler would need to produce an error, which makes programming harder. Dynamic typing together with a clearly named way of revealing secret information therefore strikes a compromise between security and ease of use.

*Basic blocks.* This concept is taken from general compiler design and denotes a sequence of instructions without branching. The compiler performs the round-minimizing optimization only in the context of a basic block because it requires rearranging instructions.

### 6.1 Minimizing the Number of Rounds

This is the main optimization conducted by the compiler. It analyses a basic block to find instructions of the same kind that can be merged into a single instruction because they are independent. Recall that such instructions allow an infinite number of arguments. An easy example of a possible merge are two multiplications that can be run in parallel. MP-SPDZ differs from SPDZ-2 in that it merges different operations independently while SPDZ-2 reduces multiplication to openings using Beaver multiplications. Such a reduction clearly excludes protocols that do not use Beaver multiplications. As a result, MP-SPDZ uses an instruction for secure multiplication, which is merged independently of opening operations. Another difference to SPDZ-2 is that the latter splits the opening process in two instructions called `startopen` and `stopopen`, which enables executing local computation while waiting for information from the network. However, this comes at the cost of an increased complexity in the compiler as well as the fact that the number of parallel opens is limited by the communication buffer, which has to be considered on the compiler level. MP-SPDZ uses a single instruction for opening, which simplifies the handling.

The algorithm proceeds in three steps:

- (1) First it creates a dependency graph of all instructions in a basic block. This graph takes the form of a directed acyclic graph where the nodes stand for instructions. There are a few

reasons to create an edge between two nodes, most notably if the output of one instruction is an input to another. Other dependencies include various types of interaction with the environment, e.g., the order of input read from a party should not change.

- (2) The algorithm assigns instructions that can be merged to rounds. All instructions in the same round have the same type. This is done by a variant of the longest-path algorithm:
  - (a) All nodes without predecessors are assigned to round zero.
  - (b) Every non-mergeable node is assigned the maximum of the round number of its predecessors.
  - (c) Every mergeable node is assigned the minimal round that is larger than all of its predecessors and that is compatible, that is, it is not occupied by an instruction of another type. Since this algorithm only considers predecessors of nodes, it can be run together with the dependency graph creation.
- (3) All instructions in the same round are merged. This involves merging the arguments such that the semantics are preserved, and merging all edges in the dependency graph.
- (4) The instructions are output in topological order according to the changed dependency graph.

*Memory instruction dependency.* Since memory instructions allow runtime addresses, it is not straightforward how to treat them in the dependency graph. If every read instruction is made to depend on a write instruction and vice versa, a possibly large potential for minimizing rounds is lost. Consider an unrolled loop where a memory address is read followed by some computation on the read value, the result of which is stored at the same address, and the same is repeated for more addresses. All computations can clearly be executed in parallel, but a dependency of every memory instruction on the previous one prevents this. The compiler therefore only considers dependencies if they involve the same address, be it a compile-time address or the same register for run-time addresses. In combination with caching registers when accessing data structures such as multi-dimensional arrays, this strikes a balance between efficiency and correctness guarantees. Furthermore, the compiler offers a command to start a new basic block, which inherently preserves the order of instructions.

*Dead code elimination.* The dependency graph created above also allows eliminating instructions with unused results. An instruction is considered obsolete if it is not considered inherently essential (because of side-effects on the environment or the memory) and if all successors are obsolete. A simple backward pass suffices to determine and eliminate obsolete instructions.

### 6.2 Register Allocation

As in SPDZ-2, the compiler initially uses an unlimited supply of write-once registers, which are allocated to a minimal number of registers at the end. For a straightline program without branches, this is trivial by passing through backwards, allocating a register whenever it is read for the last time, and deallocating it at the single time it is written to. For a program with branches there is the difficulty that some registers have to be live throughout a loop to prevent overwrites after the last read. This is solved by allocating

```

inputmixed 18, 0, s1, 0, 0, s3, 0, 0, s5, 0, 0, s2, 1, 0, s4, 1, 0, s6, 1 # 0
dotprods 8, 8, s0, s1, s2, s3, s4, s5, s6 # 1
asm_open 2, c0, s0 # 2
print_reg_plain c0 # 3
print_char 10 # 4

```

Figure 3: Virtual machine code of our inner product example with vectors of length three.

registers that are written to before the loop starts before processing the parts of the loop.

### 6.3 Loops

While the nature of multi-party computation makes it non-trivial to implement loops that depend on secret data, loops depending on public data naturally reduce the representation of computation. Similar to SPDZ-2, MP-SPDZ supports loops depending both on compile-time and run-time public data. The former still allows compile-time analysis of the computation cost. In the high-level language loops can be executed using function decorators:

```

@for_range(n)
def _(i):
    a[i] = ...

```

While this is certainly unusual, it enables compilation by running the high-level Python code within the scope of the compiler. Furthermore, it allows a variety of loops without creating a domain-specific language. For example, `@for_range` creates a strict run-time loop executed consecutively, `@for_range_opt` implements the dynamic optimization described below, and `@for_range_multithread` and `@for_range_opt_multithread` execute a loop in a fixed number of threads.

*Dynamic loop optimization.* Büscher et al. [17] have described a trade-off between unrolling loops in order to merge communication rounds as in Section 6.1 and limiting memory usage during compilation caused by the increased space to represent the unrolled computation. They propose a dynamic approach of unrolling until a time budget is exhausted. MP-SPDZ adapts their approach by using a budget on the number of instructions produced, which serves as a proxy for the overall compilation cost.

### 6.4 Repetitive Code

Since the basic computation in multi-party computation is limited to additions and multiplications in a domain, even seemingly simple computations such as comparisons translate into a non-trivial composition of basic operations. The design principle of MP-SPDZ is to break down computation into basic operations in order to limit the complexity of the virtual machine. However, this means that repetitive code leads to a repetitive expansion of the same building block. To avoid that cost, the compiler offers to treat certain basic computations atomically and merging them as in Section 6.1. While speeding up compilation this has the downside that some parallel computations will be executed sequentially. For example, a comparison and an equality test on independent data can be executed in parallel (thus reducing the number of rounds), but treating the two separately will execute them sequentially. Because of the trade-off, this optimization is only used when requested.

Instances of building blocks treated in the way above range from integer operations such as comparison and truncation to mathematical functions such as the trigonometric functions.

As a concrete example, consider the case of computing the maximum of several values. To minimize the round complexity, this has to be computed as a binary tree, selecting the maximum at every node, and all comparisons on every level should be computed in parallel. It is somewhat onerous to program this as in a vectorized manner. MP-SPDZ allows using a simple recursive approach instead, merging all parallel comparisons. With the repetitive code optimization, this happens before creating the circuits used for comparison, which speeds up the compilation considerably.

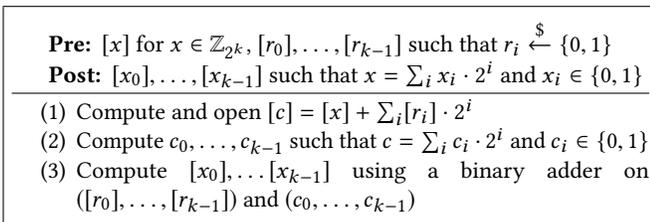
## 7 HIGH-LEVEL LIBRARY

In this section, we will describe the capabilities for secret computation implemented on top of the compiler. These implementations generally involve breaking down the desired functionality to the basic operations supported by the protocols such as input, output, and arithmetic in the computation domain. As these operations are modeled by the arithmetic black-box, it is usually straightforward to prove the security of the extended functionality.

### 7.1 Integer Operations

Beyond basic arithmetic in the respective domain, the library implements comparison, equality, left/right shift, modulo power of two, and power of two with secret exponent. A core component of most of these operations in larger computation domains is a “mask-and-open” approach. This involves adding or subtracting a secret random value of a certain form to or from a secret value and opening the result, which can then be processed as a public value, for example to extract single bits. See Figure 4 for the basic example of a bit decomposition. It uses these public bits together with the secret bits of the masking to compute a bit decomposition of a secret value, that is, an individual secret sharing for every bit instead of a secret sharing for the whole value. Unlike the latter, the former allows computing all integer operations above directly. Catrina and de Hoogh [19] have shown how to optimize this simple approach with computation modulo a prime, later adapted by Dalskov et al. [26] for computation modulo a power of two and by Escudero et al. [34] for switching to binary computation for parts of it.

Mask-and-open with computation modulo a prime is only statistically secure and requires the secret value to be within an assumed interval. For example, assume that  $x \in [0, 2^l]$  and that  $r \in [0, 2^{l+s}]$  is randomly chosen for a security parameter  $s$ . Then,  $x + r$  is statistically indistinguishable from a uniform number in  $[0, 2^{l+s}]$ . In computation modulo a power of two this is not an issue because



**Figure 4: Bit decomposition modulo  $2^k$  using random bits.**

overflow bits can simply be “erased” by multiplying with a power of two.

## 7.2 Fractional Numbers

As SPDZ-2, MP-SPDZ offers two ways of representing fractional numbers, fixed-point and floating-point. The former denotes representing a fractional number  $x$  by an integer  $y$  such that  $x = y \cdot 2^{-f}$  for a fixed precision  $f$ . The latter is similar to floating-point representations such as IEEE 754 in that  $x$  is represented by the four-tuple  $(v, p, z, s)$  such that

$$x = \begin{cases} (-1)^s \cdot v \cdot 2^p & z = 0 \\ 0 & z = 1. \end{cases}$$

The additional bits  $s$  and  $z$  simplify the computation given that secure computation does not directly allow one to access single bits of a larger value.

*Fixed-point numbers.* Due to the larger efficiency, this is the preferred approach for fractional numbers in MP-SPDZ. Addition and subtraction are straight-forward by the linearity of the representation, and multiplication corresponds to integer multiplication followed by truncation of  $f$  bits. The truncation can either be computed as a left shift or a more efficient probabilistic truncation. The latter involves randomized rounding based on the input. For example, 0.25 would be rounded to 0 with probability 0.75 and to 1 with probability 0.25 when rounding to integers. Catrina and Saxena first suggested this for computation modulo a prime [20], Dalskov et al. [26] adapted it for computation modulo a power of two, and Escudero et al. [34] to mixed computation. Catrina and Saxena also presented how to use Goldschmidt’s algorithm [40] to compute division in secure computation.

*Floating-point numbers.* Aliasgari et al. [2] have shown how to implement floating-point numbers in the context of secure computation modulo a prime. Their approach translates directly to  $\mathbb{Z}_{2^k}$  with the exception of multiplying by  $(2^m)^{-1}$  in order to compute a left shift for  $x$  when  $x = 0 \pmod{2^m}$ . This can be achieved through bit decomposition, however. The library implements addition, subtraction, division, and comparisons of floating-point numbers. Note that the approach by Aliasgari et al. is not fully compliant to IEEE754. However, it is possible to use the binary circuits for compliant computation of some operations (addition, multiplication, division, and square root) in the Bristol Fashion format provided by SCALE-MAMBA [22]. See Section 7.3.

*Mathematical functions.* Aly and Smart [3] have implemented a number of mathematical functions in secure computation, ranging

from trigonometric functions over square root to exponential and logarithm. MP-SPDZ integrates this by reusing code provided by the authors through SCALE-MAMBA [22]. All functions are implemented for fixed-point numbers while the implementations for floating-point numbers are restricted to sine, cosine, and tangent.

## 7.3 Further Functionality

The library extends beyond basic mathematics as shown in the following paragraphs.

*Machine learning.* The library provides functionality for logistic regression [54] and deep-learning inference [26]. The former allows choosing between an accurate implementation of the sigmoid function and an approximation similar to ABY<sup>3</sup> [61]. The latter supports the quantization scheme used in MobileNets [46] and several ImageNet solutions such as DenseNet, ResNet, and SqueezeNet. These are compiled automatically from TensorFlow with the help of CryptFlow [56], which opens the possibility for further networks to be used.

*Oblivious data structures.* MP-SPDZ retains the code used by Keller and Scholl [52]. This includes an oblivious array, queue, and stack implementation. Oblivious here means that all accesses are done secretly, that is, without revealing indices (where applicable) and whether an access is reading or writing. However, it is inherent to secure computation that an upper limit to the amount of data in the structure is revealed. Oblivious RAM [72] is the core technique to achieve efficient data structures for larger sizes. Based on the data structures above, the library contains example implementations of Dijkstra’s algorithm for shortest path in graphs, and the Gale-Shapley algorithm for stable matching.

*Binary circuits.* The library allows processing binary circuits in the so-called Bristol Fashion format. This text-based format has been established by the authors of SCALE-MAMBA [22], inspired by similar formats used for other secure computation frameworks. SCALE-MAMBA comes a selection of example circuits such as the Keccak sponge function [13]. Based on the latter, the library provides the computation of SHA3 for short inputs. Unlike SCALE-MAMBA, MP-SPDZ processes binary circuits in the compiler in the compiler in order to optimize them for secret sharing computation and Yao’s garbled circuits. SCALE-MAMBA only uses Bristol Fashion circuits in the context of BMR-style garbled circuits where such optimizations matter less because the key used in AES is not fixed.

## ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers of CCS 2020 for their helpful suggestions and comments.

## REFERENCES

- [1] Alexandra Institute. 2020. FRESCO - A Framework for Efficient Secure Computation. (2020). <https://github.com/aicis/fresco>.
- [2] Mehrdad Aliasgari, Marina Blanton, Yihua Zhang, and Aaron Steele. 2013. Secure Computation on Floating Point Numbers. In *NDSS 2013*. The Internet Society.
- [3] Abdelrahman Aly and Nigel P. Smart. 2019. Benchmarking Privacy Preserving Scientific Operations. In *ACNS 19 (LNCS)*, Robert H. Deng, Valérie Gauthier-Umaña, Martín Ochoa, and Moti Yung (Eds.), Vol. 11464. Springer, Heidelberg, 509–529. [https://doi.org/10.1007/978-3-030-21568-2\\_25](https://doi.org/10.1007/978-3-030-21568-2_25)
- [4] Toshinori Araki, Assi Barak, Jun Furukawa, Marcel Keller, Yehuda Lindell, Kazuma Ohara, and Hikaru Tsuchida. 2018. Generalizing the SPDZ Compiler For Other Protocols. In *ACM CCS 2018*, David Lie, Mohammad Mannan, Michael Backes,

- and XiaoFeng Wang (Eds.). ACM Press, 880–895. <https://doi.org/10.1145/3243734.3243854>
- [5] Toshinori Araki, Jun Furukawa, Yehuda Lindell, Ariel Nof, and Kazuma Ohara. 2016. High-Throughput Semi-Honest Secure Three-Party Computation with an Honest Majority. In *ACM CCS 2016*, Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi (Eds.). ACM Press, 805–817. <https://doi.org/10.1145/2976749.2978331>
  - [6] Carsten Baum, Daniele Cozzo, and Nigel P. Smart. 2019. Using TopGear in Overdrive: A More Efficient ZKPoK for SPDZ. In *SAC 2019 (LNCS)*, Kenneth G. Paterson and Douglas Stebila (Eds.), Vol. 11959. Springer, Heidelberg, 274–302. [https://doi.org/10.1007/978-3-030-38471-5\\_12](https://doi.org/10.1007/978-3-030-38471-5_12)
  - [7] Donald Beaver. 1992. Efficient Multiparty Protocols Using Circuit Randomization. In *CRYPTO'91 (LNCS)*, Joan Feigenbaum (Ed.), Vol. 576. Springer, Heidelberg, 420–432. [https://doi.org/10.1007/3-540-46766-1\\_34](https://doi.org/10.1007/3-540-46766-1_34)
  - [8] Donald Beaver, Silvio Micali, and Phillip Rogaway. 1990. The Round Complexity of Secure Protocols (Extended Abstract). In *22nd ACM STOC*. ACM Press, 503–513. <https://doi.org/10.1145/100216.100287>
  - [9] Mihir Bellare, Viet Tung Hoang, Sriram Keelveedhi, and Phillip Rogaway. 2013. Efficient Garbling from a Fixed-Key Blockcipher. In *2013 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, 478–492. <https://doi.org/10.1109/SP.2013.39>
  - [10] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. 1988. Completeness Theorems for Non-Cryptographic Fault-Tolerant Distributed Computation (Extended Abstract). In *20th ACM STOC*. ACM Press, 1–10. <https://doi.org/10.1145/62212.62213>
  - [11] Josh Cohen Benaloh and Jerry Leichter. 1990. Generalized Secret Sharing and Monotone Functions. In *CRYPTO'88 (LNCS)*, Shafi Goldwasser (Ed.), Vol. 403. Springer, Heidelberg, 27–35. [https://doi.org/10.1007/0-387-34799-2\\_3](https://doi.org/10.1007/0-387-34799-2_3)
  - [12] Rikke Bendlin, Ivan Damgård, Claudio Orlandi, and Sarah Zakarias. 2011. Semi-homomorphic Encryption and Multiparty Computation. In *EUROCRYPT 2011 (LNCS)*, Kenneth G. Paterson (Ed.), Vol. 6632. Springer, Heidelberg, 169–188. [https://doi.org/10.1007/978-3-642-20465-4\\_11](https://doi.org/10.1007/978-3-642-20465-4_11)
  - [13] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. 2009. Keccak sponge function family main document. *Submission to NIST (Round 2)* 3, 30 (2009), 320–337.
  - [14] Dan Bogdanov, Sven Laur, and Jan Willemson. 2008. Sharemind: A Framework for Fast Privacy-Preserving Computations. In *ESORICS 2008 (LNCS)*, Sushil Jajodia and Javier López (Eds.), Vol. 5283. Springer, Heidelberg, 192–206. [https://doi.org/10.1007/978-3-540-88313-5\\_13](https://doi.org/10.1007/978-3-540-88313-5_13)
  - [15] Peter Bogetoft, Ivan Damgård, Thomas Jakobsen, Kurt Nielsen, Jakob Pagter, and Tomas Toft. 2006. A Practical Implementation of Secure Auctions Based on Multiparty Integer Computation. In *EUROCRYPT 2006 (LNCS)*, Giovanni Di Crescenzo and Avi Rubin (Eds.), Vol. 4107. Springer, Heidelberg, 142–147.
  - [16] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. 2012. (Leveled) fully homomorphic encryption without bootstrapping. In *ITCS 2012*, Shafi Goldwasser (Ed.). ACM, 309–325. <https://doi.org/10.1145/2090236.2090262>
  - [17] Niklas Büscher, Daniel Demmler, Stefan Katzenbeisser, David Kretzmer, and Thomas Schneider. 2018. HyCC: Compilation of Hybrid Protocols for Practical Secure Computation. In *ACM CCS 2018*, David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang (Eds.). ACM Press, 847–861. <https://doi.org/10.1145/3243734.3243786>
  - [18] Niklas Büscher, Andreas Holzer, Alina Weber, and Stefan Katzenbeisser. 2016. Compiling Low Depth Circuits for Practical Secure Computation. In *ESORICS 2016, Part II (LNCS)*, Ioannis G. Askoxylakis, Sotiris Ioannidis, Sokratis K. Katsikas, and Catherine A. Meadows (Eds.), Vol. 9879. Springer, Heidelberg, 80–98. [https://doi.org/10.1007/978-3-319-45741-3\\_5](https://doi.org/10.1007/978-3-319-45741-3_5)
  - [19] Octavian Katrina and Sebastiaan de Hoogh. 2010. Improved Primitives for Secure Multiparty Integer Computation. In *SCN 10 (LNCS)*, Juan A. Garay and Roberto De Prisco (Eds.), Vol. 6280. Springer, Heidelberg, 182–199. [https://doi.org/10.1007/978-3-642-15317-4\\_13](https://doi.org/10.1007/978-3-642-15317-4_13)
  - [20] Octavian Katrina and Amitabh Saxena. 2010. Secure Computation with Fixed-Point Numbers. In *FC 2010 (LNCS)*, Radu Sion (Ed.), Vol. 6052. Springer, Heidelberg, 35–50.
  - [21] David Chaum, Claude Crépeau, and Ivan Damgård. 1988. Multiparty Unconditionally Secure Protocols (Extended Abstract). In *20th ACM STOC*. ACM Press, 11–19. <https://doi.org/10.1145/62212.62214>
  - [22] KU Leuven COSIC. 2019. SCALE-MAMBA. (2019). <https://github.com/KULeuven-COSIC/SCALE-MAMBA>.
  - [23] Ronald Cramer, Ivan Damgård, Daniel Escudero, Peter Scholl, and Chaoping Xing. 2018. SPD  $\mathbb{Z}_k$ : Efficient MPC mod  $2^k$  for Dishonest Majority. In *CRYPTO 2018, Part II (LNCS)*, Hovav Shacham and Alexandra Boldyreva (Eds.), Vol. 10992. Springer, Heidelberg, 769–798. [https://doi.org/10.1007/978-3-319-96881-0\\_26](https://doi.org/10.1007/978-3-319-96881-0_26)
  - [24] Ronald Cramer, Ivan Damgård, and Yuval Ishai. 2005. Share Conversion, Pseudorandom Secret-Sharing and Applications to Secure Computation. In *TCC 2005 (LNCS)*, Joe Kilian (Ed.), Vol. 3378. Springer, Heidelberg, 342–362. [https://doi.org/10.1007/978-3-540-30576-7\\_19](https://doi.org/10.1007/978-3-540-30576-7_19)
  - [25] Ronald Cramer, Ivan Damgård, and Ueli M. Maurer. 2000. General Secure Multiparty Computation from any Linear Secret-Sharing Scheme. In *EUROCRYPT 2000 (LNCS)*, Bart Preneel (Ed.), Vol. 1807. Springer, Heidelberg, 316–334. [https://doi.org/10.1007/3-540-45539-6\\_22](https://doi.org/10.1007/3-540-45539-6_22)
  - [26] Anders Dalskov, Daniel Escudero, and Marcel Keller. 2019. Secure Evaluation of Quantized Neural Networks. Cryptology ePrint Archive, Report 2019/131. (2019). <https://eprint.iacr.org/2019/131>.
  - [27] Ivan Damgård, Daniel Escudero, Tore Kasper Frederiksen, Marcel Keller, Peter Scholl, and Nikolaj Volgushev. 2019. New Primitives for Actively-Secure MPC over Rings with Applications to Private Machine Learning. In *2019 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, 1102–1120. <https://doi.org/10.1109/SP.2019.00078>
  - [28] Ivan Damgård, Marcel Keller, Enrique Larraia, Valerio Pastro, Peter Scholl, and Nigel P. Smart. 2013. Practical Covertly Secure MPC for Dishonest Majority - Or: Breaking the SPDZ Limits. In *ESORICS 2013 (LNCS)*, Jason Crampton, Sushil Jajodia, and Keith Mayes (Eds.), Vol. 8134. Springer, Heidelberg, 1–18. [https://doi.org/10.1007/978-3-642-40203-6\\_1](https://doi.org/10.1007/978-3-642-40203-6_1)
  - [29] Ivan Damgård, Jesper Buus Nielsen, Michael Nielsen, and Samuel Ranellucci. 2017. The TinyTable Protocol for 2-Party Secure Computation, or: Gate-Scrambling Revisited. In *CRYPTO 2017, Part I (LNCS)*, Jonathan Katz and Hovav Shacham (Eds.), Vol. 10401. Springer, Heidelberg, 167–187. [https://doi.org/10.1007/978-3-319-63688-7\\_6](https://doi.org/10.1007/978-3-319-63688-7_6)
  - [30] Ivan Damgård, Claudio Orlandi, and Mark Simkin. 2018. Yet Another Compiler for Active Security or: Efficient MPC Over Arbitrary Rings. In *CRYPTO 2018, Part II (LNCS)*, Hovav Shacham and Alexandra Boldyreva (Eds.), Vol. 10992. Springer, Heidelberg, 799–829. [https://doi.org/10.1007/978-3-319-96881-0\\_27](https://doi.org/10.1007/978-3-319-96881-0_27)
  - [31] Ivan Damgård, Valerio Pastro, Nigel P. Smart, and Sarah Zakarias. 2012. Multiparty Computation from Somewhat Homomorphic Encryption. In *CRYPTO 2012 (LNCS)*, Reihaneh Safavi-Naini and Ran Canetti (Eds.), Vol. 7417. Springer, Heidelberg, 643–662. [https://doi.org/10.1007/978-3-642-32009-5\\_38](https://doi.org/10.1007/978-3-642-32009-5_38)
  - [32] Daniel Demmler, Thomas Schneider, and Michael Zohner. 2015. ABE - A Framework for Efficient Mixed-Protocol Secure Two-Party Computation. In *NDSS 2015*. The Internet Society.
  - [33] Hendrik Eerikson, Marcel Keller, Claudio Orlandi, Pille Pullonen, Joonas Puura, and Mark Simkin. 2020. Use Your Brain! Arithmetic 3PC for Any Modulus with Active Security. In *ITC 2020*, Yael Tauman Kalai, Adam D. Smith, and Daniel Wichs (Eds.), Schloss Dagstuhl, 5:1–5:24. <https://doi.org/10.4230/LIPIcs.ITC.2020.5>
  - [34] Daniel Escudero, Satrajit Ghosh, Marcel Keller, Rahul Rachuri, and Peter Scholl. 2020. Improved Primitives for MPC over Mixed Arithmetic-Binary Circuits. In *CRYPTO 2020, Part II (LNCS)*, Daniele Micciancio and Thomas Ristenpart (Eds.), Vol. 12171. Springer, Heidelberg, 823–852. [https://doi.org/10.1007/978-3-030-56880-1\\_29](https://doi.org/10.1007/978-3-030-56880-1_29)
  - [35] Agner Fog et al. 2011. Instruction tables: Lists of instruction latencies, throughputs and micro-operation breakdowns for Intel, AMD and VIA CPUs. *Copenhagen University College of Engineering* 93 (2011), 110.
  - [36] Tore Kasper Frederiksen, Marcel Keller, Emanuela Orsini, and Peter Scholl. 2015. A Unified Approach to MPC with Preprocessing Using OT. In *ASIACRYPT 2015, Part I (LNCS)*, Tetsu Iwata and Jung Hee Cheon (Eds.), Vol. 9452. Springer, Heidelberg, 711–735. [https://doi.org/10.1007/978-3-662-48797-6\\_29](https://doi.org/10.1007/978-3-662-48797-6_29)
  - [37] Jun Furukawa, Yehuda Lindell, Ariel Nof, and Or Weinstein. 2017. High-Throughput Secure Three-Party Computation for Malicious Adversaries and an Honest Majority. In *EUROCRYPT 2017, Part II (LNCS)*, Jean-Sébastien Coron and Jesper Buus Nielsen (Eds.), Vol. 10211. Springer, Heidelberg, 225–255. [https://doi.org/10.1007/978-3-319-56614-6\\_8](https://doi.org/10.1007/978-3-319-56614-6_8)
  - [38] Martin Geisler. 2007. VIFF: Virtual Ideal Functionality Framework. <http://viff.dk/> (2007).
  - [39] Craig Gentry, Shai Halevi, and Nigel P. Smart. 2012. Homomorphic Evaluation of the AES Circuit. In *CRYPTO 2012 (LNCS)*, Reihaneh Safavi-Naini and Ran Canetti (Eds.), Vol. 7417. Springer, Heidelberg, 850–867. [https://doi.org/10.1007/978-3-642-32009-5\\_49](https://doi.org/10.1007/978-3-642-32009-5_49)
  - [40] Robert E. Goldschmidt. 1964. *Applications of division by convergence*. Master's thesis. MIT.
  - [41] Marcella Hastings, Brett Hemenway, Daniel Noble, and Steve Zdancewicz. 2019. SoK: General Purpose Compilers for Secure Multi-Party Computation. In *2019 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, 1220–1237. <https://doi.org/10.1109/SP.2019.00028>
  - [42] Marcella Hastings, Brett Hemenway, Daniel Noble, and Steve Zdancewicz. 2020. Sample code and build environments for MPC frameworks. (2020). <https://github.com/MPC-SoK/frameworks>.
  - [43] Carmit Hazay, Peter Scholl, and Eduardo Soria-Vazquez. 2017. Low Cost Constant Round MPC Combining BMR and Oblivious Transfer. In *ASIACRYPT 2017, Part I (LNCS)*, Tsuyoshi Takagi and Thomas Peyrin (Eds.), Vol. 10624. Springer, Heidelberg, 598–628. [https://doi.org/10.1007/978-3-319-70694-8\\_21](https://doi.org/10.1007/978-3-319-70694-8_21)
  - [44] Brett Hemenway, Steve Lu, Rafail Ostrovsky, and William Welsch IV. 2016. High-Precision Secure Computation of Satellite Collision Probabilities. In *SCN 16 (LNCS)*, Vassilis Zikas and Roberto De Prisco (Eds.), Vol. 9841. Springer, Heidelberg, 169–187. [https://doi.org/10.1007/978-3-319-44618-9\\_9](https://doi.org/10.1007/978-3-319-44618-9_9)

- [45] Yuval Ishai, Joe Kilian, Kobbi Nissim, and Erez Petrank. 2003. Extending Oblivious Transfers Efficiently. In *CRYPTO 2003 (LNCS)*, Dan Boneh (Ed.), Vol. 2729. Springer, Heidelberg, 145–161. [https://doi.org/10.1007/978-3-540-45146-4\\_9](https://doi.org/10.1007/978-3-540-45146-4_9)
- [46] Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew G. Howard, Hartwig Adam, and Dmitry Kalenichenko. 2017. Quantization and Training of Neural Networks for Efficient Integer-Arithmetic-Only Inference. *CoRR abs/1712.05877* (2017). <http://arxiv.org/abs/1712.05877>
- [47] Marcel Keller. 2017. The Oblivious Machine - Or: How to Put the C into MPC. In *LATINCRYPT 2017 (LNCS)*, Tanja Lange and Orr Dunkelman (Eds.), Vol. 11368. Springer, Heidelberg, 271–288. [https://doi.org/10.1007/978-3-030-25283-0\\_15](https://doi.org/10.1007/978-3-030-25283-0_15)
- [48] Marcel Keller, Emmanuela Orsini, and Peter Scholl. 2015. Actively Secure OT Extension with Optimal Overhead. In *CRYPTO 2015, Part I (LNCS)*, Rosario Gennaro and Matthew J. B. Robshaw (Eds.), Vol. 9215. Springer, Heidelberg, 724–741. [https://doi.org/10.1007/978-3-662-47989-6\\_35](https://doi.org/10.1007/978-3-662-47989-6_35)
- [49] Marcel Keller, Emmanuela Orsini, and Peter Scholl. 2016. MASCOT: Faster Malicious Arithmetic Secure Computation with Oblivious Transfer. In *ACM CCS 2016*, Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi (Eds.). ACM Press, 830–842. <https://doi.org/10.1145/2976749.2978357>
- [50] Marcel Keller, Valerio Pastro, and Dragos Rotaru. 2018. Overdrive: Making SPDZ Great Again. In *EUROCRYPT 2018, Part III (LNCS)*, Jesper Buus Nielsen and Vincent Rijmen (Eds.), Vol. 10822. Springer, Heidelberg, 158–189. [https://doi.org/10.1007/978-3-319-78372-7\\_6](https://doi.org/10.1007/978-3-319-78372-7_6)
- [51] Marcel Keller, Dragos Rotaru, Peter Scholl, and Nigel P. Smart. 2018. Multiparty computation with SPDZ, MASCOT, and Overdrive offline phases. (2018). <https://github.com/bristolcrypto/SPDZ-2>.
- [52] Marcel Keller and Peter Scholl. 2014. Efficient, Oblivious Data Structures for MPC. In *ASIACRYPT 2014, Part II (LNCS)*, Palash Sarkar and Tetsu Iwata (Eds.), Vol. 8874. Springer, Heidelberg, 506–525. [https://doi.org/10.1007/978-3-662-45608-8\\_27](https://doi.org/10.1007/978-3-662-45608-8_27)
- [53] Marcel Keller, Peter Scholl, and Nigel P. Smart. 2013. An architecture for practical actively secure MPC with dishonest majority. In *ACM CCS 2013*, Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung (Eds.). ACM Press, 549–560. <https://doi.org/10.1145/2508859.2516744>
- [54] Marcel Keller and Ke Sun. 2019. A Note on Our Submission to Track 4 of iDASH 2019. *Cryptology ePrint Archive, Report 2019/1246*. (2019). <https://eprint.iacr.org/2019/1246>.
- [55] Marcel Keller and Avishay Yanai. 2018. Efficient Maliciously Secure Multiparty Computation for RAM. In *EUROCRYPT 2018, Part III (LNCS)*, Jesper Buus Nielsen and Vincent Rijmen (Eds.), Vol. 10822. Springer, Heidelberg, 91–124. [https://doi.org/10.1007/978-3-319-78372-7\\_4](https://doi.org/10.1007/978-3-319-78372-7_4)
- [56] Nishant Kumar, Mayank Rathee, Nishanth Chandran, Divya Gupta, Aseem Rastogi, and Rahul Sharma. 2020. CryptFlow: Secure TensorFlow Inference. In *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, Los Alamitos, CA, USA, 336–353. <https://doi.org/10.1109/SP40000.2020.00092>
- [57] Andrei Lapets, Nikolaj Volgushev, Azer Bestavros, Frederick Jansen, and Mayank Varia. 2016. Secure MPC for analytics as a web application. In *2016 IEEE Cybersecurity Development (SecDev)*. IEEE, 73–74.
- [58] Yehuda Lindell and Ariel Nof. 2017. A Framework for Constructing Fast MPC over Arithmetic Circuits with Malicious Adversaries and an Honest-Majority. In *ACM CCS 2017*, Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu (Eds.). ACM Press, 259–276. <https://doi.org/10.1145/3133956.3133999>
- [59] Yehuda Lindell, Benny Pinkas, Nigel P. Smart, and Avishay Yanai. 2015. Efficient Constant Round Multi-party Computation Combining BMR and SPDZ. In *CRYPTO 2015, Part II (LNCS)*, Rosario Gennaro and Matthew J. B. Robshaw (Eds.), Vol. 9216. Springer, Heidelberg, 319–338. [https://doi.org/10.1007/978-3-662-48000-7\\_16](https://doi.org/10.1007/978-3-662-48000-7_16)
- [60] Chang Liu, Xiao Shaun Wang, Kartik Nayak, Yan Huang, and Elaine Shi. 2015. OblivM: A Programming Framework for Secure Computation. In *2015 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, 359–376. <https://doi.org/10.1109/SP.2015.29>
- [61] Payman Mohassel and Peter Rindal. 2018. ABY<sup>3</sup>: A Mixed Protocol Framework for Machine Learning. In *ACM CCS 2018*, David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang (Eds.). ACM Press, 35–52. <https://doi.org/10.1145/3243734.3243760>
- [62] Payman Mohassel and Yupeng Zhang. 2017. SecureML: A System for Scalable Privacy-Preserving Machine Learning. In *2017 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, 19–38. <https://doi.org/10.1109/SP.2017.12>
- [63] Peter L. Montgomery. 1985. Modular multiplication without trial division. *Mathematics of computation* 44, 170 (1985), 519–521.
- [64] B. Mood, D. Gupta, H. Carter, K. Butler, and P. Traynor. 2016. Frigate: A Validated, Extensible, and Efficient Compiler and Interpreter for Secure Computation. In *2016 IEEE European Symposium on Security and Privacy (EuroSP)*. 112–127.
- [65] Jesper Buus Nielsen, Peter Sebastian Nordholt, Claudio Orlandi, and Sai Sheshank Burra. 2012. A New Approach to Practical Active-Secure Two-Party Computation. In *CRYPTO 2012 (LNCS)*, Reihaneh Safavi-Naini and Ran Canetti (Eds.), Vol. 7417. Springer, Heidelberg, 681–700. [https://doi.org/10.1007/978-3-642-32009-5\\_40](https://doi.org/10.1007/978-3-642-32009-5_40)
- [66] Benny Pinkas, Thomas Schneider, Nigel P. Smart, and Stephen C. Williams. 2009. Secure Two-Party Computation Is Practical. In *ASIACRYPT 2009 (LNCS)*, Mitsuru Matsui (Ed.), Vol. 5912. Springer, Heidelberg, 250–267. [https://doi.org/10.1007/978-3-642-10366-7\\_15](https://doi.org/10.1007/978-3-642-10366-7_15)
- [67] Aseem Rastogi, Matthew A. Hammer, and Michael Hicks. 2014. Wysteria: A Programming Language for Generic, Mixed-Mode Multiparty Computations. In *2014 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, 655–670. <https://doi.org/10.1109/SP.2014.48>
- [68] Dragos Rotaru and Tim Wood. 2019. MARBled Circuits: Mixing Arithmetic and Boolean Circuits with Active Security. In *INDOCRYPT 2019 (LNCS)*, Feng Hao, Sushmita Ruj, and Sourav Sen Gupta (Eds.), Vol. 11898. Springer, Heidelberg, 227–249. [https://doi.org/10.1007/978-3-030-35423-7\\_12](https://doi.org/10.1007/978-3-030-35423-7_12)
- [69] Barry Schoenmakers. 2020. MPyC: Secure Multiparty Computation in Python. <https://github.com/lshoe/mpyc>. (2020).
- [70] Adi Shamir. 1979. How to Share a Secret. *Communications of the Association for Computing Machinery* 22, 11 (Nov. 1979), 612–613.
- [71] Ebrahim M. Songhori, Siam U. Hussain, Ahmad-Reza Sadeghi, Thomas Schneider, and Farinaz Koushanfar. 2015. TinyGarble: Highly Compressed and Scalable Sequential Garbled Circuits. In *2015 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, 411–428. <https://doi.org/10.1109/SP.2015.32>
- [72] Emil Stefanov, Marten van Dijk, Elaine Shi, Christopher W. Fletcher, Ling Ren, Xiangyao Yu, and Srinivas Devadas. 2013. Path ORAM: an extremely simple oblivious RAM protocol. In *ACM CCS 2013*, Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung (Eds.). ACM Press, 299–310. <https://doi.org/10.1145/2508859.2516660>
- [73] Multiparty.org Development Team. 2020. JavaScript Implementation of Federated Functionalities. (2020). <https://github.com/multiparty/jiff>.
- [74] Xiao Wang, Alex J. Malozemoff, and Jonathan Katz. 2016. EMP-toolkit: Efficient Multiparty computation toolkit. <https://github.com/emp-toolkit>. (2016).
- [75] Xiao Wang, Samuel Ranellucci, and Jonathan Katz. 2017. Global-Scale Secure Multiparty Computation. In *ACM CCS 2017*, Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu (Eds.). ACM Press, 39–56. <https://doi.org/10.1145/3133956.3133979>
- [76] Samee Zahur and David Evans. 2015. Obliv-C: A Language for Extensible Data-Oblivious Computation. *Cryptology ePrint Archive, Report 2015/1153*. (2015). <https://eprint.iacr.org/2015/1153>.
- [77] Samee Zahur, Mike Rosulek, and David Evans. 2015. Two Halves Make a Whole - Reducing Data Transfer in Garbled Circuits Using Half Gates. In *EUROCRYPT 2015, Part II (LNCS)*, Elisabeth Oswald and Marc Fischlin (Eds.), Vol. 9057. Springer, Heidelberg, 220–250. [https://doi.org/10.1007/978-3-662-46803-6\\_8](https://doi.org/10.1007/978-3-662-46803-6_8)
- [78] Yihua Zhang, Aaron Steele, and Marina Blanton. 2013. PICCO: a general-purpose compiler for private distributed computation. In *ACM CCS 2013*, Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung (Eds.). ACM Press, 813–826. <https://doi.org/10.1145/2508859.2516752>

## A FURTHER BENCHMARKS

**Table 5: Time in seconds to compute an inner product of two 100,000-element vectors of 64-bit integers in a WAN setting. Protocols skipping input sharing are denoted by “\*”.**

	SH 2-party (OT)	SH replicated 3-party	Shamir	Malicious Shamir	SPDZ (HighGear)	Yao’s GC (32-bit)
ABY	27	⊥	⊥	⊥	⊥	⊥
ABY <sup>3</sup>	⊥	1.5	⊥	⊥	⊥	⊥
EMP-toolkit	⊥	⊥	⊥	⊥	⊥	314
MPyC	⊥	⊥	17.8	⊥	⊥	⊥
Obliv-C	⊥	⊥	⊥	⊥	⊥	343
OblivM	⊥	⊥	⊥	⊥	⊥	1,712
PICCO	⊥	⊥	0.2*	⊥	⊥	⊥
SCALE-MAMBA	⊥	⊥	⊥	115	570	⊥
MP-SPDZ	29	0.9	0.2*/1.8	14	170	319

**Table 6: Total communication in MB to compute an inner product of two 100,000-element vectors of 64-bit integers. Protocols skipping input sharing are denoted by “\*”.**

	SH 2-party (OT)	SH replicated 3-party	Shamir	Malicious Shamir	SPDZ (HighGear)	Yao’s GC (32-bit)
ABY	312	⊥	⊥	⊥	⊥	⊥
ABY <sup>3</sup>	⊥	4.8	⊥	⊥	⊥	⊥
EMP-toolkit	⊥	⊥	⊥	⊥	⊥	3,395
MPyC	⊥	⊥	36	⊥	⊥	⊥
Obliv-C	⊥	⊥	⊥	⊥	⊥	3,523
OblivM	⊥	⊥	⊥	⊥	⊥	6,312
PICCO	⊥	⊥	0.006*	⊥	⊥	⊥
SCALE-MAMBA	⊥	⊥	⊥	157	4,547	⊥
MP-SPDZ	312	1.6	6.4	83	592	3,616

**Table 7: Time and communication for logistic regression with various protocols tolerating two corrupted parties.**

Security	Parties	Modulo	Protocol	Time (s)	Comm. (MB)
Semi-honest	5	128-bit prime	[CDM00]	17.07	1,452
		$2^{64}$	[DEF <sup>+</sup> 19]-	63.19	27,725
	3	128-bit prime	[KOS16]-	562.73	147,735
		128-bit prime	[KPR18]- (semi HE)	366.22	7,895
Covert	3	128-bit prime	[KPR18]- (somewhat HE)	550.08	5,058
			[KPR18]* (semi HE)	1769.68	30,621
Malicious	5	128-bit prime	[KPR18]* (somewhat HE)	3842.80	40,689
			[LN17] (Shamir)	44.02	4,188
	3	$2^{64}$	[DEF <sup>+</sup> 19]	1063.82	218,089
128-bit prime		[KOS16]	2559.30	451,019	

**Table 8: Time and communication per party for inner product of length 100,000 with various binary circuit protocols tolerating two corrupted parties. Numbers with “\*” are scaled up from inner products of length 1,000.**

Security	Parties	Type	Protocol	Time (s)	Comm. (GB)
Semi-honest	5	Secret sharing	[CDM00][CCD88]	107	3.2
		Garbled circuit	[LPSY15][CDM00]	3,416	578.7
	3	Secret sharing	[FKOS15]-	65	8.4
		Garbled circuit	[LPSY15][KOS16]-	24,219*	11,981.6*
Malicious	5	Secret sharing	[LN17][CCD88]	232	25.4
		Garbled circuit	[LPSY15][LN17]	12,087*	1,624.7*
	3	Secret sharing	[FKOS15]	1,029	140.9
		Secret sharing	[DEF <sup>+</sup> 19]*	40,891*	5,209.5*
		Garbled circuit	[LPSY15][KOS16]	249,508*	70,586.3*