MOTION – A Framework for Mixed-Protocol Multi-Party Computation

Lennart Braun
TU Darmstadt
lennart.braun@stud.tu-darmstadt.de

Daniel Demmler
Universität Hamburg
demmler@informatik.uni-hamburg.de

Thomas Schneider
TU Darmstadt
schneider@encrypto.cs.tu-darmstadt.de

Oleksandr Tkachenko
TU Darmstadt
tkachenko@encrypto.cs.tu-darmstadt.de

ABSTRACT

We present MOTION, an efficient and generic framework for mixed-protocol secure multi-party computation (MPC). Our framework is built from the ground up and incorporates several important engineering decisions such as full communication serialization which enables MPC over arbitrary messaging interfaces and removes the need of owning network sockets. It is available under the liberal MIT license and independent of external MPC libraries, which often have stricter licenses. MOTION is extensive and thoroughly tested: it currently consists of more than 36 000 lines of code, 20% of which are unit and component tests. It is built in a user-friendly, modular, and extensible way, intended to be used as tool in MPC research and to increase adoption of MPC protocols in practice. MOTION incorporates several novel performance optimizations that improve the communication complexity and latency, e.g., 2× better online round complexity of precomputed correlated Oblivious Transfer (OT).

We instantiate our framework with protocols for N parties and security against up to $N-1$ passive corruptions: the MPC protocols of Goldreich-Micali-Wigderson (GMW) in its arithmetic and Boolean version and oblivious transfer (OT)-based BMR (Ben-Efraim et al., CCS'16), as well as novel and highly efficient conversions between them, including a non-interactive conversion from BMR to arithmetic GMW. Moreover, we design a novel garbling technique that saves 20% of communication in the BMR protocol.

MOTION is highly efficient, which we demonstrate in our experiments by measuring its run-times in various network settings with different numbers of parties. For secure evaluation of AES-128 with N=3 parties in the high-latency network setting from the OT-based BMR paper, we achieve a 16× better throughput of 16 AES/s using BMR. This shows that the BMR protocol is much more competitive than previously assumed. For N=3 parties and full-threshold protocols in the LAN setting, MOTION is 10×–18× faster than the previous best passively secure implementation from the MP-SPDZ framework, and 190×–586× faster than the actively secure SCALE-MAMBA framework. Finally, we show that our framework is highly efficient for privacy-preserving neural network inference.

1 INTRODUCTION

Secure Multi-Party Computation (MPC) allows multiple parties to jointly compute a function on their private inputs without revealing anything but the function's output. This concept was first introduced in the 1980s by Yao [72] and Goldreich-Micali-Wigderson [33] and was initially considered of merely theoretical interest. The seminal work of Fairplay [55] was the first to implement MPC protocols and showed that MPC can indeed be practical. A long line of research has followed since then and has shown MPC to be a viable solution for preserving privacy in applications such as auctions [15], stable matching [30], set intersection [41], biometric matching [13, 56], and machine learning [57, 75].

This was facilitated by implementations of generic MPC frameworks that can be used for multiple applications. However, many MPC frameworks have a somewhat limited scope: they allow computations only for a fixed number of parties, e.g., two [19, 29, 39] or three [14, 20, 58, 60, 62] of which at most one can be corrupted, they only implement a single MPC protocol [55, 68], or are custom-tailored towards a few use-cases [54]. Furthermore, most of the publicly available code of MPC frameworks as surveyed in [36] is in a rather prototypical state and its main purpose is to generate performance measurements. This is a big problem, as these implementations cannot be used in practice and building on top of them in future research is often a tedious procedure that involves a significant amount of time that needs to be spent and expertise in understanding poorly documented code, fixing a multitude of existing problems and limitations.

In this work we present MOTION, an MPC framework that overcomes these limitations and aims to be a piece of software of high quality and usability. MOTION is object-oriented, developer-friendly, and well-documented, 20% of its code base are unit and component tests, and it is designed in a modular and extensible way, serving as a powerful tool for future MPC research and implementations. MOTION is a *generic solution* for implementing mixed-protocol MPC with two or more parties and guarantees security against all but one passively corrupted parties (full-threshold). It supports multiple MPC protocols and can securely and efficiently convert between them in order to benefit from each protocol's strengths. Previous mixed-protocol MPC frameworks are either limited to two parties [29], require an honest majority [14, 20, 58, 60, 62], or are full-threshold but with stronger active security only and hence are less efficient [24]. The architecture of MOTION allows to implement further MPC protocols in different security models.

Our motivation is to enable MPC protocols in practical application scenarios with an *arbitrary* number of parties and full threshold security. By this, we increase the number of parties in order to strengthen the *security* guarantees of the protocols. The limitation to exactly two or three parties of previous works might be problematic for practical settings, where a larger number of participants

want to jointly compute on private data. Also, in outsourcing scenarios [47], where a very large number of clients securely outsource computation to a smaller number of non-colluding computing parties, it might be desirable to increase the number of computing parties to achieve better security guarantees. In contrast, the goal of many previous works, e.g., [14, 20, 58, 60, 62], was to improve *performance* by increasing the number of parties, while simultaneously effectively reducing the security guarantees because only a single party can be corrupted.

Our protocols include novel performance optimizations in order to enable privacy-preserving computations for large real-world applications. With AES and private inference using a convolutional neural network, we present examples of such applications. Biometric identification using the Euclidean distance, which we also demonstrate, is an example for a multi-party application where a client wants to privately check if a data point is included in a large data set that is provided by multiple data owners. One could imagine a security check where a fingerprint is tested against databases of known fingerprints supplied by different security agencies. This scenario also finds application in other domains, e.g., statistics or financial analysis.

We integrate MOTION with the HyCC compiler [17] that generates optimized circuits for hybrid MPC protocols. With this, MOTION can directly perform efficient MPC of functionalities that have been specified in the C programming language. This enables MPC also for developers with limited domain knowledge and for a large range of applications.

We provide a detailed performance evaluation of both the low-level building blocks and protocol parts, as well as our full-scale applications on large data sets. Our performance results provide new insights, such as better performance of BMR compared to GMW (often even in low-latency networks) for deep natural circuits, e.g., for integer division, which could be of independent interest for protocol designers.

Paper Organization and Our Contributions

We summarize related work in §2 and provide preliminary information about notation, our setting, and security assumptions in §3. In §5, we describe the used MPC building blocks in detail. Our main contributions are the following:

MOTION framework for mixed-protocol MPC. In §4, we describe the design rationales behind our MOTION framework. It is a well-engineered and modular framework for MPC which is extensively tested and well-documented. We aim at high usability and will release our code[1] as open-source software under the permissive MIT license[2]. Our novel important features are the asynchronous evaluation of gates and oblivious transfers, which allows for a high level of abstraction in designing the protocols and new circuit evaluation strategies. Full communication serialization allows the use of MOTION in server and web applications. Our implementation can optionally interleave ('pipeline') the evaluation of the input-independent setup and the input-dependent online phase and allows to implement arbitrary circuit evaluation strategies in an abstract way.

[1]The code will be published shortly at https://encrypto.de/code/MOTION.
[2]https://choosealicense.com/licenses/mit/

Full-threshold hybrid MPC with passive security. We implement the existing full-threshold passively secure MPC protocols Boolean and arithmetic GMW [33], and BMR [8, 11] for multiple parties that guarantee a high level of security because all but one party can be corrupted (full-threshold). In §6, we introduce optimizations of these protocols, including an improvement of precomputed correlated OTs (C-OTs) from 2 to 1 rounds. This is of independent interest and makes the direct use of precomputed C-OTs for AND gates in GMW even more efficient than the use of Multiplication Triples (MTs) [6], having slightly lower communication and takes 1 instead of 2 rounds in the input-independent setup phase. In §7, we provide efficient protocols for converting between all of the above mentioned MPC protocols. MOTION is the first framework that efficiently combines these protocols. We support direct processing of hybrid circuits generated using the HyCC compiler [17].
Performance and Applications. In §8, we evaluate the performance of MOTION and demonstrate its practical relevance by showing that securely computing real-world applications such as biometric identification, AES-128, SHA-256, and convolutional neural network inference with N parties and full-threshold security is highly efficient, especially with our protocol conversions.
We compare MOTION's performance to other full-threshold frameworks. For biometric matching with N=3 parties, MOTION outperforms the passively secure implementations in MP-SPDZ [48] by 10.4×–17.6×, and the actively secure implementation in SCALE-MAMBA [2] by more than two orders of magnitude.

2 RELATED WORK

Practical MPC has been a very active field of research, especially in the past decade. Here, we provide an overview of the results most relevant to MOTION. Besides several theoretical foundations, we also discuss related implementations.

2.1 Theoretical Foundations

Yao's garbled circuits [72] and the protocol by Goldreich, Micali and Wigderson (GMW) [33] were the seminal works that introduced MPC with two and multiple parties, respectively. The protocol of Beaver, Micali and Rogaway (BMR) [8] can be seen as a multi-party variant of Yao's protocol. We provide a more detailed overview of these protocols in §6. We denote Yao's protocol and BMR with Y, the GMW protocol using Boolean sharing with B and arithmetic sharing with A.

2.2 MPC Implementations and Frameworks

A thorough overview and categorization of MPC implementations is given in [36]. In this section, we summarize *mixed-protocol* MPC implementations (that support multiple protocols) and compare them with MOTION in Tab. 1.

2-Party Solutions. Fairplay [55] was one of the first works that showed practical feasibility of MPC by providing an implementation of Yao's protocol. The TASTY framework [40] was the first mixed-protocol framework, and it combined Yao's GCs and Homomorphic Encryption (HE). The ABY framework [29] is a mixed-protocol framework for secure two-party computation based on Oblivious Transfer (OT) [4]. ABY showed that using OT yields better efficiency than using HE in the online phase. Researchers from Baidu have

Table 1: Related mixed-protocol MPC frameworks with N parties, threshold t, and active (●) or passive (◑) security. We denote the license of unpublished source code as '—'.

Framework	N	t	Security	Protocols	License
ABY [29]	2	1	◑	$A/B/Y$	LGPL-3.0
PrivC [39]	2	1	◑	A/Y	—
TASTY [40]	2	1	◑	A/Y	no license
EzPC [19]	2	1	◑	A/Y	MIT
OPA Mixing [46]	2	1	◑	any 2 of $A/B/Y$	MIT
ABY^3 [58]	3	1	● or ◑	$A/B/Y$	MIT
Sharemind [14]	3	1	● or ◑	A/B	payware[3]
ASTRA [20]	3	1	● or ◑	A/B	—
BLAZE [60]	3	1	● or ◑	A/B	—
Trident [62]	4	1	●	$A/B/Y$	—
SCALE-MAMBA [2]	≥ 2	$N-1$	●	A/Y	MIT-like
MP-SPDZ [48]	≥ 2	$N-1$	● or ◑	A/B or Y	MIT-like
MOTION (this work)	≥ 2	$N-1$	◑	$A/B/Y$	MIT

recently re-implemented two-party arithmetic and Yao sharing protocols from ABY in their product-level framework PrivC [39]. Recently, a hybrid solution partitioned protocols into a part that is executed using classical MPC primitives and a part that is evaluated in an Intel SGX enclave [21].

3- and 4-Party Solutions. ABY^3 [58] is a mixed-protocol implementation with a focus on privacy-preserving machine learning with exactly 3 parties. BLAZE [60] and ASTRA [20] further improve upon the performance of ABY^3 in the same setting. Trident [62] proposes hybrid 4-party protocols. Sharemind [14] is a framework for both integer arithmetic and Boolean operations. All of these frameworks only allow up to a single corruption, whereas MOTION provides full-threshold security, which is stronger given the same adversary model (cf. §3.4). A comparison between different number of corruptions against non-matching adversary models (e.g., passively secure full-threshold vs. actively-secure honest-majority protocols) is non-trivial and out of scope of this work.

N-Party Solutions. FairplayMP [10] is an extension of Fairplay that implements the BMR protocol [8] with the setup phase computed using the honest-majority BGW protocol [12]. Choi et al. [22] provide an N-party passively secure Boolean GMW implementation. SDPZ [24] is an actively secure MPC protocol based on arithmetic sharing in prime fields. SPD$\mathbb{Z}_{2^k}$ [23] introduced integer computations modulo 2^k for this approach. [25] gave an efficient implementation of SPD$\mathbb{Z}_{2^k}$ with applications to machine learning. Zaphod [3] allows to efficiently combine BMR [8, 38] with the SPDZ protocol, which is integrated in the SCALE-MAMBA implementation [2]. However, SCALE-MAMBA currently implements only actively secure MPC protocols, which have substantially higher overhead than passively secure ones. An alternative implementation of the SPDZ protocol is provided by MP-SPDZ [48] that also includes implementations of other protocols. MP-SPDZ has recently also included conversions between arithmetic and Boolean sharing. MPC protocols with a large number of parties were implemented in [71] as part of the EMP toolkit [69, 70], which contains also implementations of other MPC protocols. The BMR protocol was implemented in [11].

[3]Only a Sharemind-*emulator* is available for free.

2.3 Compilers for MPC Protocols

Another line of research focuses on directly compiling existing code into an MPC protocol. The PCF compiler [51] processes C code and creates a compact intermediate circuit format that was evaluated by an interpreter for the actively secure Yao-based two-party protocol of [52]. Wysteria [63] is a multi-party framework that implements the GMW protocol and offers type-based security and correctness checks. Frigate [59] is a verified compiler for creation of circuits for MPC protocols that can be securely evaluated with MPC implementations. PICCO [74] is a source-to-source compiler for C programs that builds on arithmetic secret sharing using bit decomposition for bit operations. Obliv-C [73] compiles a special-purpose C-based language to plain C for evaluating it in Yao's garbled circuits. EzPC [19] is a secure 2-party computation framework that allows to generate an efficient partitioning for mixed computations based on Yao's garbled circuits and arithmetic sharing. The authors of [46] show efficient algorithms for computing an optimal partitioning for mixing any two MPC protocols. The HyCC compiler [17] allows compilation of C code into efficient mixed-protocol MPC. MOTION directly supports the evaluation of circuits generated by HyCC. A combination of private memory access using Oblivious RAM (ORAM) and Yao's garbled circuits is implemented in ObliVM [53] that compiles programs from a Java-like language.

3 PRELIMINARIES

In this section, we provide background information about our setting, the adversary model, and define the notation.

3.1 Notation

We abbreviate $[i] = \{1, \ldots, i\}$. We denote the number of parties as N and the parties themselves as $P_1, \ldots, P_N$. A value x that is shared between N parties, is denoted as tuple $\langle x \rangle^S = (\langle x \rangle_1^S, \ldots, \langle x \rangle_N^S)$, where the superscript $S \in \{A, B, Y\}$ denotes the sharing type (cf. §6), and the subscript $i \in [N]$ denotes the i-th share of x that is held by party P_i. We write $\langle \mathbf{x} \rangle^B$ or $\langle \mathbf{x} \rangle^Y$ in bold font for a vector of ℓ shared bits, which we interpret as an ℓ-bit unsigned integer or element of $\mathbb{Z}_{2^\ell}$. $\text{Share}_i^S(x)$ denotes party P_i sharing their private value x in sharing S with all parties, while $\text{Share}^S(x)$ denotes that all parties create a sharing of a public value x in sharing S. $\text{Rec}_i^S(\langle x \rangle^S)$ denotes the reconstruction of a shared value $\langle x \rangle^S$ such that only party P_i receives x and $\text{Rec}^S(\langle x \rangle^S)$ denotes the reconstruction of a shared value $\langle x \rangle^S$ such that all parties receive x. We use the symmetric security parameter κ. With $x \in_R X$ we denote that x is drawn uniformly at random from the set X. We denote $\mathbf{x}[i]$ as the i-th element from $\mathbf{x}$.

3.2 Secure Multi-Party Computation (MPC)

MPC protocols are run by N parties and typically divided into a *setup phase*, that is independent of the parties' inputs and can be precomputed, and an *online phase*, that starts when the parties supply their private inputs.

3.3 Outsourcing Scenario

Alternatively to running our protocols directly between the N parties, they can also be used in an outsourcing scenario in a natural

way. As described in more detail in [47], in an outsourcing scenario an arbitrary number of input parties secret-share their private data to N non-colluding computing parties, who have no insight into that data. These computing parties evaluate our N-party MPC protocols and send the secret-shared output to a set of output parties, who can then reconstruct the plaintext output. Input and output parties can be (partially) identical or distinct. This allows for a large number of input/output parties without significantly increasing communication complexity of the protocols, since input sharing and output reconstruction are cheap 1-round operations and also provide security against actively corrupted input/output parties. The number of computing parties N can be chosen in accordance to performance and security requirements.

3.4 Adversary Model

The protocols we consider in this work are secure against passively corrupted (semi-honest) adversaries that follow the protocol specification but try to infer additional information about the other parties' private inputs by inspecting the protocol transcript. This passive attacker model is useful in scenarios where the involved parties trust each other but are legally constrained to keep information confidential, e.g., when computing statistics on sensitive health records. In the outsourcing scenario (cf. §3.3), a prime use case for our protocols, the computing parties are trusted to not collude with each other and run in a secured network. Discovering active attacks would lead to an immense loss of reputation and would hurt the business model of offering privacy-preserving services. From a research perspective, advances in the passive security model often lead to advances in stronger adversary models and serve as a performance baseline to show general feasibility of MPC-based applications. Moreover, techniques like attestation and several protocol extensions can be used for ensuring security against stronger adversaries, which we leave as future work.

4 ARCHITECTURE OF OUR MOTION FRAMEWORK

MOTION implements mixed-protocol MPC with an *arbitrary* number of $N \geq 2$ parties with full-threshold security against passive adversaries. Since communication complexity inherently scales with N, the focus of this work is to achieve practical performance for relatively small N, e.g., $N \leq 16$, as also commonly used in an outsourcing setting (cf. §3.3). Our framework allows to implement MPC protocols in different security models by design, such as honest majority and/or actively secure protocols.

MOTION is implemented in C++ and uses many of the modern features that were introduced in the C++17 standard. In the first place, it is a library and can easily be used in external projects. Our implementation has only few dependencies, making it OS-independent and fully compatible with the very liberal MIT license. MOTION requires only the following third-party libraries: `Boost`[4] (for network communication, logging, parsing command line arguments, fibers, and statistics), `flatbuffers`[5] (for communication serialization), `fmt`[6] (for string processing), optionally `Google Test`[7] (for unit and component tests), and `OpenSSL`[8] (for cryptographic primitives). MOTION does not depend on any third-party OT or MPC libraries. We take the software development process of our framework very seriously and designed the whole system with great care. Extensive component tests are included for all of the framework parts in order to ensure the correctness and security of our implementation. We routinely test MOTION for memory leaks. We will make our code public on GitHub[9] and will actively support and extend it in the future. Currently, our codebase consists of 179 source files, totaling in 36 000 lines of code, 20% of which are tests.

4.1 Novel Design Aspects

Besides designing MOTION with great care, we enhance its architecture by including several novel design aspects that improve its usability and facilitate new use cases. Although a few of our design aspects have at least partially been addressed in previous frameworks, *their combination* is novel. The extensibility of our framework paves the way to integrate also other optimizations such as different circuit evaluation strategies and new MPC protocols in the future.

4.1.1 Communication Serialization. As mentioned by Shai Halevi in his keynote talk at ACM CCS'18 [35], the requirement of MPC frameworks to own a TCP socket has in the past hindered their adoption, e.g., in server applications which often run under constrained permissions or have proprietary messaging interface. This restriction is solved in MOTION by using communication serialization. In MOTION, *all* the transferred messages are serialized and contain metadata sufficient to make messages recognizable without having to rely on an order preserving channel, e.g., a TCP socket. To the best of our knowledge, MOTION is the first MPC framework, whose communication is completely serialized. The most important benefit of the communication serialization is that our framework neither needs to own a separate TCP connection, nor does it rely on TCP (or similar protocols) as transport protocol, as it was the case for all previous MPC frameworks. Also, communication serialization allows for MPC to be based on low-latency network protocols (e.g., QUIC [42] or RUDP [43]), which can yield substantial performance improvements in MPC as shown in [16, 67], and also facilitates the real-world MPC use in previously infeasible scenarios, such as MPC in many proprietary networks, Web Services, Remote Procedure Calls (RPCs), or even via peer-to-peer messengers without establishing separate connections for the MPC framework. In order to demonstrate and evaluate the functionality of our framework we use `Boost` for TCP connections between the parties, but we stress that exchanging the communication parts with a different networking protocol is intended by design and would only involve minimal code changes.

Furthermore, we use `flatbuffers`' schema files to define how the serialized communication is structured in MOTION. The schema

[4] https://www.boost.org
[5] https://github.com/google/flatbuffers
[6] https://github.com/fmtlib/fmt
[7] https://github.com/google/googletest
[8] https://www.openssl.org
[9] https://encrypto.de/code/MOTION

files are independent of the programming language and can be compiled to a variety of different programming languages such as Java, Go, JavaScript, Rust, and even Swift[10]. These schema files significantly reduce the overhead of implementing MOTION or parts of it in a different programming language, e.g., to enable MPC on iPhones, while still supporting the original messaging interface. This approach enables multi-party protocols communication between our original C++ framework and many heterogeneous devices, operating systems, and programming languages.

4.1.2 Provider-based use of MPC Primitives. In our framework, the developer does not need to know how the MPC primitives interact with the network stack. Instead of the synchronized use of different MPC primitives on a network interface directly, we provide convenient provider interfaces for Oblivious Transfer (OT), Multiplication Triples (MTs) [6], Shared Bits (SBs), and Square Pairs (SPs), where the requirements for a program run can be registered, and are then automatically handled without any further action from the developer. The providers return pointers to the registered objects, which provide a separate interface to execute the desired protocol, e.g., set inputs to the OT functionality or wait for a batch of MTs to be computed.

Besides the convenient user interface, our providers enable the developer to easily replace the computation procedure of a primitive partially or completely, e.g., use a semi-trusted third party that generates correlated randomness (e.g., MTs) and distributes it among the computing parties instead of computing it using expensive crypto (cf. [64]).

4.1.3 Multiple Layers of Abstraction. MOTION is both developer-friendly and function-rich. On the one hand, it provides a convenient way of developing MPC solutions without significant MPC knowledge using secure type classes, which provide overloaded C++ operators and can be used just as the classic C++ types. Moreover, all of our most abstract APIs operate directly in C++. This simplifies error handling and omits the need for the developer to learn a new domain-specific language that is accepted by the framework. On the other hand, the developer can also use our API to get access to the low-level MPC primitives and analyze or modify the underlying routines, e.g., add new optimizations or even use the MPC primitives standalone. The latter allows for using our framework, e.g., to only compute base OTs or OT extension. When using our MPC protocols or the underlying primitives, developers do not require any knowledge about how the protocols or primitives work together or interact with the message passing interface. To the best of our knowledge, MOTION is the first MPC framework that provides such a high level of abstraction while allowing to use all primitives directly. The other MPC frameworks either translate a special-purpose language to circuits (e.g., [17]), or low-level code (e.g., [73]), or they process commands/circuits by a compiled interpreter (e.g., [2]).

We give a small code snippet in List. 1 to illustrate the simplicity of the code in MOTION. Note that the example depicted in List. 1 is optimized for efficiency, and also a straightforward implementation would be fully functional but less efficient. For the convenient use of MOTION by non-experts in MPC, we recommend to utilize our HyCC adapter for efficiency reasons (cf. §4.1.4). Also, we show an advanced example of functionality extension of MOTION in List. 2 in App. A, where we implement a secure two-party protocol for multiplying an integer known by one party by a secret-shared bit.

[10] https://github.com/mzaks/FlatBuffersSwift

Listing 1: Code excerpt for efficiently computing minimum squared Euclidean distance in MOTION.

```
using namespace MOTION;
using suint = SecureUnsignedInteger;
using vec = std::vector<suint>;

// variable a is an arithmetic GMW share
// variable v is a vector of arithmetic GMW shares
// computes squared Euclidean distance between
// a and each element in v
// returns BMR share of min. sqr. Euclidean distance
suint MinSqrEuclideanDistance(suint& a, vec& v){
  vec res(v.size()); // result vector

  // automatic use of more efficient squaring
  auto a_sqr = a*a, two_a = 2*a;

  // compute squared Euclidean distance
  // (a-v[i])^2 = a^2 + v[i]^2 - 2a*v[i]
  for(unsigned int i = 0; i < res.size(); ++i)
    res[i] = a_sqr + v[i]*(v[i] - two_a);

  // convert each distance to BMR sharing
  for(auto& e : res)
    e = e->Convert<MPCProtocol::BMR>();

  // select initial minimum as 0-th element
  auto min = res[0];

  // find the minimum distance
  for(unsigned int i = 1; i < res.size(); ++i){
    auto smaller = res[i] < min;
    //    smaller ? res[i] : min
    min = smaller.MUX(res[i], min);
  }
  return min;
}
```

4.1.4 HyCC Integration. If the use of an abstract language instead of the direct use of C++ classes is desired, e.g., if the developer has no expert knowledge in MPC and thus is not able to manually implement and optimize MPC protocols, the developer can import efficient hybrid circuits generated by the HyCC compiler [17] from a subset of the C programming language using our HyCC adapter. MOTION fully supports the features of HyCC and follows its partitioning guidelines, which results in protocols that are tailored according to a user-specified optimization goal. *Previous works:* HyCC is integrated in the ABY framework [29] for N=2-party MPC.

4.1.5 Asynchronous Gate Evaluation. MOTION allows the secure evaluation of arbitrary circuits without additional information about their structure. Each gate depends on its parents, and some gates require network communication for their evaluation. Thus, we are faced with a complex dependency graph consisting of possibly millions of interdependent tasks, which need to be scheduled on the available CPU cores. Evaluation of each gate in a separate thread is clearly infeasible if not impossible due to constraints of the operating system. Our solution is to use fibers, i.e., threads implemented in userspace, which are run on a fixed number of

worker threads. For this, we use the Boost.Fiber library[11]. When a fiber is blocked, e.g., because a message has not yet been received, the worker thread switches to a different fiber and continues to evaluate a different gate. Compared to the overhead of a context switch between threads on the operating system level, switching between fibers is lightweight and possible in less than 100 CPU cycles, which is typically about an order of magnitude less than for a context switch between threads. Another advantage is that fibers can be used in the same way as usual threads. Thus, the implementation of a protocol is straightforward since the developer has access to all common synchronization mechanisms. We adapted a work-stealing scheduling algorithm from Boost.Fiber to our own thread pool which creates one worker thread per CPU core by default. Also, we designed and implemented a number of efficient synchronization primitives and asynchronous access mechanisms for fibers to handle interactions between gates and MPC primitives. Note that we do not put any constraints on the number of physical processor cores used by a party. Moreover, MOTION allows parties to run with a different number of threads, which is a common problem of synchronized MPC where the work is scheduled for a static number of threads and/or communication channels before the protocol evaluation, and an inconsistent number of threads either yields wrong results or causes a program crash, e.g., [29, 65].

We evaluate all gates separately as soon as their parent gates become ready. In the following, we highlight two benefits of this approach. Firstly, the asynchronous gate evaluation decreases the online time for evaluating unbalanced circuits. Consider a scenario with high network latency and two major subcircuits, as shown in Fig. 1: one consisting of only few data-dependent layers with many costly non-XOR gates (blue), and the other subcircuit containing much fewer non-XOR gates but consisting of a large number of interactive layers (green). If evaluated layer-wise (as it is done in many current MPC frameworks), the large subcircuit has a blocking effect on the deep circuit due to the longer evaluation time. In our framework, the default scheduler evaluates gates in first come first serve order. On the other hand, the possibility to replace the default scheduler by a custom one with a different evaluation strategy is intended by design. The goal of a custom scheduler can be to prioritize gates according to the maximum subcircuit depth (i.e., gates that lead to the deepest subcircuit are evaluated first) to minimize communication latency, or to synchronize the evaluation layer-wise to evaluate in a batch all gates in a layer to possibly save bandwidth.

This asynchrony is especially beneficial in networks with high latency, e.g., for trans-continental Internet connections. Also, we batch operations for all input wires and the contained SIMD values of each gate in order to reduce communication. Secondly, integration of new protocols into MOTION becomes easier, since gate evaluation is independent of other protocols by design (in contrast to most existing MPC frameworks). Thus, the developer does not need to know how the other parts of the framework work to integrate the new protocols. *Previous Works.* Asynchronous gate evaluation was implemented in the no longer supported VIFF framework [26] using callbacks in Python. The callbacks, however, made the code

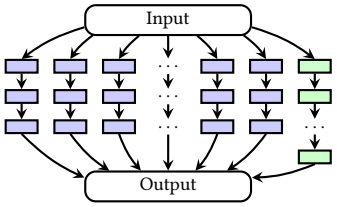

Figure 1: Circuit with a large number of parallel gates (blue) and many data-dependent, sequential gates (green), that benefits from asynchronous gate evaluation.

unnecessarily complicated. The more recent VIFF derivate MPyC[12], uses native coroutines in Python. However, Python is a suboptimal choice for highly efficient MPC, because it is a scripting language and thus substantially less efficient than lower-level programming languages such as C and C++.

4.1.6 Code Vectorization. We design our code with vectorization of CPU instructions in mind to improve its efficiency. This goal is different from the MPC-level SIMD instructions (cf. §4.1.7) and affects the compiled code directly. Yet, explicit vectorization using architecture-specific instructions would limit the number of architectures MOTION supports. To achieve both better efficiency through (better) vectorization of CPU instructions and rich support for various architectures, most of our code is optimized such that the compiler can vectorize it automatically using the native instructions of the underlying architecture. This is achieved by using multiple techniques such as eliminating loop dependencies and branching, enforcing buffer alignment that matches the cache line size, and giving the compiler various hints to produce better code, e.g., using the `restrict` type qualifier. The few SSE instructions we used are supported by many architectures, and for those that do not support them, we automatically provide a slightly slower pure C++ code drop-in replacement. *Previous Works.* To the best of our knowledge, for today the only MPC framework implemented in a low-level programming language with possibility of cross-platform compilation is ABY [29].

4.1.7 Single Instruction Multiple Data (SIMD). We intentionally design the API of MOTION in a way that encourages the use of MPC-level Single Instruction Multiple Data (SIMD) instructions that process vectors of data instead of single data, e.g., vectors of bits instead of a single bit. This not only drastically reduces the memory footprint but also the required communication, since sending 1-bit values has significant overhead. This optimization results in a much better amortized efficiency and throughput, which we detail in §8. SIMD instructions are especially relevant for the outsourcing setting, where the outsourcing servers often simultaneously process data of many users. In our experiments, extensive use of SIMD instructions improved the throughput of MOTION by about an order of magnitude in both the LAN and WAN setting (cf. Tab. 3). *Previous Works.* SIMD instructions have been used for $N \in \{2, 3\}$ in [14, 29, 66].

[11] https://www.boost.org/doc/libs/1_73_0/libs/fiber/

[12] https://github.com/lschoe/mpyc

4.1.8 Interleaved Setup and Online Phase. Circuit evaluation in MPC happens in one of two modes: sequential or interleaved ('pipelined'). The sequential mode runs the input-dependent online phase only after the input-independent setup has completed. This allows precise measurements of the setup and online phase communication and computation requirements, or full precomputation ahead of the online phase. Frameworks like ABY [29] support only this evaluation mode. The interleaved mode, on the other hand, runs both phases in parallel and facilitates possibly more efficient evaluation of the circuit in terms of load balancing, since the gates that otherwise would have been waiting for the setup phase to finish can be evaluated faster, thus improving the protocol latency. *Previous Works.* A similar approach is used in SCALE-MAMBA [2]. However, their approach often overproduces correlated randomness in the setup phase, which is disadvantageous for small applications. In contrast, MOTION produces exactly the required amount of correlated randomness. To our knowledge we are the first framework to offer both sequential and interleaved circuit evaluation, giving the user the freedom of choice according to the use case.

4.2 Implemented Building Blocks

To make MOTION easy to use and extend, we design it in a completely different way compared to prior work. We did not use any existing implementations of the cryptographic protocols, since they would need to be completely redesigned. Below, we list the main components implemented in our framework that can also be of interest for other applications.

- OT* by Hauck and Loss [37] for base OTs (cf. §5.1.1).
- Providers for OT extension including Beaver's OT precomputation [7] and different OT flavors [5] (cf. §5.1.2): general, random (cf. §5.1.3), additively correlated, and XOR-correlated OTs (C-OTs) including our optimization of precomputed C-OT w.r.t. round complexity (cf. §5.1.3).
- Providers for Multiplication Triples (cf. §5.2), Squaring Pairs (cf. §5.3), and Shared Bits (cf. §5.4) using C-OTs.
- N-party full-threshold passively secure MPC protocols: Arithmetic GMW (cf. §6.1), Boolean GMW [33] (cf. §6.2), and BMR [8, 11] (cf. §6.3), and secure conversions between them (cf. §7).
- Plenty of utility classes, e.g., adapted `Boost.Fiber` for fibers, (aligned) bit vector, bit span, logger, run-time and communication statistics over multiple runs, function-encapsulating conditions, and reusable promises and futures.
- Unit and component tests for all of the implemented components and most of the utility classes.
- Application examples (cf. §8.2) that use MOTION as a library, e.g., MPC protocols for AES-128, SHA-256, minimum Euclidean distance, and Convolutional Neural Networks (CNNs).

5 MPC BUILDING BLOCKS

In this section, we describe the primitives that our protocols rely on, as well as our improvements to them.

5.1 Oblivious Transfer

Oblivious Transfer (OT) [61] is the basic building block for various generic and custom MPC protocols. It involves two parties, a sender $\mathcal{S}$ that inputs two messages (m_0, m_1), and a receiver $\mathcal{R}$ that inputs the choice bit $c \in \{0,1\}$. The functionality outputs $\perp$ to $\mathcal{S}$ and only m_c to $\mathcal{R}$. It is guaranteed, that $\mathcal{S}$ does not learn c and that $\mathcal{R}$ does not learn m_{1-c}. This kind of OT is called 1-out-of-2 OT and it can be generalized to 1-out-of-n OT where $\mathcal{S}$ inputs $(m_0, \ldots, m_{n-1})$, $\mathcal{R}$ inputs c, and the functionality outputs $(\perp, m_c)$. Inherently, OT requires public-key cryptography [44], which is computationally expensive. However, the important *OT extension* technique proposed by Ishai et al. [45] allows to use a small number of "base OTs" and use them as seeds to compute a much larger number of OTs using significantly faster symmetric cryptography. OT can also be precomputed [7], moving a significant part of computation and communication from the online phase to the input-independent setup phase (cf. §3.2). To precompute an OT, $\mathcal{R}$ starts the OT extension protocol using a random $r \in_R \{0,1\}$. In the online phase, $\mathcal{R}$ sends $p := r \oplus c$ to $\mathcal{S}$, who swaps the messages if $p = 1$ and does nothing otherwise. Then, the parties proceed as in the original OT extension protocol. OT precomputation adds one sequential message to the OT extension protocol, thus increasing the number of communication rounds.

5.1.1 Base OTs. Abstractly speaking, the base OTs are computed as follows: for each $j \in [\kappa]$ (where κ is the symmetric security parameter, e.g., 128 bit) $\mathcal{S}$ inputs $(s_{j,0}, s_{j,1}) \in_R \{0,1\}^{2\kappa}$ and $\mathcal{R}$ inputs $c_j \in_R \{0,1\}$. $\mathcal{R}$ obtains s_{j,c_j} for each $j \in [\kappa]$. In this work, we use the base OT protocol by Hauck and Loss [37] (denoted as OT*). We use it in the random OT setting, i.e., (1) the choice bits of $\mathcal{R}$ are random, and (2) we emit the last step of the protocol for sending the messages to $\mathcal{R}$.

5.1.2 OT Extension. We use the OT extension protocol by Ishai et al. [45] with optimizations from [4, 5] and denote it as General OT (G-OT). The protocol is defined as follows. First, the parties run a base OT protocol with inverted roles. In the setup phase, $\mathcal{R}$ uses the sent messages from base OTs to generate $\mathbf{T} \in \{0,1\}^{m \times \kappa}$ with $\mathbf{T}[j] = \mathrm{PRG}(s_{j,0})$ for each $j \in [\kappa]$ and sets $u_j = \mathrm{PRG}(s_{j,1}) \oplus \mathbf{T}[j] \oplus \mathbf{r}$, where m is the number of required OTs, $\mathbf{r}$ are $\mathcal{R}$'s real choices, and PRG is a pseudo-random generator. Then, $\mathcal{S}$ creates $\mathbf{V} \in \{0,1\}^{m \times \kappa}$ with $\mathbf{V}[j] = c_j u_j \oplus \mathrm{PRG}(s_{j,c_j})$ for each $j \in [\kappa]$, where c_j is the choice bit in the j-th base OT. Finally, both parties transpose their matrices: $\mathcal{S}$ sets $\mathbf{V}' = \mathbf{V}^T$ and $\mathcal{R}$ sets $\mathbf{T}' = \mathbf{T}^T$.

In the online phase, $\mathcal{S}$ sends to $\mathcal{R}$ $y_{i,0} := x_{i,0} \oplus H(i, \mathbf{V}'[i])$ and $y_{i,1} := x_{i,0} \oplus H(i, \mathbf{V}'[i] \oplus \mathbf{c})$ for each $i \in [m]$, where $H(\cdot)$ is a one-way pseudo-random function. $\mathcal{R}$ sets the output of the OT $i \in [m]$ as $x_{i,r_i} := y_{i,r_i} \oplus H(i, \mathbf{T}'[i])$. We instantiate both PRG and H using AES (cf. §5.5).

5.1.3 OT Flavors. In many cases, OT-based MPC protocols need to compute very specific functionalities using OT. Asharov et al. [4, 5] have first shown that OT extension can be done significantly more efficiently for specific tasks.

Random OT (R-OT). Random OT can essentially be seen as a truncated OT extension protocol with no inputs. The parties run the same protocol steps as in OT extension, but omitting the last step where $\mathcal{S}$ masks his messages and sends them to $\mathcal{R}$. Instead, the parties only compute their masks and set them as the output of the protocol. Slightly more formally, $\mathcal{S}$ sets $(H(i, \mathbf{V}'[i]), H(i, \mathbf{V}'[i] \oplus \mathbf{c}))$ and $\mathcal{R}$ sets $H(i, \mathbf{T}'[i])$ as his output. R-OT can be used to compute

other OT functionalities such as G-OT and correlated OT (C-OT) or to compute MTs in 2PC [5].

Correlated OT (C-OT). Correlated OT is a special OT flavor that is very well-suited for MPC. Its main use case is multiplication of a (secret-shared) bit or string by a (secret-shared) bit yielding a secret-shared multiplication result. The functionality of C-OT is as follows: $\mathcal{S}$ inputs bit-string x and $\mathcal{R}$ bit r. The functionality outputs $(x_0, x_0 \odot x)$ to $\mathcal{S}$ and $x_0 \odot rx$ to $\mathcal{R}$, where x_0 is random and $\odot$ is usually bit-wise XOR, which we denote as XOR-correlated C-OT ($C^{\oplus}$-OT), or addition mod 2^ℓ, which we denote as additively correlated C-OT (C^{+}-OT). The $C^{\oplus}$-OT results in an XOR-sharing of the multiplication, and C^{+}-OT is an additively shared multiplication. The difference to G-OT is that instead of sending two masked messages, $\mathcal{S}$ sends only one message $y_i = x_{i,1} \odot H(i, \mathbf{V}'[i] \oplus \mathbf{c}) = x_{i,0} \odot x_i \odot H(i, \mathbf{V}'[i] \oplus \mathbf{c}) = H(i, \mathbf{V}'[i]) \odot x_i \odot H(i, \mathbf{V}'[i] \oplus \mathbf{c})$ and $\mathcal{R}$ sets $x_{i,r_i} = r_i y_i \odot H(i, \mathbf{T}'[i])$ for each $i \in [m]$.

Our optimization for precomputed C-OT. We introduce a generic factor two improvement of the online round complexity of precomputed C-OT. To simplify the description of the C-OT precomputation, we reuse the R-OT protocol. First, the parties compute R-OT with the choice bit $r \in_R \{0,1\}$. In order to obtain the correct message mask for his real choice c, $\mathcal{R}$ sends $p = 0$ if $r = c$ and $p = 1$ otherwise. After obtaining p, $\mathcal{S}$ swaps the messages iff $p = 1$, and proceeds with the original protocol. Our improvement is based on the observation that in C-OT (in contrast to G-OT) the message of $\mathcal{S}$ is equal in both cases ($p = 0$ and $p = 1$), i.e., $y_i := H(i, \mathbf{V}'[i]) \odot x_i \odot H(i, \mathbf{V}'[i] \oplus \mathbf{c}) = H(i, \mathbf{V}'[i] \oplus \mathbf{c}) \odot x_i \odot H(i, \mathbf{V}'[i])$. However, p may change the *output of* $\mathcal{S}$. This is why $\mathcal{S}$ sends y_i independent of the choice bit and waits for receiving p to determine his own correct output, whereas $\mathcal{R}$ sends p, waits for y_i and unmasks it. Our improved precomputed C-OT protocol results in two messages that are sent *independently*, which improves the online phase latency by a factor of 2. Its latency is equal to the original C-OT protocol without precomputation.

5.2 Multiplication Triples

Multiplication triples (MTs) proposed by Beaver [6] allow to reduce the online complexity of MPC protocols. In fact, given a secret-shared MT in form of $\langle c \rangle^S = \langle a \rangle^S \langle b \rangle^S$, where party P_i holds $\langle a \rangle_i^S$, $\langle b \rangle_i^S$, and $\langle c \rangle_i^S$, and $S \in \{A, B\}$, the parties require only cleartext operations and very little communication in the online phase to privately compute a multiplication.

5.2.1 Arithmetic MTs (A-MTs). For ℓ-bit A-MTs, we generalize the C^{+}-OT-based A-MT generation protocol by Demmler et al. [29] from two to N parties. Namely, each party P_i locally generates two random shares $\langle a \rangle_i^A \in_R \mathbb{Z}_{2^\ell}$ and $\langle b \rangle_i^A \in_R \mathbb{Z}_{2^\ell}$, and the parties interactively compute $\langle c \rangle^A = \langle a \rangle^A \langle b \rangle^A = \sum_i (\langle a \rangle_i^A \langle b \rangle_i^A) + \sum_{i,j \neq i} (\langle a \rangle_i^A \langle b \rangle_j^A) \mod 2^\ell$. For each $P_{j \neq i}$, P_i and P_j run the following secure multiplication protocol to multiply $\langle a \rangle_i^A \cdot \langle b \rangle_j^A$ and $\langle a \rangle_j^A \cdot \langle b \rangle_i^A$. To perform a secure multiplication of two values $x, y \in \mathbb{Z}_{2^\ell}$ owned by P_i and P_j, respectively, they run ℓ parallel C^{+}-OTs. Here, P_i acts as the sender and inputs x to each C^{+}-OT and obtains $r_k \in \mathbb{Z}_{2^\ell}$ from the k-th C^{+}-OT with $k \in [\ell]$. P_j acts as $\mathcal{R}$ and for each $k \in [\ell]$ inputs the k-th bit of y (denoted by $\mathbf{y}[k]$) to the k-th C^{+}-OT, and obtains $\mathbf{y}[k]x + r_k$ as output. Finally, P_i sets $z_i = -\sum_{k=1}^{\ell} r_k 2^{k-1} \pmod{2^\ell}$ and P_j sets $z_j = \sum_{k=1}^{\ell} 2^{k-1}(\mathbf{y}[k]x + r_k) \pmod{2^\ell}$ with $z_i + z_j = x \cdot y \pmod{2^\ell}$. As observed in [29], this multiplication can be done with half the online communication for the C^{+}-OTs by not sending the most significant bits of the values that will be cut off by the bit-shifts in the subsequent computation. The total communication of ℓ-bit A-MT generation with N parties and symmetric security parameter κ is $\approx N(N-1)\ell(\kappa + \ell/2)$ bits, and requires two sequential messages.

5.2.2 Boolean MTs (B-MTs). In this work, we use $C^{\oplus}$-OT to compute B-MTs. The protocol is analogous to the one for A-MT computation (cf. §5.2.1) with the difference that here we use $C^{\oplus}$-OT instead of C^{+}-OT, and that we cannot utilize the 2× communication reduction. The total communication of the B-MT generation protocol is $N(N-1)(\kappa + 1)$ bits, and requires two sequential messages.

5.3 Square Pairs

In addition to A-MTs, we also compute *square pairs* (SPs), introduced in [27], which are pairs of random secret-shared values $\langle a \rangle^A, \langle c \rangle^A$ such that $\langle c \rangle^A = \langle a \rangle^A \cdot \langle a \rangle^A$. They can be generated analogously to A-MTs (cf. §5.2.1) but require only a single secure multiplication between each pair of parties and, thus, only half the number of C^{+}-OTs (and hence communication) compared to MTs in the same sharing.

5.4 Shared Bits

Another form of precomputation are *shared bits* (SBs) [27], which are arithmetic sharings $\langle b \rangle^A$ over $\mathbb{Z}_{2^\ell}$ of random bits $b \in \{0,1\}$. We generate these with an adapted version of Π_{RandBit} from [25], and use square pairs (cf. §5.3) to compute the required squaring in $\mathbb{Z}_{2^{\ell+2}}$ (cf. §6.1). Hence, $N(N-1)(\ell+2)/2$ C^{+}-OTs with additive correlation in the ring $\mathbb{Z}_{2^{\ell+2}}$ are needed per shared bit over $\mathbb{Z}_{2^\ell}$.

5.5 Fixed-Key AES

The implemented OT extension protocol (cf. §5.1.2) and the garbling scheme in the BMR protocol (cf. §6.3) make extensive use of hash and pseudorandom functions. Thus, if instantiated inefficiently, these can become the main bottleneck of MPC. Modern CPUs have dedicated instruction sets for performing cryptographic operations such as AES in hardware (e.g., AES-NI on x86). As these instructions are substantially faster than a software implementation, it is natural to utilise these primitives to speed up higher level protocols. Since the AES key schedule is still quite inefficient, constructions using a fixed key have been used for garbling schemes [9].

In our framework, we instantiate hash and pseudorandom functions for OT extension and BMR garbling (cf. §6.3) with fixed-key AES following the approach of Guo et al. [34] and Ben-Efraim et al. [11].

5.6 Bandwidth-Saving Broadcast

Broadcast is a common element among many of the protocols implemented here (cf. §6,§7). Often, each party sends some data of size ℓ bit to every other party, whereupon the shares are accumulated, e.g., via XOR. Using point-to-point channels, this results in

$N(N-1)\ell$ bit of total communication. Since we consider the passive security setting, we can reduce this to $2(N-1)\ell$ bit by letting everyone send its part to a designated party who performs the accumulation and broadcasts the result. Now the communication is no longer quadratic but instead linear in the number of parties. This comes at the cost of one additional round of communication.

6 MPC PROTOCOLS

In this section, we describe the established passively secure full-threshold MPC base protocols Arithmetic sharing (§6.1), Boolean sharing with GMW (§6.2), and Yao sharing with BMR (§6.3). We indicate their use in protocols as A, B, and Y, respectively. The costs of all primitive operations are summarized in Tab. 6 in App. A.

6.1 Arithmetic Sharing (A)

Arithmetic sharing is based on additive secret sharing in the ring $\mathbb{Z}_{2^\ell}$ where a value x is shared as $\langle x\rangle^A = \left(\langle x\rangle_1^A, \ldots, \langle x\rangle_N^A\right) \in \mathbb{Z}_{2^\ell}^N$ such that $x = \sum_{j=1}^N \langle x\rangle_j^A \pmod{2^\ell}$ and party P_i holds share $\langle x\rangle_i^A$. In the following, we assume a fixed bit length ℓ.

For input sharing $\mathsf{Share}_i^A(x)$, party P_i samples $\langle x\rangle_1^A, \ldots, \langle x\rangle_N^A \in_R \mathbb{Z}_{2^\ell}$ such that $x = \sum_{j=1}^N \langle x\rangle_j^A \pmod{2^\ell}$, and sends $\langle x\rangle_j^A$ to party P_j. The communication can be avoided by sampling $\langle x\rangle_j^A$ from the output of a PRG whose seed is known only to parties P_i and P_j. For output reconstruction $\mathsf{Rec}_i^A(\langle x\rangle^A)$, each party P_j sends $\langle x\rangle_j^A$ to party P_i who computes $x \leftarrow \sum_{j=1}^N \langle x\rangle_j^A$. For $\mathsf{Rec}^A(\langle x\rangle^A)$, each party P_j broadcasts $\langle x\rangle_j^A$ and computes $x \leftarrow \sum_{j=1}^N \langle x\rangle_j^A$. Alternatively, each P_j could send $\langle x\rangle_j^A$ to P_1, who reconstructs $x \leftarrow \sum_{j=1}^N \langle x\rangle_j^A$ and sends it back to all parties. This requires an additional round, but in total only $O(N\ell)$ instead of $O(N^2\ell)$ bits of communication, which can be used as a trade-off for low-latency networks with limited bandwidth.

Linear operations can be computed locally, i.e., without communication: E.g., for $\langle z\rangle^A \leftarrow a \cdot \langle x\rangle^A + \langle y\rangle^A + b$ with public $a, b \in \mathbb{Z}_{2^\ell}$, party P_1 computes $\langle z\rangle_1^A \leftarrow a \cdot \langle x\rangle_1^A + \langle y\rangle_1^A + b$, and all other parties $P_2, \ldots, P_N$ compute $\langle z\rangle_i^A \leftarrow a \cdot \langle x\rangle_i^A + \langle y\rangle_i^A$.

Multiplication can be computed using multiplication triples (MTs) (cf. §5.2.1): Let $(\langle a\rangle^A, \langle b\rangle^A, \langle c\rangle^A)$ be an MT for $\mathbb{Z}_{2^\ell}$. For $\langle z\rangle^A \leftarrow \langle x\rangle^A \cdot \langle y\rangle^A$, the parties compute $d \leftarrow \mathsf{Rec}^A(\langle x\rangle^A - \langle a\rangle^A)$, $e \leftarrow \mathsf{Rec}^A(\langle y\rangle^A - \langle b\rangle^A)$, and $\langle z\rangle^A \leftarrow \langle c\rangle^A + e \cdot \langle x\rangle^A + d \cdot \langle y\rangle^A - d \cdot e$. Multiplications can also be computed with less communication at cost of an additional round by applying the communication-saving reconstruction method described above to the computation of d and e.

Squares are computed more efficiently with Square Pairs (SPs) (cf. §5.3) [27] using only half of the communication: Let $(\langle a\rangle^A, \langle c\rangle^A)$ be an SP for $\mathbb{Z}_{2^\ell}$. For $\langle z\rangle^A \leftarrow \langle x\rangle^A \cdot \langle x\rangle^A$, the parties compute $d \leftarrow \mathsf{Rec}^A(\langle x\rangle^A - \langle a\rangle^A)$, and $\langle z\rangle^A \leftarrow \langle c\rangle^A + 2 \cdot d \cdot \langle x\rangle^A - d^2$.

6.2 Boolean Sharing with GMW (B)

Boolean GMW [33] uses XOR-based secret sharing, which is equivalent to additive secret sharing in the ring $\mathbb{Z}_2$, where addition and multiplication correspond to XOR ($\oplus$) and AND ($\wedge$), respectively. Hence, this is a special case of the arithmetic sharing (cf. §6.1) with $\ell = 1$. A value $x \in \{0,1\}$ is shared as $\langle x\rangle^B = \left(\langle x\rangle_1^B, \ldots, \langle x\rangle_N^B\right)$ such that $x = \bigoplus_{i=1}^N \langle x\rangle_i^B$ and party P_i holds $\langle x\rangle_i^B$. All basic operations are computed analogously to those in arithmetic sharing.

We write $\langle \mathbf{x}\rangle^B$ for a vector of ℓ shared bits interpreted as an ℓ-bit integer or element of $\mathbb{Z}_{2^\ell}$. In this context $\langle \mathbf{x}\rangle^B + \langle \mathbf{y}\rangle^B$ denotes addition and $\langle \mathbf{x}\rangle^B \cdot \langle \mathbf{y}\rangle^B$ denotes multiplication in $\mathbb{Z}_{2^\ell}$. Basic operations are done using depth-optimized Boolean circuits [28, 66].

More efficient AND Gates without MTs. The use of Boolean Multiplication Triples (B-MTs) instead of direct secure bit multiplication is motivated by their very cheap online phase with exactly *one* communication round, low communication, and only cleartext operations. Computation of a B-MT requires a secure bit multiplication, which makes $2\binom{N}{2}$ calls to precomputed C-OT (cf. §5.1.3) in the setup phase, and $\mathsf{Rec}^B(\langle e\rangle^B, \langle d\rangle^B)$ in the online phase (cf. §5.2.2). The secure multiplication protocol using our optimized C-OT also requires only one round in the online phase (instead of two with non-optimized C-OT), but only *one* round in the setup phase (instead of two for MT). Moreover, it also transfers exactly 4 bits in the online phase between each (P_i, P_j) with $i \neq j$, but compared to MTs 4 bits *less* in the setup phase.

6.3 Yao Sharing with BMR (Y)

The BMR protocol [8] is an extension of Yao's garbled circuits protocol [72] to the multi-party case. Instead of the garbled circuit being constructed by one party and evaluated by the other, it is garbled by all parties collaboratively. During the setup phase, the parties engage in a garbling protocol such that no set of up to $N-1$ parties gains enough information to recover any intermediate values in the resulting garbled circuit. In the online phase, the garbled circuit can be evaluated by each party without further communication. In the following, we use the notation by [11]. Moreover, we implement the free-XOR technique for BMR introduced by [11], which allows to evaluate XOR gates without any communication or cryptographic operations during the setup or online phase.

In the setup phase, each party P_i generates a global key offset $R^i \in_R \{0,1\}^\kappa$, and shares λ_w^i of random permutation bits $\lambda_w := \bigoplus_{j=1}^N \lambda_w^j$ and pairs of keys $k_{w,0}^i, k_{w,1}^i$ for each wire w in the circuit: If w is an input wire of the circuit such that party P_j provides that input, then $\lambda_w^i \in_R \{0,1\}$ if $i = j$ and $\lambda_w^i = 0$ otherwise. If w is *not* the output of an XOR gate, the share of the permutation bit $\lambda_w^i \in_R \{0,1\}$ and the key $k_{w,0}^i \in_R \{0,1\}^\kappa$ are chosen randomly. If w is the output of an XOR gate with input wires a, b, then $\lambda_w^i := \lambda_a^i \oplus \lambda_b^i$ and $k_{w,0}^i := k_{a,0}^i \oplus k_{b,0}^i$. The second key $k_{w,1}^i$ is in both cases implicitly defined as $k_{w,0}^i \oplus R^i$. If w is an output wire of the circuit, then all λ_w^i are sent to the party (or the parties) collecting that output.

Furthermore, the parties invoke the following garbling functionality $\mathcal{F}_{GC}$: It takes R^i and $\lambda_w^i, k_{w,0}^i, k_{w,1}^i$ for all wires w from each party P_i as inputs. Let F^2 be a double-key PRF and let $\circ$ denote concatenation in the following. We instantiate F^2 with a fixed-key AES construction [11, 34] (cf. §5.5). The garbling functionality $\mathcal{F}_{GC}$ computes for each AND gate g, and for all $j \in [N]$ and $\alpha, \beta \in \{0,1\}$:

$$\tilde{g}^j_{\alpha,\beta} \leftarrow \left(\bigoplus_{i=1}^{N} F^2_{k^i_{a,\alpha},k^i_{b,\beta}}(g \circ j)\right) \oplus k^j_{w,0} \oplus \left(R^j \cdot ((\lambda_a \oplus \alpha)(\lambda_b \oplus \beta) \oplus \lambda_w)\right).$$

It outputs $\tilde{g}^1_{\alpha,\beta} \circ \cdots \circ \tilde{g}^N_{\alpha,\beta}$ for all g and $\alpha, \beta \in \{0,1\}$.

In our framework, we instantiate $\mathcal{F}_{GC}$ with the OT-based protocol by [11] achieving a BMR instantiation with full corruption threshold. Their garbling protocol uses one bit $\mathrm{C}^{\oplus}$-OT and three correlated $\mathrm{C}^{\oplus}$-OTs of strings of length κ per pair of parties to generate the garbled tables of an AND gate with inputs a, b and output w:

First the parties securely compute $\lambda_{ab} := \lambda_a \cdot \lambda_b$ such that each party P_j receives a random share λ^j_{ab} using two 1-bit $\mathrm{C}^{\oplus}$-OTs per pair of parties. Then they locally compute $\lambda_{a\bar{b}w} := \lambda_a \cdot \overline{\lambda_b} \oplus \lambda_w$, by setting $\lambda^j_{a\bar{b}w} := \lambda^j_{ab} \oplus \lambda^j_a \oplus \lambda^j_w$, and analogously $\lambda_{\bar{a}bw} := \overline{\lambda_a} \cdot \lambda_b \oplus \lambda_w$ and $\lambda_{\bar{a}\bar{b}w} := \overline{\lambda_a} \cdot \overline{\lambda_b} \oplus \lambda_w$. As a third step, the parties securely compute $R^j \cdot ((\lambda_a \oplus \alpha) \cdot (\lambda_b \oplus \beta) \oplus \lambda_w)$ for all $j = 1, \ldots, N$ and $\alpha, \beta \in \{0,1\}$. The multiplications can be done by eight κ-bit $\mathrm{C}^{\oplus}$-OTs per pair of parties: For all parties $P_j \neq P_i$, P_j inputs R^j as correlation and $P_i \neq P_j$ inputs λ_{abw}, $\lambda_{\bar{a}bw}$, $\lambda_{a\bar{b}w}$, and $\lambda_{\bar{a}\bar{b}w}$ as choice bits. Let $\rho^i_{j,\alpha,\beta}$ denote the resulting share of P_i of the product, i.e., the output of the corresponding $\mathrm{C}^{\oplus}$-OT. Finally the garbled tables $\{\tilde{g}^j_{0,0}, \tilde{g}^j_{0,1}, \tilde{g}^j_{1,0}, \tilde{g}^j_{1,1}\}_{j=1}^N$ are computed as follows: For $j = 1, \ldots, N$, and $\alpha, \beta \in \{0,1\}$, party P_j broadcasts

$$F^2_{k^j_{a,\alpha},k^j_{b,\beta}}(g \circ j) \oplus k^j_{w,0} \oplus \rho^j_{j,\alpha,\beta}$$

and all other parties P_i broadcast

$$F^2_{k^i_{a,\alpha},k^i_{b,\beta}}(g \circ j) \oplus \rho^i_{j,\alpha,\beta}.$$

The XOR of these messages yields the table entry $\tilde{g}^j_{\alpha,\beta}$. Ben-Efraim et al. [11] noticed that one of the κ-bit $\mathrm{C}^{\oplus}$-OTs can be saved in the third step: Instead of computing and using $\rho^i_{j,1,1}$ as described above, the table entry $\tilde{g}^j_{1,1}$ is computed as XOR of the values

$$F^2_{k^j_{a,\alpha},k^j_{b,\beta}}(g \circ j) \oplus k^j_{w,0} \oplus R^j \oplus \rho^j_{j,0,0} \oplus \rho^j_{j,0,1} \oplus \rho^j_{j,1,0}$$

from P_j and

$$F^2_{k^i_{a,\alpha},k^i_{b,\beta}}(g \circ j) \oplus \rho^i_{j,0,0} \oplus \rho^i_{j,0,1} \oplus \rho^i_{j,1,0}$$

from each other party P_j.

New garbling optimization. For a circuit consisting of m AND gates, the above garbling protocol uses $3m \cdot N(N-1)$ $\mathrm{C}^{\oplus}$-OTs of κ-bit strings in total. We show a *novel* optimization and reduce this to $\kappa \cdot N(N-1)$ $\mathrm{C}^{\oplus}$-OTs of $3m$-bit strings: First, note that we can swap the inputs of the OTs and use κ $\mathrm{C}^{\oplus}$-OTs of 3-bit strings to compute the shares $\rho^i_{j,\alpha,\beta}$. Let $\hat{\lambda}^i_g \in \{0,1\}^3$ be the triple $(\lambda^i_{abw}, \lambda^i_{a\bar{b}w}, \lambda^i_{\bar{a}bw})$ for the gate g. Then we can use the bits of R^j as choice bits and $\hat{\lambda}^i_g$ as correlation in the $\mathrm{C}^{\oplus}$-OT. By using the concatenation of all $\hat{\lambda}^i_1, \ldots, \hat{\lambda}^i_m$ as correlation, we can use the same $\kappa N(N-1)$ $\mathrm{C}^{\oplus}$-OT for all gates. This optimization results in 20/16/13/11% less communication for $N = 2/3/4/5$ parties. Using the bandwidth-saving broadcast from §5.6, the improvement is 20% and independent of the number of parties. Note that this optimization is not depicted in our benchmarks.

During the online phase, each party P_i holds for a wire w a public value $\alpha_w = \lambda_w \oplus x$ (with permutation bit λ_w and real value x), keys $k^j_{w,\alpha}$ for $j = 1, \ldots, N$ (and $k^i_{w,1-\alpha} = k^i_{w,\alpha} \oplus R^i$), and an additive share λ^i_w of the permutation bit. We can write this in the form of a sharing as $\langle x \rangle^Y = (\lambda^1_x, \ldots, \lambda^N_x; (\alpha, k^1_{x,\alpha}, \ldots, k^N_{x,\alpha}))$ where the part after the semicolon denotes public information.

Given the setup as described above, we now describe the basic operations of this sharing during the online phase: For $\mathrm{Share}^Y_i(x)$, party P_i (holding λ) broadcasts $\alpha = x \oplus \lambda$, and each party P_j broadcasts k^j_α. For $\mathrm{Rec}^Y_i(\langle x \rangle^Y)$, party P_i (holding λ) computes $x \leftarrow \alpha \oplus \lambda$. Let $\langle x \rangle^Y = (\lambda^1_x, \ldots, \lambda^N_x; (\alpha, k^1_{x,\alpha}, \ldots, k^N_{x,\alpha}))$, and $\langle y \rangle^Y = (\lambda^1_y, \ldots, \lambda^N_y; (\beta, k^1_{y,\beta}, \ldots, k^N_{y,\beta}))$ be two shared values. The XOR of these $\langle z \rangle^Y \leftarrow \langle x \rangle^Y \oplus \langle z \rangle^Y$ can be computed using a free-XOR technique: With $\gamma \leftarrow \alpha \oplus \beta$, and $k^j_{z,\gamma} \leftarrow k^j_{x,\alpha} \oplus k^j_{y,\beta}$ for all $j \in [N]$, we get $\langle z \rangle^Y = (\lambda^1_z, \ldots, \lambda^N_z; (\gamma, k^1_{z,\gamma}, \ldots, k^N_{z,\gamma}))$. For an AND gate g, the corresponding garbled tables must be decrypted: For $j \in [N]$, the next key is computed as $k^j_{z,\gamma} \leftarrow \tilde{g}^j_{\alpha,\beta} \oplus \bigoplus_{i=1}^N F_{k^i_{x,\alpha},k^i_{y,\beta}}(g \circ j)$. Then, each party P_i can deduce γ by checking whether $k^i_{z,\gamma} = k^i_{z,0}$ or $k^i_{z,\gamma} = k^i_{z,1}$ holds. Basic operations can be done using size-optimized Boolean circuits [50, 68].

7 MPC PROTOCOL CONVERSIONS

We present secure and efficient conversions between the three protocols (cf. §6) to enable passively secure hybrid MPC, which allows to benefit from the respective advantages of the underlying protocols. We summarize the conversion costs in Tab. 6 in App. A.

7.1 Boolean to Arithmetic Sharing – *B2A*

For converting a Boolean sharing $\langle \mathbf{x} \rangle^B$ of ℓ bits into an arithmetic sharing $\langle x \rangle^A$ over $\mathbb{Z}_{2^\ell}$ such that x equals $\mathbf{x}$ when interpreted as an element of $\mathbb{Z}_{2^\ell}$, we present two variants

Straightforward: Additive Masking. As described in prior work [3, 29, 40] in different settings, the conversion can be computed as follows: A random mask is added to the input value in the Boolean sharing. The result is reconstructed and shared again in the arithmetic sharing where the mask is subtracted again.

The mask can be generated in the online phase by letting each party share a random value. However, since the mask is input-independent, it could also be generated by running a subprotocol during the setup phase. Here, we assume that we have a pair $(\langle r \rangle^A, \langle \mathbf{r} \rangle^B)$ of sharings of the same value $r \in_R \mathbb{Z}_{2^\ell}$. To convert the sharings, the parties compute $\langle \mathbf{t} \rangle^B \leftarrow \langle \mathbf{x} \rangle^B - \langle \mathbf{r} \rangle^B$, $t \leftarrow \mathrm{Rec}^B(\langle \mathbf{t} \rangle^B)$, $\langle t \rangle^A \leftarrow \mathrm{Share}^A(t)$, and $\langle x \rangle^A \leftarrow \langle t \rangle^A + \langle r \rangle^A$.

This requires at least $\Omega(\log \ell)+1$ rounds of communication in the online phase for computing the subtraction circuit in GMW [18, 66] and the subsequent reconstruction. Moreover, one pair of sharings $(\langle r \rangle^A, \langle \mathbf{r} \rangle^B)$ for $r \in_R \mathbb{Z}_{2^\ell}$, generated in the setup phase, is required.

Optimized: Using Shared Bits. In our implementation, we adapt the approach from [25] for $\mathrm{SPD}\mathbb{Z}_{2^k}$ to our setting and use so called

shared bits. A shared bit is a pair of sharings $(\langle r\rangle^A, \langle r\rangle^B)$ of a random bit $r \in_R \{0,1\}$.

Let $\langle \mathbf{x}\rangle^B = (\langle x_0\rangle^B, \ldots, \langle x_{\ell-1}\rangle^B)$ with least significant bit $\langle x_0\rangle^B$. Given shared bits $(\langle r_i\rangle^A, \langle r_i\rangle^B)$ for $i = 0, \ldots, \ell-1$, we can convert $\langle \mathbf{x}\rangle^B$ into an arithmetic sharing as follows: For each bit $i = 0, \ldots, \ell-1$, the parties compute in parallel $\langle t_i\rangle^B \leftarrow \langle x_i\rangle^B \oplus \langle r_i\rangle^B$, $t_i \leftarrow \mathsf{Rec}^B(\langle t_i\rangle^B)$, and $\langle x_i\rangle^A \leftarrow t_i + \langle r_i\rangle^A - 2t_i\langle r_i\rangle^A$. Thereafter, the output sharing is computed as $\langle x\rangle^A \leftarrow \sum_{i=0}^{\ell-1} 2^i \cdot \langle x_i\rangle^A$.

This costs only one round of communication for the reconstruction of the t_i during which $N(N-1)\ell$ bits are transmitted, and ℓ shared bits, which are generated during the setup phase (cf. §5.4).

7.2 Boolean to Yao Sharing – *B2Y*

The straight-forward way to do the *B2Y* conversion of a shared value $\langle x\rangle^B$ would be that each party P_i reshares its Boolean share $\langle x\rangle_i^B$ in Yao sharing as $\langle x_i\rangle^Y \leftarrow \mathsf{Share}_i^Y(\langle x\rangle_i^B)$ and the parties compute $\langle x\rangle^Y \leftarrow \bigoplus_{j=1}^N \langle x_i\rangle^Y$. The sharing requires two rounds of communication and has a total communication cost of $N(N-1)(N\kappa+1)$ bits, which is in $O(N^3\kappa)$.

The properties of the BMR sharing allow the following natural optimization for the *B2Y* conversion (also implemented by [3]): Let w be the BMR wire that is supposed to obtain the value x. Note that party P_i holds in addition to its Boolean share $\langle x\rangle_i^B$ also a share λ_w^i of the random permutation bit $\lambda_w = \bigoplus_{j=1}^N \lambda_w^j$, and keys $k_{w,0}^i, k_{w,1}^i = k_{w,0}^i \oplus R^i$, which are generated during the BMR setup phase (cf. §6.3). For the conversion, each party P_i first broadcasts $\alpha_i \leftarrow \langle x\rangle_i^B \oplus \lambda_w^i$. Then, every party P_i computes $\alpha \leftarrow \bigoplus_{j=1}^N \alpha_j$ and broadcasts $k_{w,\alpha}^i$. Then $\langle x\rangle^Y := (\lambda_w^1, \ldots, \lambda_w^N; (\alpha, k_{w,\alpha}^1, \ldots, k_{w,\alpha}^N))$ is a valid Yao sharing of x since $\alpha = \bigoplus_{j=1}^N \alpha_j = \bigoplus_{j=1}^N \langle x\rangle_i^B \oplus \lambda_w^j = x \oplus \lambda_w$. This optimized conversion requires also two rounds but only $N(N-1)(\kappa+1)$ bits of communication, which is in $O(N^2\kappa)$. This is an improvement by a factor of $\frac{(N\kappa+1)}{(\kappa+1)} \approx N$ over the straight-forward solution.

7.3 Yao to Boolean Sharing – *Y2B*

Let $\langle x\rangle^Y = (\lambda_x^1, \ldots, \lambda_x^N; (\alpha, k_{x,\alpha}^1, \ldots, k_{x,\alpha}^N))$ be the Yao sharing of a value $x \in \{0,1\}$. As described in §6.3, the public value α is the real value x masked with the random permutation bit $\lambda_x = \bigoplus_{j=1}^N \lambda_x^j$, i.e., $\alpha = x \oplus \lambda_x$. Hence, the shared permutation bit is already a Boolean sharing $\langle \alpha \oplus x\rangle^B = (\lambda_x^1, \ldots, \lambda_x^N)$, and the parties compute $\langle x\rangle^B \leftarrow \langle \alpha \oplus x\rangle^B \oplus \alpha$ (cf. §6.1 and §6.2), i.e., party P_1 computes $\langle x\rangle_1^B \leftarrow \lambda_1 \oplus \alpha$ and all other parties $P_2, \ldots, P_N$ set $\langle x\rangle_j^B := \lambda_j$. Then, we have obtained a Boolean sharing $\langle x\rangle^B$ of x since $\bigoplus_{j=1}^N \langle x\rangle_j^B = \alpha \oplus \lambda = x$. *Y2B* can be computed locally and hence is for free.

7.4 Arithmetic to Yao Sharing – *A2Y*

Given an arithmetic sharing $\langle x\rangle^A = (\langle x\rangle_1^A, \ldots, \langle x\rangle_N^A)$ over $\mathbb{Z}_{2^\ell}$ we want to obtain a Yao sharing $\langle \mathbf{x}\rangle^Y$ of ℓ bits such that $\mathbf{x}$ equals x when interpreted as element of $\mathbb{Z}_{2^\ell}$. To achieve this, every party P_i first shares its additive share of $\langle x\rangle^A$ in the Yao sharing: $\langle \mathbf{x}_i\rangle^Y \leftarrow \mathsf{Share}_i^Y(\langle x\rangle_i^A)$. Then, they compute $\langle \mathbf{x}\rangle^Y \leftarrow \sum_{j=1}^N \langle \mathbf{x}_j\rangle^Y$ using a Boolean circuit for addition. The conversion requires two rounds of communication for the sharing (cf. §6.3), and the evaluation of $N-1$ addition circuits in BMR with $O(\ell N)$ AND gates in total.

7.5 Arithmetic to Boolean Sharing – *A2B*

There are two options for converting an arithmetic sharing $\langle x\rangle^A$ over $\mathbb{Z}_{2^\ell}$ into a Boolean sharing $\langle \mathbf{x}\rangle^B$ of ℓ bits, such that $\mathbf{x}$ equals x when interpreted as element of $\mathbb{Z}_{2^\ell}$.

We can do the analogue of *A2Y* (cf. §7.4) in Boolean sharing. However, this requires $O(\log N \cdot \log \ell)$ rounds of communication to compute the $N-1$ additions using depth-optimized addition circuits [18, 66].

In order to avoid the additional communication rounds, we first convert $\langle x\rangle^A$ to $\langle \mathbf{x}\rangle^Y$ (cf. §7.4), and then $\langle \mathbf{x}\rangle^Y$ to $\langle \mathbf{x}\rangle^B$ for free (cf. §7.3). Hence, *A2B* has the same costs as *A2Y*.

7.6 Yao to Arithmetic Sharing – *Y2A*

Straightforward: Via *Y2B* and *B2A*. We implemented the conversion of a Yao sharing $\langle \mathbf{x}\rangle^Y$ of ℓ bits into an arithmetic sharing $\langle x\rangle^A$ over $\mathbb{Z}_{2^\ell}$ by first converting $\langle \mathbf{x}\rangle^Y$ into a Boolean sharing $\langle \mathbf{x}\rangle^B$ for free with *Y2B* (cf. §7.3), and then applying *B2A* (cf. §7.1) to obtain $\langle x\rangle^A$. Hence, the costs of *Y2A* are the same as for *B2A*: one round of communication during the online phase.

Optimized: Without Online Communication. Furthermore, we present a *novel* conversion protocol that computes *Y2A without any online communication*: This conversion requires a precomputed pair $(\langle r\rangle^A, \langle \mathbf{r}\rangle^Y)$ consisting of an arithmetic sharing $\langle r\rangle^A$ and a Yao sharing $\langle \mathbf{r}\rangle^Y$ of the same randomly chosen value $r \in_R \mathbb{Z}_{2^\ell}$. Since r is sampled independently of the overall protocol's inputs, it can be generated beforehand in the setup phase (see below). We first describe the *online* phase of the conversion. Given a sharing $\langle \mathbf{x}\rangle^Y$ of a value x and a pair as described above, we compute an arithmetic sharing $\langle x\rangle^A$ of x as follows: First, the input value x is masked with r in Yao sharing $\langle \mathbf{t}\rangle^Y \leftarrow \langle \mathbf{x}\rangle^Y - \langle \mathbf{r}\rangle^Y$. Then, the masked value is reconstructed $t \leftarrow \mathsf{Rec}^Y(\langle \mathbf{t}\rangle^Y)$, and shared arithmetically $\langle t\rangle^A \leftarrow \mathsf{Share}^A(t)$. Finally, the mask is removed in arithmetic sharing $\langle x\rangle^A \leftarrow \langle t\rangle^A + \langle r\rangle^A$. Note that each of these steps can be computed without any communication in the online phase: The subtraction circuit in Yao sharing and $\mathsf{Rec}^Y(\cdot)$ can be computed locally due to the properties of the BMR protocol (cf. §6.3). Also, $\mathsf{Share}^A(\cdot)$ and addition in arithmetic sharing do not require any online communication (cf. §6.1). The input-independent pair $(\langle r\rangle^A, \langle \mathbf{r}\rangle^Y)$ can be generated in the *setup* phase as follows: Every party P_i samples $\langle r\rangle_i^A \in_R \mathbb{Z}_{2^\ell}$ resulting in a sharing $\langle r\rangle^A = \sum_{j=0}^N \langle r\rangle_j^A$. Then $\langle \mathbf{r}\rangle^Y$ is obtained by applying a *A2Y* conversion to $\langle r\rangle^A$ (cf. §7.4). To reduce communication costs, we can also compute the summation in Boolean sharing (cf. §7.5) and convert to BMR afterwards (cf. §7.2). This improves the setup phase by a factor of $O(N)$ in communication at the expense of more rounds in the setup phase.

8 PERFORMANCE EVALUATION

Along with the code base of MOTION, we provide the code and benchmarks for multiple applications that use MOTION as a C++ library. In this section, we evaluate the performance of MOTION and of these applications. We compare the applications' performance

with other MPC implementations that also offer full-threshold security, i.e., protocols that *increase* their level of security by adding more parties. We do not compare with frameworks such as [14, 20, 58, 60, 62] that involve multiple parties for performance improvements, but offer only security against a single corruption. We run benchmarks in two different environments: our own servers connected via a local network and several AWS servers.

Our servers: Five servers each equipped with an Intel Core i9-7960X processor and 128 GB RAM, connected via a 10 Gbps network. For this benchmark environment, we define two network settings to analyze how our framework behaves in different scenarios.

- LAN: The network is used as is with 10 Gbps bandwidth and 0.25 ms RTT. This setting represents parties in a fast LAN or an outsourcing scenario (cf. §3.3) with servers located in a network with low latency and high bandwidth, e.g., computing parties connected at an Internet Exchange Point (IXP).
- WAN: `tc`[13] is used to limit the network bandwidth to 1 Gbps and simulate an average RTT of 100 ms, simulating parties connected over the Internet. The scenarios covered by this setting are, for example, ad-hoc MPC over the Internet run by normal users and outsourcing computation to servers located in distinct locations, e.g., each server is owned by a different company in a different country.

AWS servers: To perform experiments with a larger number of parties, we use 10, 15, or 20 `r5.8xlarge` instances on AWS EC2, located in the same availability zone[14]. Each instance has 32 vCPUs using Intel Xeon Platinum 8175 or worse processors with 256 GiB memory and a 10 Gbps network connection[15]. We measured a bandwidth between 4.8 Gbps and 9.6 Gbps and an RTT between 0.043 ms and 0.079 ms among the instances. This setting represents two use cases: (1) a direct use of MPC between many parties, e.g., for privacy-preserving auctions, and (2) outsourcing to many servers of which all but one can be passively corrupted for a very high level of privacy, e.g., for privacy-preserving computation on genomic data.

We average most of our benchmarks over 100 iterations. On AWS, we run 10 to 25 iterations to reduce the required time and costs. MOTION includes the functionality of automatically collecting extensive run-time and communication statistics. These numbers can be viewed for individual executions, and separate parts of protocols and primitives (e.g., OTs, MTs, etc.), as well as aggregated numbers for an entire batch of executions, including average numbers and their standard deviation. While MOTION can easily support other communication protocols (cf. §4.1.1), we used TCP in our benchmarks. We have run several benchmarks with the TCP traffic tunneled through a TLS channel using `stunnel`[16] and did not observe any noticeable performance overhead.

8.1 Microbenchmarks

We provide extensive microbenchmarks and communication requirements for primitive MPC operations and conversions, as well as microbenchmarks for integer operations in MOTION that can serve as guidelines for protocol design and cost estimation. Due to space limitations, we provide these results in App. A.

Table 2: Run-times in nanoseconds for one OT, amortized over 10 million parallel evaluations, averaged over 100 runs.

	G-OT		R-OT	$C^{\oplus}$-OT		C^{+}-OT			
Bit size	1	128	128	1	128	8	16	32	64
libOTe [65] LAN	–	120	–	–	130	–	–	–	–
MOTION LAN (this work)	151	196	77	131	147	119	121	126	134
libOTe [65] WAN	–	820	–	–	874	–	–	–	–
MOTION WAN (this work)	1 069	1 221	957	932	973	955	960	972	980

8.1.1 Performance of OT Extension (Tab. 2). In Tab. 2, we compare our OT extension implementation with the libOTe library by Peter Rindal [65]. The major part of libOTe is written in assembly and is, hence, very efficient. libOTe provides interfaces for single OT batches, which are explicitly associated with a communication channel, and operates directly on network sockets without message serialization and thus requires to manually synchronize all the uses of different OTs. Taking the above into account, libOTe is easy to use and efficient in small MPC applications but, unfortunately, is often inconvenient for constructing complex MPC protocols. Our `OTProvider` class implemented in MOTION provides an abstract non-blocking API to request and use OTs without any knowledge about the underlying communication channel or other OTs.

We compare the efficiency of our OT extension implementation on 128-bit $C^{\oplus}$-OT, which is one of the core components of the BMR protocol (cf. §6.3) and 128-bit G-OT, which can be used for implementing other MPC protocols. For a total of 10 million parallel 128-bit $C^{\oplus}$-OT evaluations averaged over 100 runs, libOTe is only 10% faster than our `OTProvider`. In the same setting, libOTe's 128-bit G-OT implementation is 1.6× faster than our 128-bit G-OT implementation. In the WAN setting, the performance difference is slightly smaller: libOTe outperforms our `OTProvider` for 128-bit C-OT by factor 1.1 and G-OT by factor 1.5. Taking into account the additional overhead for the communication serialization and the much higher level of abstraction in MOTION, the performance difference between the implementations is very small. To further improve the efficiency of our `OTProvider`, it is possible to replace parts of our code with assembly code as was done in libOTe. However, we aim to avoid this by design to make our code portable to different platforms like ARM.

8.1.2 Boolean Circuits: BMR vs. GMW. Since the state-of-the-art BMR protocol [11] undoubtedly incurs higher overhead than GMW, the authors of [11] created artificial circuits to show that circuits with very high depth can be evaluated faster in BMR than in GMW in high-latency networks. Here, we give a *real-world* example where BMR is more efficient than GMW even in the LAN setting with *low* network latency.

In the experiments on our servers, evaluation of integer division circuits generated using HyCC [17] was always faster in BMR than in GMW. In the LAN setting, the difference was 1.1×–1.4×, whereas in the WAN setting the factors were between 3× and 5.3×. For the 3-party 64-bit integer division, the run-time difference between BMR and GMW in the WAN setting was 296 ms, which

[13] http://man7.org/linux/man-pages/man8/tc.8.html
[14] Exact location omitted for anonymous submission.
[15] https://aws.amazon.com/ec2/instance-types/
[16] https://www.stunnel.org/

Table 3: Total (online+setup) run-times in seconds for biometric matching, comparing several implementations and protocols over various domains for N parties, bitlength ℓ and multiple database sizes. We benchmarked MOTION and ABY [29] with circuits generated with the HyCC compiler [17]. In MOTION, the runtimes with SIMD are amortized over 192 / 32 parallel circuits for DB sizes of 1 024 / 4 096. Best runtimes are in bold.

							LAN		WAN	
							DB Size		DB Size	
Implementation	Protocol	Domain	Security	N	Thresh.	ℓ	1 024	4 096	1 024	4 096
ABY [29]	A+B	$\mathbb{Z}_{2^\ell}$	passive	2	1	32	0.26	0.89	2.6	4.1
ABY [29]	A+Y	$\mathbb{Z}_{2^\ell}$	passive	2	1	32	**0.24**	**0.76**	**2.5**	**3.6**
MP-SPDZ [48]	MASCOT [49]	$\mathbb{F}_p$	active	3	2	32	45.78	174.90	1 150.4	4 596.0
MP-SPDZ [48]	MASCOT [49]	$\mathbb{F}_p$	passive	3	2	32	9.56	36.78	935.0	3 746.0
MP-SPDZ [48]	SPD$\mathbb{Z}_{2^k}$ [23]	$\mathbb{Z}_{2^\ell}$	active	3	2	64	57.25	231.53	1 643.8	6 580.2
MP-SPDZ [48]	SPD$\mathbb{Z}_{2^k}$ [23]	$\mathbb{Z}_{2^\ell}$	passive	3	2	64	3.89	13.80	1 126.1	4 500.5
MP-SPDZ [48]	FKOS15 [31]	binary	active	3	2	32	104.76	413.25	3 456.6	13 772.6
MP-SPDZ [48]	OT-based	binary	passive	3	2	32	4.17	14.80	1 346.4	5 289.3
SCALE-MAMBA [2]	Full-Threshold	$\mathbb{F}_p$	active	3	2	32	128.95	253.19	858.9	2 033.0
MOTION (**this work**)	A+B	$\mathbb{Z}_{2^\ell}$	passive	3	2	32	5.74	19.99	15.4	41.3
MOTION (**this work**) w/ SIMD	A+B	$\mathbb{Z}_{2^\ell}$	passive	3	2	32	**0.22**	**1.33**	**1.2**	**5.2**
MOTION (**this work**)	A+Y	$\mathbb{Z}_{2^\ell}$	passive	3	2	32	5.71	21.78	10.2	29.2
MOTION (**this work**) w/ SIMD	A+Y	$\mathbb{Z}_{2^\ell}$	passive	3	2	32	0.26	1.55	1.8	7.5

is equivalent to the run-time of 127 secure 64-bit additions or 30 secure 64-bit multiplications in 3-party GMW, and thus is significant. This substantial difference is due to the very high depth of the division circuit, which ranges from depth 65 for 8-bit division to depth 2 218 for 64-bit division. However, on the AWS servers with high bandwidth, low latency, and 10 to 20 parties, BMR performs worse and scales poorer than GMW due to its substantially higher run-time and communication overhead. More detailed results are provided in Tab. 8 in App. A.

8.1.3 Comparison with N-Party GMW [22]. Compared to the passively secure N-party Boolean GMW implementation by Choi et al. [22], which requires amortized 4.61 µs to evaluate one AND gate by three parties, MOTION requires only 0.55 µs (cf. Tab. 7 in App. A), which is 8.4× faster. Our better run-times can be explained by our more efficient OT extension implementation, and the use of MPC-level SIMD instructions (cf. §4.1.7). Both implementations were benchmarked on the same hardware.

8.2 Applications

In this section, we benchmark the runtimes for secure evaluation of real-world applications in MOTION and compare them with other full-threshold MPC frameworks. We run all these implementations on identical hardware using the same network conditions.

8.2.1 Biometric Matching (Tab. 3). Here, we analyze the overhead of moving from passively secure full-threshold MPC to actively-secure full-threshold MPC by comparing our framework with the SCALE-MAMBA framework [2] and with multiple protocols implemented in MP-SPDZ [48] and also compare with the passively secure 2-party ABY framework [29]. As function, we use biometric matching that computes the Euclidean distance between a 4-dimensional sample and a database of biometric samples and then determines the minimum distance. A code example for the 2-dimensional case is provided in Listing 1 on page 5. We give the run-times in Tab. 3 for 2^{10} and 2^{12} database entries. Apart from benchmarks for the HyCC biomatch circuit [17] that is evaluated in a non-SIMD fashion, we provide a native MOTION implementation for the biometric matching with equivalent functionality but utilizing SIMD instructions evaluating 200 parallel circuits for 1 024 elements and 40 parallel circuits for 4 096 elements. The latter results in 16×–34× amortized speedup in the LAN setting and in a 42×–221× amortized speedup in the WAN setting (cf. Tab. 3).

Comparison with SCALE-MAMBA [2] & MP-SPDZ [48]. For SCALE-MAMBA, we set up a 3-party scenario with full-threshold security. For the passively secure versions of the MASCOT [49] and SPD$\mathbb{Z}_{2^k}$ [23] protocol, we compiled mixed circuits which are more efficient, whereas the actively secure versions of these protocols turned out to be more efficient when running plain, non-hybrid circuits in the respective sharing. As a default we used values with a bitlength ℓ=32 bits, but had to run some measurements with ℓ=64 bit values, due to limitations of the respective implementation.

Comparing MOTION's runtimes from Tab. 3 with those of the passively secure protocols of MP-SPDZ, we can see that the HyCC biometric matching circuit in MOTION is from 1.6× slower to 1.7× faster in the LAN setting, and 61×–255× faster in the WAN setting. With enabled SIMD support, MOTION outperforms MP-SPDZ and SCALE-MAMBA in all settings: it is at least 8.9× / 271× faster in the LAN / WAN setting than the fastest protocol implemented in MP-SPDZ (passive) or SCALE-MAMBA (active).

Comparison with ABY [29]. The passively secure two-party ABY framework outperforms most other implementations. As shown in Tab. 3, biometric matching in ABY is from slightly slower to 1.8× faster than in MOTION in the LAN setting. This is mainly due to the higher cost of the *A2B* conversion in MOTION, which requires multiple addition circuits instead of one, and because BMR incurs higher communication and computation costs than two-party garbled circuits. However, in the WAN setting where the communication plays a greater role than in the LAN setting, we measured from 2.0× faster to 1.4× slower runtimes in MOTION

Table 4: Run-times in milliseconds for the evaluation of the Bristol Fashion circuits [1] for AES-128 with key scheduling and SHA-256 in GMW (B) and BMR (Y) executed with MOTION and Choi et al.'s GMW [22]. The run-times are amortized over 512 / 256 executions of AES-128 / SHA-256.

			LAN			WAN		
	Implementation		N=2	N=3	N=5	N=2	N=3	N=5
AES	Choi et al. [22]	B	27.3	42.2	84.8	28.9	44.4	90.0
	ABY [29]	B	0.2	—	—	8.5	—	—
	online	B	<0.1	—	—	6.6	—	—
	ABY [29]	Y	0.2	—	—	1.9	—	—
	online	Y	0.1	—	—	0.1	—	—
	MOTION (**this work**)	B	1.9	2.5	3.8	13.8	14.9	18.9
	online	B	0.4	0.5	0.8	7.3	7.7	8.1
	MOTION (**this work**)	Y	4.7	8.0	17.1	61.1	87.8	141.7
	online	Y	0.2	0.2	0.4	0.5	0.7	0.9
SHA	Choi et al. [22]	B	80.7	128.5	254.5	104.1	150.7	287.6
	ABY [29]	B	1.5	—	—	339.8	—	—
	online	B	0.7	—	—	334.1	—	—
	ABY [29]	Y	1.5	—	—	8.1	—	—
	online	Y	0.5	—	—	0.6	—	—
	MOTION (**this work**)	B	8.3	10.8	16.0	500.2	572.4	614.1
	online	B	2.7	3.2	4.5	479.6	547.9	538.4
	MOTION (**this work**)	Y	19.0	29.6	61.8	201.9	279.3	492.6
	online	Y	1.3	1.6	1.8	1.4	2.1	2.7

than in ABY, which is due to the more efficient communication using SIMD instructions in MOTION.

8.2.2 AES-128 and SHA-256 (Tab. 4). Here, we provide a comparison of the overhead needed to move from passively secure Secure Two-Party Computation (2PC) to passively secure full-threshold Secure N-Party Computation by comparing MOTION with the ABY framework [29]. Amortized run-times for securely evaluating 1 000 parallel invocations of AES-128 and SHA-256 in MOTION are given in Tab. 4. An important observation from this table is that for both AES and SHA the run-time of GMW (B) in the WAN setting is dominated by the online time, which cannot be precomputed, whereas the online time of BMR (Y) is only a small fraction of the total run-time. BMR has substantially higher run-times in the LAN setting, but has a much faster online phase in the WAN setting.

Comparison with N-Party GMW [22] A comparison of our GMW (B) implementation with the passively secure GMW implementation of Choi et al. [22] is given in Tab. 4. In the LAN setting, MOTION is 14×–22× faster for AES and 10×–16× for SHA. In the WAN setting, MOTION is 2×–5× faster for AES and 1.7×–2× slower for SHA. Surprisingly, their implementation is almost as fast in the WAN as in the LAN setting, and is faster in the WAN setting even for N=5 parties than ABY [29] that implements highly efficient N=2-party protocols, which we cannot explain. Note that we did not try to verify the correctness of their implementation. Also, MOTION is able to evaluate 10× as many SHA circuits in the *same* time, while Choi et al.'s implementation only supports circuits of very limited size. Moreover, their circuits have to be provided in a custom file format, whereas MOTION has builtin support for multiple circuit formats such as the commonly used Bristol (Fashion) format. Also, Choi et al.'s implementation does not distinguish between setup and online phase.

Comparison with ABY [29]. In ABY, we securely computed 100 000 AES evaluations in the same LAN setting, using two-party

Table 5: Total run-times in seconds for CryptoNets [32] with HyCC hybrid circuits [17] for N parties and full threshold.

		LAN				WAN			
Implementation		N=2	N=3	N=4	N=5	N=2	N=3	N=4	N=5
ABY [29]	A+B	0.5	—	—	—	3.2	—	—	—
	A+Y	0.5	—	—	—	3.4	—	—	—
MOTION (**this work**)	A+B	3.5	4.2	4.9	5.7	6.7	8.0	9.8	13.2
	A+Y	3.6	4.2	4.9	5.8	6.7	8.6	12.6	13.1

Boolean GMW (B) in 20.0 s. In contrast, MOTION requires 183.4 s using N=3-party GMW which is 9.2× slower, and 303.8 s using N=5-party GMW, which is 15.2× slower. This difference results from the more efficient N=2-party protocols implemented in ABY and the substantially higher level of abstraction in MOTION. Although the workload of each party increases with the total number of parties, the difference between three and five-party GMW in MOTION is only minor (1.65×) due to the substantially better load balancing with more parties. As expected, the high-depth SHA-256 circuit can be evaluated faster in Y than in B. However, for N=2 parties Yao's GCs are 25× faster than BMR in the WAN setting, which indicates the significant gap between the efficiency of both protocols. The MOTION runtimes here are extrapolated from the runtimes in Tab. 4.

Comparison with BMR [11]. The original passively secure OT-based BMR implementation by Ben-Efraim et al. [11] requires approximately 1 s (698±930 ms setup and 138±88 ms online time) for a single AES-128 evaluation by N=3 parties with 75 ms average network latency and 10 Gbps network bandwidth. MOTION takes 1.3 s in the same network setting (as their code is not publicly available, we use slightly different machines), which is similar to the run-times in [11]. By evaluating 1 000 AES circuits in parallel, we achieve an amortized run-time of 61 ms, which is at least 16× faster than [11].

8.2.3 Privacy-Preserving Machine Learning (Tab. 5). MOTION can be used for privacy-preserving machine learning. We give benchmarks for privately evaluating a convolutional neural network for handwriting recognition in Tab. 5. For our benchmarks, we use the hybrid circuits generated by HyCC [17] for CryptoNets [32] with ReLU as activation function. In the case of N=2 parties, MOTION is slower than ABY [29]: 7× in the LAN and 2× in the WAN setting, because our protocols are generic for N parties, whereas ABY has optimized protocols for exactly N=2 parties only. When increasing the number of parties and hence obtaining better security due to the full threshold protocols in MOTION, the performance of MOTION decreases only slightly, e.g., 1.6× for N=5 vs. N=2 parties in LAN and 2.0× in WAN.

Acknowledgments. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 850990 PSOTI). It was co-funded by the Deutsche Forschungsgemeinschaft (DFG) — SFB 1119 CROSSING/236615297 and GRK 2050 Privacy & Trust/251-805230, and by the German Federal Ministry of Education and Research and the Hessen State Ministry for Higher Education, Research and the Arts within ATHENE.

REFERENCES

[1] V. A. Abril, P. Maene, N. Mertens, D. Sijacic, N. Smart. "'Bristol Fashion' MPC Circuits". *https://homes.esat.kuleuven.be/~nsmart/MPC/*. 2019.

[2] A. Aly, M. Keller, D. Rotaru, P. Scholl, N. P. Smart, T. Wood. "SCALE-MAMBA". *https://homes.esat.kuleuven.be/~nsmart/SCALE/*. 2018.

[3] A. Aly, E. Orsini, D. Rotaru, N. P. Smart, T. Wood. "Zaphod: Efficiently Combing LSSS and Garbled Circuits in SCALE". In: *Workshop on Encrypted Computing & Applied Homomorphic Cryptography (WAHC'19)*. ACM, 2019, pp. 33–44.

[4] G. Asharov, Y. Lindell, T. Schneider, M. Zohner. "More Efficient Oblivious Transfer and Extensions for Faster Secure Computation". In: *CCS'13*. ACM. 2013, pp. 535–548.

[5] G. Asharov, Y. Lindell, T. Schneider, M. Zohner. "More Efficient Oblivious Transfer Extensions". In: *Journal of Cryptology (JoC)* 3 (2017), pp. 805–858.

[6] D. Beaver. "Efficient Multiparty Protocols using Circuit Randomization". In: *CRYPTO'91*. Springer. 1991, pp. 420–432.

[7] D. Beaver. "Precomputing Oblivious Transfer". In: *CRYPTO'95*. Springer. 1995, pp. 97–109.

[8] D. Beaver, S. Micali, P. Rogaway. "The Round Complexity of Secure Protocols". In: *STOC'90*. ACM, 1990, pp. 503–513.

[9] M. Bellare, V. T. Hoang, S. Keelveedhi, P. Rogaway. "Efficient Garbling from a Fixed-Key Blockcipher". In: *S&P'13*. IEEE, 2013, pp. 478–492.

[10] A. Ben-David, N. Nisan, B. Pinkas. "FairplayMP: A System for Secure Multi-Party Computation". In: *CCS'08*. ACM. 2008, pp. 257–266.

[11] A. Ben-Efraim, Y. Lindell, E. Omri. "Optimizing Semi-Honest Secure Multiparty Computation for the Internet". In: *CCS'16*. ACM, 2016, pp. 578–590.

[12] M. Ben-Or, S. Goldwasser, A. Wigderson. "Completeness Theorems for Non-Cryptographic Fault-Tolerant Distributed Computation". In: *STOC'88*. ACM. 1988, pp. 1–10.

[13] M. Blanton, P. Gasti. "Secure and Efficient Protocols for Iris and Fingerprint Identification". In: *ESORICS'11*. Springer, 2011, pp. 190–209.

[14] D. Bogdanov, S. Laur, J. Willemson. "Sharemind: A Framework for Fast Privacy-Preserving Computations". In: *ESORICS'08*. Springer, 2008, pp. 192–206.

[15] P. Bogetoft, D. L. Christensen, I. Damgård, M. Geisler, T. P. Jakobsen, M. Krøigaard, J. D. Nielsen, J. B. Nielsen, K. Nielsen, J. Pagter, M. I. Schwartzbach, T. Toft. "Secure Multiparty Computation Goes Live". In: *FC'09*. Springer, 2009, pp. 325–343.

[16] M. Brandt, C. Orlandi, K. Shishak, H. Shulman. "Optimizing Transport Layer for Secure Computation". In: *IACR Cryptology ePrint Archive, Report 2019/836* (2019). https://ia.cr/2019/836.

[17] N. Büscher, D. Demmler, S. Katzenbeisser, D. Kretzmer, T. Schneider. "HyCC: Compilation of Hybrid Protocols for Practical Secure Computation". In: *CCS'18*. ACM, 2018, pp. 847–861.

[18] N. Büscher, A. Holzer, A. Weber, S. Katzenbeisser. "Compiling Low Depth Circuits for Practical Secure Computation". In: *ESORICS'16*. Springer, 2016, pp. 80–98.

[19] N. Chandran, D. Gupta, A. Rastogi, R. Sharma, S. Tripathi. "EzPC: Programmable and Efficient Secure Two-Party Computation for Machine Learning". In: *EuroS&P'19*. IEEE, 2019, pp. 496–511.

[20] H. Chaudhari, A. Choudhury, A. Patra, A. Suresh. "ASTRA: High Throughput 3PC over Rings with Application to Secure Prediction". In: *CCSW'19*. ACM, 2019, pp. 81–92.

[21] J. I. Choi, D. Tian, G. Hernandez, C. Patton, B. Mood, T. Shrimpton, K. R. B. Butler, P. Traynor. "A Hybrid Approach to Secure Function Evaluation Using SGX". In: *ASIACCS'19*. ACM, 2019, pp. 100–113.

[22] S. G. Choi, K.-W. Hwang, J. Katz, T. Malkin, D. Rubenstein. "Secure Multi-Party Computation of Boolean Circuits with Applications to Privacy in On-Line Marketplaces". In: *CT-RSA'12*. Springer. 2012, pp. 416–432.

[23] R. Cramer, I. Damgård, D. Escudero, P. Scholl, C. Xing. "SPD$\mathbb{Z}_{2^k}$: Efficient MPC mod 2^k for Dishonest Majority". In: *CRYPTO'18*. Springer, 2018, pp. 769–798.

[24] I. Damgård, V. Pastro, N. P. Smart, S. Zakarias. "Multiparty Computation from Somewhat Homomorphic Encryption". In: *CRYPTO'12*. Springer, 2012, pp. 643–662.

[25] I. Damgård, D. Escudero, T. Frederiksen, M. Keller, P. Scholl, N. Volgushev. "New Primitives for Actively-Secure MPC over Rings with Applications to Private Machine Learning". In: *S&P'19*. IEEE, 2019, pp. 1102–1120.

[26] I. Damgård, M. Geisler, M. Krøigaard, J. B. Nielsen. "Asynchronous Multiparty Computation: Theory and Implementation". In: *CRYPTO'09*. Code: http://viff.dk. Springer. 2009, pp. 160–179.

[27] I. Damgård, M. Keller, E. Larraia, V. Pastro, P. Scholl, N. P. Smart. "Practical Covertly Secure MPC for Dishonest Majority - Or: Breaking the SPDZ Limits". In: *ESORICS'13*. Springer, 2013, pp. 1–18.

[28] D. Demmler, G. Dessouky, F. Koushanfar, A.-R. Sadeghi, T. Schneider, S. Zeitouni. "Automated Synthesis of Optimized Circuits for Secure Computation". In: *CCS'15*. ACM, 2015, pp. 1504–1517.

[29] D. Demmler, T. Schneider, M. Zohner. "ABY – A Framework for Efficient Mixed-Protocol Secure Two-Party Computation". In: *NDSS'15*. Internet Society, 2015.

[30] J. Doerner, D. Evans, A. Shelat. "Secure Stable Matching at Scale". In: *CCS'16*. ACM, 2016, pp. 1602–1613.

[31] T. K. Frederiksen, M. Keller, E. Orsini, P. Scholl. "A Unified Approach to MPC with Preprocessing Using OT". In: *ASIACRYPT'15*. Springer, 2015, pp. 711–735.

[32] R. Gilad-Bachrach, N. Dowlin, K. Laine, K. Lauter, M. Naehrig, J. Wernsing. "CryptoNets: Applying Neural Networks to Encrypted Data with High Throughput and Accuracy". In: *International Conference on Machine Learning (ICML'16)*. 2016, pp. 201–210.

[33] O. Goldreich, S. Micali, A. Wigderson. "How to Play any Mental Game or A Completeness Theorem for Protocols with Honest Majority". In: *STOC'87*. ACM, 1987, pp. 218–229.

[34] C. Guo, J. Katz, X. Wang, Y. Yu. "Efficient and Secure Multiparty Computation from Fixed-Key Block Ciphers". In:

S&P'20. To appear. Online: https://ia.cr/2019/074. IEEE, 2020.

[35] S. Halevi. "Advanced Cryptography: Promise and Challenges." In: *CCS'18*. ACM, 2018, p. 647.

[36] M. Hastings, B. Hemenway, D. Noble, S. Zdancewic. "SoK: General Purpose Compilers for Secure Multi-Party Computation". In: *S&P'19*. IEEE, 2019, pp. 1220–1237.

[37] E. Hauck, J. Loss. "Efficient and Universally Composable Protocols for Oblivious Transfer from the CDH Assumption." In: *IACR Cryptology ePrint Archive, Report 2017/1011* (2017). https://ia.cr/2017/1011.

[38] C. Hazay, P. Scholl, E. Soria-Vazquez. "Low Cost Constant Round MPC Combining BMR and Oblivious Transfer". In: *ASIACRYPT'17*. Springer. 2017, pp. 598–628.

[39] K. He, L. Yang, J. Hong, J. Jiang, J. Wu, X. Dong, Z. Liang. "PrivC—A Framework for Efficient Secure Two-Party Computation". In: *Security and Privacy in Communication Networks*. Springer, 2019, pp. 394–407.

[40] W. Henecka, S. Kögl, A.-R. Sadeghi, T. Schneider, I. Wehrenberg. "TASTY: Tool for Automating Secure Two-partY computations". In: *CCS'10*. ACM, 2010, pp. 451–462.

[41] Y. Huang, D. Evans, J. Katz. "Private Set Intersection: Are Garbled Circuits Better than Custom Protocols?" In: *NDSS'12*. Internet Society, 2012.

[42] IETF. "QUIC: A UDP-Based Multiplexed and Secure Transport". *https://tools.ietf.org/html/draft-ietf-quic-transport-23*.

[43] IETF. "Reliable UDP (RUDP) Protocol". *https://tools.ietf.org/html/draft-ietf-sigtran-reliable-udp-00*.

[44] R. Impagliazzo, S. Rudich. "Limits on the Provable Consequences of One-Way Permutations". In: *STOC'89*. ACM. 1989, pp. 44–61.

[45] Y. Ishai, J. Kilian, K. Nissim, E. Petrank. "Extending Oblivious Transfers Efficiently". In: *CRYPTO'03*. Springer. 2003, pp. 145–161.

[46] M. Ishaq, A. L. Milanova, V. Zikas. "Efficient MPC via Program Analysis: A Framework for Efficient Optimal Mixing". In: *CCS'19*. ACM, 2019, pp. 1539–1556.

[47] S. Kamara, P. Mohassel, M. Raykova. "Outsourcing Multi-Party Computation". In: *IACR Cryptology ePrint Archive, Report 2011/272* (2011). https://ia.cr/2011/272.

[48] M. Keller. "MP-SPDZ: A Versatile Framework for Multi-Party Computation". In: *IACR Cryptology ePrint Archive, Report 2020/521* (2020). https://ia.cr/2020/521.

[49] M. Keller, E. Orsini, P. Scholl. "MASCOT: Faster Malicious Arithmetic Secure Computation with Oblivious Transfer". In: *CCS'16*. ACM, 2016, pp. 830–842.

[50] V. Kolesnikov, A.-R. Sadeghi, T. Schneider. "Improved Garbled Circuit Building Blocks and Applications to Auctions and Computing Minima". In: *CANS'09*. Springer, 2009, pp. 1–20.

[51] B. Kreuter, B. Mood, A. Shelat, K. Butler. "PCF: A Portable Circuit Format for Scalable Two-party Secure Computation". In: *USENIX Security'12*. USENIX Association, 2013, pp. 321–336.

[52] B. Kreuter, A. Shelat, C.-H. Shen. "Billion-gate Secure Computation with Malicious Adversaries". In: *USENIX Security'12*. USENIX Association, 2012, pp. 285–300.

[53] C. Liu, X. S. Wang, K. Nayak, Y. Huang, E. Shi. "ObliVM: A Programming Framework for Secure Computation". In: *S&P'15*. IEEE, 2015, pp. 359–376.

[54] J. Liu, M. Juuti, Y. Lu, N. Asokan. "Oblivious Neural Network Predictions via MiniONN Transformations". In: *CCS'17*. ACM. 2017, pp. 619–631.

[55] D. Malkhi, N. Nisan, B. Pinkas, Y. Sella. "Fairplay – A Secure Two-Party Computation System". In: *USENIX Security'04*. USENIX Association, 2004, pp. 287–302.

[56] A. Mittos, B. Malin, E. D. Cristofaro. "Systematizing Genome Privacy Research: A Privacy-Enhancing Technologies Perspective". In: *PETS* 1 (2019), pp. 87–107.

[57] P. Mohassel, Y. Zhang. "SecureML: A System for Scalable Privacy-Preserving Machine Learning". In: *S&P'17*. IEEE, 2017, pp. 19–38.

[58] P. Mohassel, P. Rindal. "ABY^3: A Mixed Protocol Framework for Machine Learning". In: *CCS'18*. ACM, 2018, pp. 35–52.

[59] B. Mood, D. Gupta, H. Carter, K. Butler, P. Traynor. "Frigate: A Validated, Extensible, and Efficient Compiler and Interpreter for Secure Computation". In: *EuroS&P'16*. IEEE, 2016, pp. 112–127.

[60] A. Patra, A. Suresh. "BLAZE: Blazing Fast Privacy-Preserving Machine Learning". In: *NDSS'20*. Internet Society, 2020.

[61] M. O. Rabin. "How To Exchange Secrets with Oblivious Transfer". Tech. rep. Harvard Aiken Computation Laboratory, 1981.

[62] R. Rachuri, A. Suresh. "Trident: Efficient 4PC Framework for Privacy Preserving Machine Learning". In: *NDSS'20*. Internet Society, 2020.

[63] A. Rastogi, M. A. Hammer, M. Hicks. "Wysteria: A Programming Language for Generic, Mixed-mode Multiparty Computations". In: *S&P'14*. IEEE, 2014, pp. 655–670.

[64] M. S. Riazi, C. Weinert, O. Tkachenko, E. M. Songhori, T. Schneider, F. Koushanfar. "Chameleon: A Hybrid Secure Computation Framework for Machine Learning Applications". In: *ASIACCS'17*. ACM, 2018, pp. 707–721.

[65] P. Rindal. "libOTe: an Efficient, Portable, and Easy to Use Oblivious Transfer Library". *https://github.com/osu-crypto/libOTe*.

[66] T. Schneider, M. Zohner. "GMW vs. Yao? Efficient Secure Two-Party Computation with Low Depth Circuits". In: *FC'13*. Springer, 2013, pp. 275–292.

[67] K. Shrishak, H. Shulman, M. Waidner. "Removing the Bottleneck for Practical 2PC (Poster)". In: *CCS'18*. ACM, 2018, pp. 2300–2302.

[68] E. M. Songhori, S. U. Hussain, A.-R. Sadeghi, T. Schneider, F. Koushanfar. "TinyGarble: Highly Compressed and Scalable Sequential Garbled Circuits". In: *S&P'15*. IEEE, 2015, pp. 411–428.

[69] X. Wang. "A New Paradigm for Practical Maliciously Secure Multi-Party Computation". PhD thesis. PhD thesis. University of Maryland (College Park, Md.), 2018.

[70] X. Wang, A. J. Malozemoff, J. Katz. "EMP-toolkit: Efficient MultiParty computation toolkit". *https://github.com/emp-toolkit*. 2016.

[71] X. Wang, S. Ranellucci, J. Katz. "Global-Scale Secure Multiparty Computation". In: *CCS'17*. ACM, 2017, pp. 39–56.

[72] A. C.-C. Yao. "How to Generate and Exchange Secrets". In: *FOCS'86*. IEEE, 1986, pp. 162–167.

[73] S. Zahur, D. Evans. "Obliv-C: A Language for Extensible Data-Oblivious Computation". In: *IACR Cryptology ePrint Archive, Report 2015/1153* (2015). https://ia.cr/2015/1153.

[74] Y. Zhang, A. Steele, M. Blanton. "PICCO: A General-Purpose Compiler for Private Distributed Computation". In: *CCS'13*. ACM. 2013, pp. 813–826.

[75] W. Zheng, R. A. Popa, J. E. Gonzalez, I. Stoica. "Helen: Maliciously Secure Coopetitive Learning for Linear Models". In: *S&P'19*. IEEE, 2019, pp. 724–738.

A RUN-TIME AND COMMUNICATION COSTS

As briefly summarized in §8.1, we ran extensive microbenchmarks in MOTION.

We provide a detailed analysis of the communication and computation cost, depending on the number of parties N and bit length ℓ for primitive operations and conversions in MOTION in Tab. 6

We list the runtimes for primitive operations, sharing, reconstruction and conversion in Tab. 7. The runtimes for more complex building blocks, such as arithmetic operations and comparisons are provided in Tab. 8.

Also, we show an example of extensibility of the protocols implemented in MOTION in List. 2. Alternatively, this functionality can be implemented by inheriting from our `Gate` class to avoid the manual use of `Executor`.

Table 6: Total costs of primitive operations and conversions: the number of symmetric cryptographic operations, the number of bits sent by all parties, and the number of sequential messages required for the conversions.

	Computation [# symm. crypt. ops]		Communication [# bits]		# Messages	
	Setup	Online	Setup	Online	Setup	Online
ADD^A, XOR^B, XOR^Y	0	0	0	0	0	0
MUL^A	$2\ell N(N-1)$	0	$\ell N(N-1)(\kappa+\ell/2)$	$\ell N(N-1)$	2	1
AND^B	$2N(N-1)$	0	$N(N-1)\kappa$	$2N(N-1)$	1	1
AND^Y	$8N(2N-1)$	N^2	$N(N-1)((4N+7)\kappa+1)$	0	5	0
Share^A, Share^B	$2(N-1)$	0	0	0	0	0
Share^Y	0	0	0	$(N-1)(N\kappa+1)$	0	2
Rec^A	0	0	0	$\ell N(N-1)$	0	1
Rec^B	0	0	0	$N(N-1)$	0	1
Rec^Y	0	0	$N(N-1)$	0	1	0
Y2B	0	0	0	0	0	0
B2Y	0	0	0	$N(N-1)(\kappa+1)$	0	2
A2B, *A2Y*	$(\ell-1)\cdot 8(2N^3-3N^2+N)$	$(\ell-1)(N^3-N^2)$	$(\ell-1)\cdot N(N-1)^2((4N+7)\kappa+1)$	$\ell N(N-1)(N\kappa+1)$	5	2
B2A, *Y2A* (via *B*)	$\ell N(N-1)(\ell+2)$	0	$\ell N(N-1)(\ell+2)(\kappa+\ell+2)/2$	$N(N-1)\ell$	2	1
Y2A (w/o online comm.)	$(\ell-1)N^2(17N-9)$	$(\ell-1)N^2$	$N(N-1)\big(\ell(N\kappa+2)+(\ell-1)N((4N+7)\kappa+1)\big)$	0	5	0

Table 7: Run-times in nanoseconds in different test environments (cf. §8) for primitive operations for GMW (B) [33], arithmetic GMW (A), and the BMR protocol (Y) [8], and conversions between the protocols. For each entry, we specify the run-time of a single operation (amortized over one million operations for GMW and BMR operations, and conversions between them; over 100 thousand operations for arithmetic GMW operations, and over 1 thousand operations for the remaining conversions).

	LAN			WAN			AWS		
# Parties N	N=3	N=4	N=5	N=3	N=4	N=5	N=10	N=15	N=20
Rec_{8}^{A}	50.5	57.9	61.0	4 062.9	5 443.1	6 445.4	74.5	123.4	201.4
Rec_{16}^{A}	51.0	55.6	60.6	5 111.2	6 008.4	7 006.5	93.2	178.0	233.9
Rec_{32}^{A}	62.2	67.0	72.6	5 996.7	6 845.2	8 118.9	155.4	278.6	422.4
Rec_{64}^{A}	79.6	130.2	100.0	8 949.3	8 194.0	10 822.8	274.4	407.8	737.3
Rec^{B}	5.8	6.2	6.5	507.7	594.1	687.9	9.3	14.0	23.5
Rec^{Y}	7.9	7.9	8.6	508.0	596.6	687.1	9.9	15.9	20.1
Share_{8}^{A}	223.3	292.8	377.9	222.2	269.7	324.0	462.0	701.7	968.7
Share_{16}^{A}	180.3	240.7	302.6	203.8	240.6	279.5	467.6	730.4	980.8
Share_{32}^{A}	184.8	249.0	301.4	200.5	237.6	282.7	502.5	743.0	1 139.1
Share_{64}^{A}	161.0	219.1	272.6	184.2	222.1	255.1	524.5	568.9	1 239.6
Share^{B}	5.5	5.9	7.1	5.7	6.5	7.8	8.1	11.3	16.2
Share^{Y}	179.8	208.7	243.5	2 292.4	2 496.9	3 064.1	529.5	829.5	1 181.9
AND^{B}	545.5	747.8	829.0	3 116.9	3 725.3	4 640.8	1 729.1	2 777.0	4 743.0
AND^{Y}	5 022.4	7 690.1	10 809.1	18 218.3	21 852.6	27 557.8	50 727.0	111 913.3	237 933.2
XOR^{B}	3.6	3.6	3.3	3.9	4.1	4.1	3.9	4.2	9.3
XOR^{Y}	40.9	44.9	52.9	40.2	44.9	52.2	99.9	170.4	263.5
ADD_{8}^{A}	30.3	32.1	33.5	28.9	29.1	30.5	36.1	57.6	75.0
ADD_{16}^{A}	30.3	31.6	34.0	29.5	29.4	31.6	38.3	54.6	80.2
ADD_{32}^{A}	31.6	32.7	33.2	29.4	29.7	29.1	35.6	42.5	100.0
ADD_{64}^{A}	32.9	34.4	34.4	32.1	31.5	32.6	39.3	57.8	92.6
MUL_{8}^{A}	7 193.2	8 732.5	10 209.0	27 623.5	32 881.4	39 043.0	23 665.4	35 051.1	78 641.0
MUL_{16}^{A}	13 752.4	17 052.3	19 683.1	41 005.1	52 871.6	61 525.6	47 492.4	74 489.9	171 135.2
MUL_{32}^{A}	26 656.8	33 206.7	38 850.6	72 443.3	84 093.7	92 270.3	96 698.7	161 940.6	357 438.1
MUL_{64}^{A}	55 081.1	67 885.0	79 636.0	134 562.0	149 166.6	166 928.1	202 929.5	330 017.2	733 487.4
SQR_{8}^{A}	4 939.5	6 330.0	6 843.8	24 108.5	23 951.9	26 258.3	15 290.2	23 176.2	49 672.1
SQR_{16}^{A}	9 712.7	11 963.9	13 749.3	41 024.8	39 588.4	41 870.2	30 750.3	46 788.2	100 951.9
SQR_{32}^{A}	19 070.8	23 346.3	26 354.9	61 014.2	66 318.1	68 817.6	62 659.4	94 266.4	200 768.2
SQR_{64}^{A}	37 716.6	48 705.0	52 755.7	108 325.3	114 981.0	123 786.6	122 792.5	193 317.1	405 950.1
$A2B_8$	84 914.0	152 182.0	263 464.0	2 402 112.0	2 739 673.0	3 130 609.0	1 345 075.0	3 998 044.0	8 907 812.0
$A2B_{16}$	166 180.0	313 269.0	501 840.0	2 810 842.0	3 186 936.0	4 456 947.0	2 605 770.0	8 057 326.0	18 724 116.0
$A2B_{32}$	344 879.0	596 925.0	926 113.0	3 228 496.0	4 490 477.0	6 645 732.0	5 296 703.0	17 323 969.0	41 634 733.0
$A2B_{64}$	664 843.0	1 142 710.0	1 629 902.0	4 703 179.0	7 028 108.0	9 348 717.0	11 335 031.0	37 826 123.0	90 551 799.0
$A2Y_8$	85 514.0	260 872.0	260 897.0	2 399 031.0	2 671 459.0	3 170 296.0	1 356 501.0	3 966 832.0	9 148 202.0
$A2Y_{16}$	164 057.0	301 738.0	497 765.0	2 780 980.0	3 596 824.0	4 521 377.0	2 670 646.0	8 120 268.0	19 109 086.0
$A2Y_{32}$	330 767.0	584 865.0	894 217.0	3 207 824.0	4 453 959.0	6 369 496.0	5 347 239.0	17 262 721.0	41 456 717.0
$A2Y_{64}$	653 240.0	1 104 743.0	1 606 672.0	4 764 729.0	6 818 137.0	10 012 767.0	11 197 584.0	37 970 804.0	89 238 532.0
$B2A_8$	116 578.0	129 860.0	138 925.0	1 535 135.0	1 593 539.0	1 712 148.0	293 300.0	423 229.0	845 219.0
$B2A_{16}$	346 598.0	417 955.0	476 560.0	2 643 318.0	3 085 103.0	2 900 933.0	1 043 550.0	1 603 717.0	3 264 207.0
$B2A_{32}$	1 267 014.0	1 567 303.0	1 774 488.0	6 200 784.0	7 031 536.0	7 903 015.0	4 044 245.0	6 191 658.0	13 301 903.0
$B2A_{64}$	5 388 326.0	6 448 544.0	7 442 426.0	16 332 760.0	17 040 985.0	18 053 991.0	17 535 467.0	27 321 748.0	59 919 132.0
$B2Y_1$	149.3	194.5	210.1	1 909.6	1 956.3	2 258.3	498.7	795.9	1 111.7
$Y2A_8$	108 335.0	120 890.0	134 134.0	1 596 040.0	1 581 936.0	1 612 624.0	279 856.0	409 255.0	825 231.0
$Y2A_{16}$	330 551.0	405 120.0	458 771.0	2 689 310.0	3 009 141.0	3 318 748.0	1 057 533.0	1 557 023.0	3 243 724.0
$Y2A_{32}$	1 262 030.0	1 546 668.0	1 791 343.0	6 392 610.0	7 367 808.0	7 646 224.0	4 029 879.0	6 193 397.0	12 777 651.0
$Y2A_{64}$	5 375 035.0	6 356 610.0	7 399 503.0	16 266 624.0	17 033 616.0	18 124 902.0	17 699 223.0	27 253 770.0	58 012 547.0
$Y2B_1$	3.1	3.2	3.2	3.5	3.6	3.6	3.6	4.5	6.2

Table 8: Run-times in microseconds in different test environments (cf. §8) for a bit-string comparison (EQ) and integer operations for the GMW (B) [33] and BMR protocol (Y) [8]. For each entry, we specify the run-time of a single operation amortized over 1 000 SIMD values. We take the average over 100 protocol runs in the LAN and WAN environments and over 10 protocol runs in the AWS environment.

	LAN			WAN			AWS		
# Parties N	N=3	N=4	N=5	N=3	N=4	N=5	N=10	N=15	N=20
8-BIT INTEGERS									
EQ^B	97	114	128	995	1 086	1 106	219	240	421
EQ^Y	58	68	89	2 050	2 126	2 228	241	544	1 131
INT ADD^B	47	54	60	981	1 138	1 208	142	217	428
INT ADD^Y	52	66	93	1 461	2 051	2 274	257	525	1 113
INT DIV^B	552	1 006	1 316	11 637	12 229	12 872	2 121	2 791	3 274
INT DIV^Y	499	705	932	3 943	4 932	6 342	2 819	6 446	12 528
INT GT^B	81	94	109	1 273	1 298	1 361	190	226	411
INT GT^Y	59	69	91	1 921	2 011	2 236	272	589	1 166
INT MUL^B	100	121	131	1 899	2 218	1 966	259	379	683
INT MUL^Y	184	276	365	2 904	3 404	3 993	1 114	2 287	4 109
INT SUB^B	44	52	56	1 099	1 338	1 266	131	175	394
INT SUB^Y	53	61	84	2 039	2 125	2 213	243	496	1 080
16-BIT INTEGERS									
EQ^B	70	94	119	1 263	1 325	1 406	212	279	427
EQ^Y	97	123	172	2 178	2 389	2 557	490	1 076	2 265
INT ADD^B	75	88	97	1 447	1 605	1 789	194	294	486
INT ADD^Y	89	122	161	2 289	2 368	2 421	471	1 045	2 159
INT DIV^B	2 100	3 844	5 579	42 011	45 913	48 543	8 560	11 089	12 594
INT DIV^Y	1 659	2 302	3 007	8 678	11 092	14 292	9 795	22 840	47 875
INT GT^B	97	111	137	1 424	1 741	1 742	253	296	532
INT GT^Y	93	124	170	2 222	2 302	2 550	493	1 110	2 295
INT MUL^B	221	261	289	2 859	2 922	3 279	612	914	1 568
INT MUL^Y	650	868	1 090	4 961	6 472	7 966	4 023	8 595	15 629
INT SUB^B	82	88	92	1 284	1 460	2 050	187	290	511
INT SUB^Y	90	121	163	2 299	2 357	2 478	455	1 050	2 153
32-BIT INTEGERS									
EQ^B	103	136	164	1 540	1 640	2 113	312	551	529
EQ^Y	175	249	337	2 587	2 778	3 081	1 001	2 120	4 829
INT ADD^B	133	156	175	2 038	2 476	2 385	356	489	796
INT ADD^Y	163	245	340	2 469	2 761	3 070	947	2 196	4 410
INT DIV^B	5 596	9 671	13 044	100 048	108 743	113 700	20 950	26 834	32 313
INT DIV^Y	4 249	5 833	7 521	22 099	26 301	33 792	27 492	64 913	126 849
INT GT^B	120	151	167	1 956	2 343	2 325	316	426	633
INT GT^Y	171	263	372	2 646	2 862	3 219	993	2 252	4 445
INT MUL^B	703	826	911	4 001	5 280	6 574	1 953	2 902	5 609
INT MUL^Y	2 231	2 967	3 798	14 715	18 504	24 671	16 318	36 393	64 952
INT SUB^B	136	155	177	2 012	2 244	2 400	364	523	795
INT SUB^Y	160	236	337	2 553	2 701	3 002	981	2 109	4 277
64-BIT INTEGERS									
EQ^B	153	177	210	2 131	2 314	2 506	578	1 121	985
EQ^Y	336	475	640	3 121	3 399	4 022	2 017	4 593	12 323
INT ADD^B	269	305	353	2 209	2 563	2 613	699	998	1 965
INT ADD^Y	337	521	702	3 045	3 374	3 796	1 859	4 225	8 552
INT DIV^B	21 256	36 856	49 973	364 851	388 783	404 573	81 058	107 427	123 363
INT DIV^Y	14 806	20 042	26 360	66 339	83 557	115 663	100 945	238 141	461 013
INT GT^B	178	213	262	2 500	2 571	2 897	471	652	1 020
INT GT^Y	334	506	695	3 026	3 302	4 011	1 962	4 366	8 882
INT MUL^B	2 434	2 814	3 374	11 993	12 635	12 999	8 169	15 590	26 310
INT MUL^Y	8 588	11 618	14 485	53 813	67 671	92 458	64 278	144 322	252 703
INT SUB^B	272	309	347	2 213	2 303	2 662	725	1 004	2 029
INT SUB^Y	326	516	705	2 734	3 464	3 890	1 919	4 201	8 565

Listing 2: Code excerpt in MOTION for efficiently multiplying an integer known by one party in the clear by a secret-shared bit computed somewhere in the circuit. The result is of the multiplication is a secret-shared integer and can be used further in the circuit.

```
  using namespace MOTION;

/// \brief A secure two-party computation protocol for multiplying a
/// ring element known by one party in the clear by a bit secret-shared
/// between both parties, e.g., some intermediate result in Boolean GMW.
/// This protocol may be used, e.g., to aggregate statistics based on
/// threshold values and requires only one additively-correlated OT (ACOT).
///
/// @param s_bit secret-shared bit
/// @param val an unsigned integer number
/// \returns a secret-shared ring element in arith. GMW, i.e., s_bit * val

// knows share s_bit and integer val in the clear
template<typename T>
ArithmeticGMWWirePtr<T> MultiplyServer(BooleanGMWWirePtr s_bit, T val){
  auto& backend = s_bit->GetBackend();
  auto other_id = (backend.GetConfiguration().GetMyId() + 1) % 2;

  // register a sender OT object
  auto ot = backend.GetOTProvider(other_id).RegisterSendACOT<T>();

  // create an ArithmeticGMW wire for storing the result
  auto result = std::make_shared<ArithmeticGMWWire<T>>(backend);

  // submit the protocol as a task in form of an anonymous function
  // that will be executed in the online phase as a separate fiber
  backend.GetExecutor().EnqueueOnline(
  [s_bit, val, result, std::move(ot)]{
    ot->SetInput(val);
    ot->SendMessages();
    ot->ComputeOutput();
    // if s_bit is 1, we need to swap the output and s_bit might be a wire
    // deep in the circuit, so we need to wait until its value is computed
    s_bit->Wait();
    if(s_bit->As<bool>()) result->Set(ot->GetOutput(1));
    else result->Set(ot->GetOutput(0));
    // let other waiting routines know that this wire's value is now set
    result->SetFinished();
  });
  return result;
}

// knows only secret-shared bit s_bit
template<typename T>
ArithmeticGMWWirePtr<T> MultiplyClient(BooleanGMWWirePtr s_bit){
  auto& backend = s_bit->GetBackend();
  auto other_id = (backend.GetConfiguration().GetMyId() + 1) % 2;
  auto ot = backend.GetOTProvider(other_id).RegisterReceiveACOT<T>();
  auto result = std::make_shared<ArithmeticGMWWire<T>>(backend);

  backend.GetExecutor().EnqueueOnline(
  [s_bit, result, std::move(ot)]{
    s_bit->Wait();
    ot->SetCorrections(s_bit->As<bool>());
    ot->SendCorrections();
    ot->ComputeOutput();
    result->Set(ot->GetOutput());
    result->SetFinished();
  });
  return result;
}
```