



Research Article

MPCLan: Protocol Suite for Privacy-Conscious Computations*

Nishat Koti · Shravani Patil · Arpita Patra

Indian Institute of Science, Bangalore, India

kotis@iisc.ac.in

shravanip@iisc.ac.in

arpita@iisc.ac.in

Ajith Suresh

Technical University of Darmstadt, Darmstadt, Germany

suresh@encrypto.cs.tu-darmstadt.de

Communicated by David Pointcheval and Nigel Smart

Received 8 August 2022 / Revised 2 May 2023 / Accepted 3 May 2023

Online publication 24 May 2023

Abstract. The growing volumes of data being collected and its analysis to provide better services are creating worries about digital privacy. To address privacy concerns and give practical solutions, the literature has relied on secure multiparty computation techniques. However, recent research over rings has mostly focused on the small-party honest-majority setting of up to four parties tolerating single corruption, noting efficiency concerns. In this work, we extend the strategies to support higher resiliency in an honest-majority setting with efficiency of the online phase at the centre stage. Our semi-honest protocol improves the online communication of the protocol of Damgård and Nielsen (CRYPTO'07) without inflating the overall communication. It also allows shutting down almost half of the parties in the online phase, thereby saving up to 50% in the system's operational costs. Our maliciously secure protocol also enjoys similar benefits and requires only half of the parties, except for one-time verification towards the end, and provides security with fairness. To showcase the practicality of the designed protocols, we benchmark popular applications such as deep neural networks, graph neural networks, genome sequence matching, and biometric matching using prototype implementations. Our protocols, in addition to improved communication, aid in bringing up to 60–80% savings in monetary cost over prior work.

Keywords. Secure multiparty computation, Honest majority, Semi-honest, Malicious, Rings.

*This paper was reviewed by Jesper Buus Nielsen, Emmanuela Orsini, and an anonymous reviewer

1. Introduction

Today's world is seeing a visible transition from offline services to a heavy dependency on online platforms for banking, socializing, healthcare, etc. This is leading to an increased user presence online, which leaves a trail of online activity and personal data over the Internet. The availability of such user-specific data opens up possibilities for its misuse. For instance, there has been a lot of concern raised regarding advertisement service providers such as Google and Facebook breaching user privacy for targeted advertisement services [38]. In the process of providing enhanced targeted advertisement services, service providers are allegedly learning more information about their users than they are entitled to (e.g. user's shopping activity, browsing history) from various data collection entities. These entities collect user data via website cookies, loyalty cards, etc. [72]. While such targeted advertisements offer a personalized online experience, they may come at the cost of revealing unauthorized user data to these service providers. Such a challenge is also encountered in the healthcare sector. Collaborative analysis among healthcare institutes over patient data is known to facilitate better diagnosis and improved treatment. However, laws such as General Data Protection Regulation (GDPR) and the Consumer Privacy Act (CCPA), which prevent sharing of patient records, hinder such collaborations, thereby re-emphasizing the need for mechanisms that enable privacy-preserving computations.

Such mechanisms that ensure privacy-preserving computations can be facilitated via several privacy-enhancing technologies such as homomorphic encryption [18,47], differential privacy [42], and secure multiparty computation [12,49,98]. We focus on secure multiparty computation (MPC) as it has been the cornerstone of research lately, showcasing its effectiveness in various applications such as privacy-preserving machine learning [63,75,94], secure collaborative analytics [82], and secure genome matching [6,86]. Essentially, it offers a solution to the potential privacy issues which may arise in collaborative computation scenarios such as targeted advertisements described earlier.

MPC allows mutually distrusting parties to perform computations on their private inputs such that they learn nothing beyond the output of the computation. The distrust among the parties is captured by the notion of a centralized adversary, which is said to corrupt up to t out of the n participating parties. Depending on its behaviour, the adversary can be categorized as either *semi-honest* or *malicious* [48]. Semi-honest adversary models the corruption scenario where the corrupt parties are restricted to follow the protocol and cannot deviate arbitrarily, as in the stronger notion of malicious corruption.

MPC with honest majority, where only a minority of the parties are corrupt, enables construction of efficient protocols for multiple parties [2,13,15,35,54,82]. The recent concretely efficient protocols, operating over rings, have only considered small number of parties [26,32,63,73,75,81,93,94], which restricts the number of corruptions to at most one. Although the small-party setting has found application in the outsourced computation paradigm too, the *multiparty* setting is a better fit for real-world deployments due to its resiliency to a higher number of corruptions (up to $t < n/2$). Thus, for larger n , the number of corruptions that can be tolerated is also higher, thereby increasing the trust in the system. Moreover, the multiparty setting allows for privacy-conscious computations even in a non-outsourced deployment scenario, such as in providing targeted advertisement services (described in Fig. 1 and elaborated below), when outsourcing the

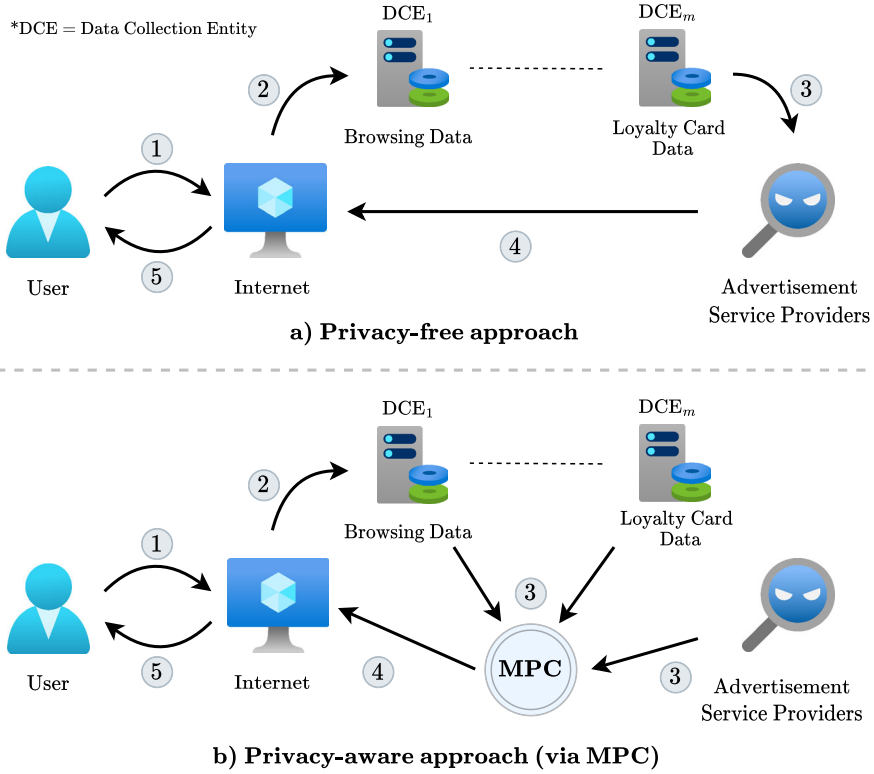


Fig. 1. Use case for privacy-conscious solutions.

computation is not feasible/preferable. Hence, to design efficient protocols, we focus on honest-majority multiparty computation.

Use Case. Consider the scenario of targeted advertisement services depicted in Fig. 1a. Typically, data collection entities track a user’s online activities via website cookies while browsing the Internet (①). Also known as cookie profiling, such data collection allows the entities to create a “profile” for each user, which may contain information such as browsing habits, gender, marital status, and age, as shown in (②). These profiles can facilitate targeted advertisements via specialized algorithms (④), which is leveraged by the advertisement service providers such as Google and Facebook. While such services offer a personalized experience, it comes at the expense of users’ private data being revealed to the service providers, as indicated in (③). A feasible solution (Fig. 1b) instead is to place a solution box at the interface between these service providers and the data collection entities such that it provides mechanisms to ensure the privacy of user data while also facilitating the required computations over the same (to provide targeted advertisement services). MPC being a technology that supports privacy-preserving computations lends itself well to such tasks. Instead of the data collection entities directly revealing the user data to the advertisement service providers, they can engage in an instance of MPC protocol (③) which securely runs the required algorithm on the user

data while maintaining its privacy. Moreover, such a computation does not require the data collection entities to reveal their data to each other, thus offering a viable solution. Furthermore, as studied in [55], the effectiveness of targeted advertisements can greatly benefit from the use of machine learning algorithms. In particular, neural networks and more recently graph neural networks [28, 69, 79, 97, 99] have shown the potential to better analyse the data available via user profiles, in turn allowing for a refined personalized experience. We thus focus on protocols for securely evaluating the standard neural networks such as VGG16 [91] (deep neural network) and graph neural network and provide benchmarks for the same in Sect. 6.

1.1. Related Work

Despite the interest in MPC for small population [2, 4, 5, 20, 25–27, 32, 45, 63, 64, 81, 94], MPC protocols for arbitrary number of parties (n) have been studied largely [3, 7, 9, 11, 13, 15, 17, 19, 21, 35, 44, 51, 54, 85] in the honest-majority ($t < n/2$) as well as the dishonest-majority ($t < n$) setting, where t denotes the maximum amount of allowed corruption. We restrict the related work to MPC protocols in the honest-majority setting, which is the focus of this work. In the honest-majority setting, protocols can be categorized as working over the field algebraic structure [35, 46, 52, 54] or rings [7, 13, 15, 17, 44]. The field-based protocols, which mostly operate over Shamir secret sharing [88] scheme, have the advantage of having the share size linear in the number of parties. On the other hand, ring-based protocols are proven to be practically more efficient since they can leverage CPU optimizations [14, 33, 36, 40, 89]. In the following, we cover the field-based protocols first, followed by ring-based ones.

Field-based protocols: In the semi-honest case, [35, 46] provide MPC protocols over fields in the information-theoretic setting. ATLAS [52] further improves upon the communication complexity of [35] in the information-theoretic setting from $12t$ field elements to $8t$ field elements per multiplication gate. ATLAS [52] also provides another protocol variant, which improves the round complexity of [35] by $2\times$ but requires slightly higher communication of $9t$ field elements. In the computational setting, the two protocol variants in ATLAS [52] roughly require communication of $4t$ and $5t$ field elements, respectively. The work of [44] demonstrates MPC protocols in the computational setting in the preprocessing model with malicious security. We observe that the semi-honest protocol derived from [44] requires communicating $2t$ elements in the online and $3t$ elements in the preprocessing phase.

In the malicious setting, the semi-honest protocol of [35] has served as the basis for obtaining malicious security for free (i.e. amortized communication cost of $3t$ field elements per multiplication gate) in the computational setting [17] as well as in the information-theoretic setting [52, 54]. These works follow the approach of executing a semi-honest protocol, followed by a verification phase to check the correctness of multiplication which involves heavy polynomial interpolation operations. As mentioned earlier, the work of [44] focuses on maliciously secure protocols for honest-majority setting in the preprocessing model. Their protocol relies on an instantiation of [54] in the preprocessing phase that requires communicating $3t$ field elements while requiring another $3t$ field elements communication in the online phase. However, their protocol is inefficient due to a consistency check required after each level of multiplication and

introduces depth-dependent overhead in communication complexity. The absence of this check results in a privacy breach as described in [53] and is elaborated in Sect. 4.

Ring-based protocols: Operating over rings is challenging because they do not have inverses for every element, which fields do. One way to work with rings is to adapt a field-based protocol to work over rings, but this can be computationally intensive due to the use of an extension field [1,44]. Another option is to use replicated secret shares (RSS) [57], which allows for direct operation over a ring without the need for extensions. However, this method results in share size becoming exponential in the number of parties due to the replication, but can be more efficient when the number of parties is constant.

The work of [15] shows how the honest-majority semi-honest field-based MPC protocol of [35] can be optimized to work over rings using RSS. Operating in the computational setting using a one-time setup for correlated randomness, this optimized version of [35] has a communication cost of $3t$ ring elements per multiplication gate. We will refer to this optimized honest-majority semi-honest protocol given in [15] (see section 7.4, page 63 of ePrint version) as DN07*. This protocol forms the state-of-the-art semi-honest protocol for honest-majority in the computational setting over rings and uses RSS. The work of [7,13] also provides semi-honest MPC protocols over rings in the computational setting, which require *each* party to communicate roughly t elements per multiplication gate, resulting in quadratic communication in the number of parties. The work of [44], as described earlier, showcases how their field protocols can be extended to work over *rings* using Galois ring extensions. The semi-honest protocol derived from their maliciously secure variant requires communicating $2t$ and $3t$ extended ring elements in the online and preprocessing phases, respectively. In the malicious setting, [44] suffers from the privacy breach over rings as well (see Sect. 4 for details). Further, both [15,17] provide protocols over rings. However, they rely on computationally heavy zero-knowledge machinery where expensive polynomial interpolation operations are carried out in the online phase.

1.2. Towards Practically Efficient Protocols

Before stating our contributions, we elaborate on the choices made in this work towards designing a practically efficient protocol for n parties, tolerating at most $t < n/2$ corruptions.

1. *Preprocessing paradigm.* With the goal of attaining as fast a response time as possible, the protocols are cast in the preprocessing paradigm [8,26,29,34,36,37,59–61,81,83]. Here, expensive *data-independent* computations are carried out in a preprocessing phase, thereby making way for a fast and efficient *data-dependent* online phase. We thus focus on improving the online phase without hampering the overall protocol complexity. Such a paradigm is apt for applications like ML-as-a-Service, where the same function is executed several times and is known beforehand.
2. *Algebraic structure.* Although operating over fields allows one to use techniques like Shamir secret sharing [88], where the number of shares are linear in the number of parties, we note that field operation bring in efficiency overhead [36,89]. This is because operating over fields requires reliance on external libraries, and

overloading basic operations such as addition and multiplication, since computer architectures are designed to operate on 32 or 64 bit rings. Thus, in an attempt to further enhance efficiency by utilizing the underlying CPU architecture, several protocols work over rings [26,63,64,73,75,94]. We follow this approach and design MPC protocols operating over the ring \mathbb{Z}_{2^ℓ} and rely on replicated secret sharing (RSS). Note that usage of RSS inherently results in exponential blow-up in the number of shares for an arbitrary number of parties. Hence, it is well suited for the practically oriented scenarios comprising a constant number of parties [15,17], which we restrict to for benchmarking our protocols.

3. *Masked evaluation.* To make our protocols efficient in the preprocessing paradigm, we use the masked evaluation paradigm, a variant of the replicated secret sharing scheme. The secret data are masked using a masking value in this case, and the mask is RSS shared. The computation is done on the publicly available masked values and the shared masks. This technique was first introduced in the context of circuit garbling schemes (see [71,96]) and was then adapted to secret sharing-based protocols in dishonest majority (see [10,58]). It was later applied to small-population honest-majority settings such as [25,32,50,63,81] and [64] to aid in the development of practically efficient protocols.
4. *Adversarial strategy.* Based on the deployment scenario, different levels of security may be desired. While semi-honest security suffices for several applications as shown in [5,6,24,25,66,77,86,92], malicious security is always desirable. When considering malicious security, we note that different levels of security notions can be attained such as (i) abort,¹ (ii) fairness,² and (iii) guaranteed output delivery.³ Thus, to cater to different scenarios, our protocols are designed to provide semi-honest and malicious security with fairness, where each security goal has its merit.
5. *Monetary cost.* To reduce the operational costs in the online phase, several recent works [26,63,64,81] reduce the number of (online) computing parties. This is useful in long computations such as those involved in privacy-preserving machine learning (PPML) applications, which span several days or even weeks. Reducing the number of online parties is especially advantageous for protocols deployed in the secure outsourced computation (SOC) setting since one has to pay for the uptime of every hired server. Shutting down even a single server significantly helps in reducing the monetary cost [64,74] of the system. We thus focus on ensuring the participation of a minimal number of parties during the online computation in our protocols. This is achieved for the first time in the multiparty (malicious) setting.⁴ Specifically, all the protocols for the semi-honest setting in our framework benefit from using only $t + 1$ parties in the online phase. The protocols in the malicious setting also enjoy this benefit except that the remainder t parties are required to come online for a short verification phase at the end. The reduction in online parties

¹Honest parties may not receive the output while corrupt parties do.

²Guarantees either all parties receive the output or none do.

³Guarantees output delivery to all parties regardless of adversarial misbehaviour.

⁴A recent work [44] also claims to achieve this reduction in online parties. However, their protocol suffers from a privacy breach as explained in Sect. 4, demanding a verification using all the parties to be carried out after each multiplicative level of the circuit.

aids in improving the operational cost of the framework by almost 50%. This is unlike prior works [15, 17, 35, 52, 54] which require active participation from all parties throughout the computation.

1.3. Our Contributions

We begin with a quick overview of the contributions of this work, followed by the details.

- We construct an n -party semi-honest protocol, tolerating at most $t < n/2$ corruptions, in the preprocessing paradigm which offers an improved online phase than the (optimized) protocol of DN07* [15], without inflating its total cost. Moreover, our protocol reduces the number of active parties in the online phase, thereby improving the system’s operational cost.
- We extend our semi-honest protocol to the malicious setting, while retaining the benefits of requiring reduced number of parties in online phase for majority of the computation. Our offer over state-of-the-art protocol of [44] is a stronger security guarantee of fairness, and at least $2\times$ improvement in round complexity via a one-time verification at the end of protocol evaluation.
- We provide support for 3 and 4 input multiplication, at the same online complexity as that of the 2 input multiplication. In addition to improving the communication cost over the approach of sequential multiplications, multi-input multiplication offers a $2\times$ improvement in the round complexity which is beneficial for high latency networks.

Moreover, the approach can be extended to an arbitrary number of inputs while retaining the same online communication, albeit requiring exponential communication in the preprocessing phase [80].

- We design building blocks for a range of applications such as deep neural networks, graph neural networks, genome sequence matching, and biometric matching. When the applications are benchmarked, our semi-honest protocol witnesses a saving of up to 69% in monetary cost and has $3.5\times$ to $4.6\times$ improvements in online run-time and throughput over DN07*. Interestingly, our maliciously secure protocols outperform the semi-honest protocol of DN07* in terms of online run-time and throughput for the applications under consideration, achieving the goal of fast online phase.

We now elaborate on the contributions and highlight the technical details and novelty of our work.

Our protocol suite follows a 3-tier architecture (Fig. 2) to attain the final goal of privacy-conscious computations. The first tier comprises fundamental primitives such as input sharing, reconstruction, multiplication (with truncation), and multi-input multiplication. The second tier includes building blocks such as dot product, matrix multiplication, conversion between Boolean and arithmetic worlds, comparison, equality, and nonlinear activation functions, as required in the applications considered. Finally, the third tier is applications. Our main contributions lie in Tier I, and these are detailed below. Going ahead, we use ‘multiparty protocols’ to mean honest-majority n -party protocols that tolerate $t > 1$ corruptions and thus do not include the tailor-made protocols in the 3 and 4 party setting.

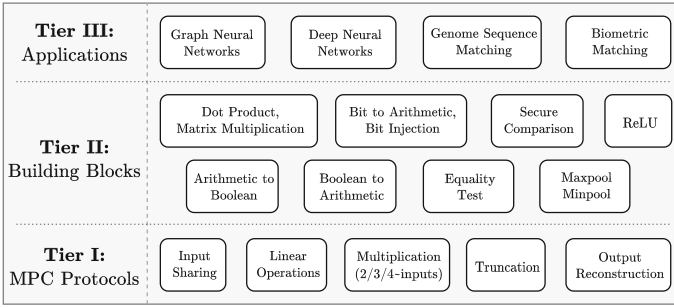


Fig. 2. Hierarchy of primitives in our 3-tier framework.

Tier I—MPC protocols Our goal is to design protocols with a fast online phase. Thus, working over \mathbb{Z}_{2^ℓ} and relying on RSS, we design a semi-honest MPC protocol in the computational setting assuming a one-time shared-key setup for correlated randomness.

Note that the straightforward extension of semi-honest multiplication protocol of DN07* to the preprocessing model, which can also be derived from the recent work of [44], incurs a communication of $3t$ elements in the preprocessing phase while communicating $2t$ elements in the online. This amounts to a $1.6\times$ overhead in the total cost over DN07*. Our contribution lies in ensuring a fast online phase, without inflating the total communication cost of the protocol. Specifically, our protocol requires communicating only $2t$ ring elements in the online phase and t in the preprocessing, for a multiplication gate. Thus, in the honest-majority multiparty setting over rings, we are the first to achieve a communication cost of $2t$ in the online phase (unlike $3t$ in the prior works [35,46]), without incurring any overhead in the total cost, i.e. our total cost still matches that of the best known (optimized) semi-honest honest-majority protocol [35,46].

We extend our protocol to provide malicious security with *fairness* at the cost of additionally communicating t elements in the online phase and $2t$ in the preprocessing phase. Although (*abort*) protocol of [44] has the same communication as our maliciously secure protocol, we achieve a stronger security notion of fairness. Moreover, [44] requires an additional round of communication for consistency checks after each level, the absence of which results in a privacy breach (described in [53] and elaborated in Sect. 4) and necessitates participation from all parties. However, by relying on a variant of RSS, our protocol avoids the consistency check after each level of circuit evaluation and ensures privacy. Notably, we only require participation from all parties for a one-time verification at the end of evaluation, thus reducing the number of rounds by d . (d denotes circuit depth.)

3- and 4-input multiplications Following [64,78,80], to reduce the online communication cost and round complexity, we design protocols to enable the multiplication of 3 and 4 inputs in a single shot. Compared to the naive approach of performing sequential multiplications to multiply 3/4 inputs, the *multi-input multiplication* protocol enjoys the benefit of having the same online phase complexity as that of the 2-input multiplication protocol. This brings in a $2\times$ improvement in the online round complexity and improves the online communication cost. Support for multi-input multiplication enables usage

of optimized adder circuits [80] for secure comparison and Boolean addition, thereby resulting in a faster online phase. The recent work of [52] also proposes a method to improve the round complexity of circuit evaluation by evaluating all gates in two consecutive layers in a circuit in parallel. We observe that their method can be viewed as a variant of multi-input multiplication with 3 and 4 inputs. Thus, our protocols need not be limited to facilitate faster comparison and Boolean additions alone (as described above), but can be used to reduce the round and communication complexity of any general circuit evaluation. Note that [52] only improves the round complexity ($2\times$) without inflating the communication cost when compared to DN07*. However, we focus on improving round complexity ($2\times$) *as well as* communication of the online phase by trading off an increase in the preprocessing.

Tier II—Building blocks We design efficient protocols for several building blocks in semi-honest and malicious settings, which are stepping stones for Tier III applications. These are extensions from the small-party setting [63,75,80,81], and we provide the details in Sect. 5.1 (semi-honest) and Sect. 5.2 (malicious).

Tier III—Applications To showcase the practicality of our framework and improvements of our protocols, we benchmark a range of applications such as neural networks (NN), which also includes the popular deep NN called VGG16 [91], graph neural network, genome sequence matching, and biometric matching and is considered for the first time in the n -party honest-majority setting. We benchmark the applications in the WAN setting using Google Cloud instances. As mentioned, owing to the inherent restrictions of RSS and keeping the focus on practical scenarios, we showcase the performance of our protocols for $n = 5, 7, \text{ and } 9$ and compare with the state-of-the-art semi-honest protocol of DN07*.

1. *Deep neural networks.* We benchmark inference phases of deep neural networks such as LeNet [67] and VGG16 [91]. We observe savings of up to 69% in monetary cost, and improvements of up to $4.3\times$ in online run-time and throughput, in comparison with DN07*.
2. *Graph neural network.* We benchmark the inference phase of graph neural network [39,90] on MNIST [68] data set. In comparison with DN07*, our protocol improves up to $3.5\times$ in online run-time and sees up to 15% savings in monetary cost.
3. *Genome sequence matching.* We demonstrate an efficient protocol for similar sequence queries (SSQ), which can be used to perform secure genome matching. Our protocol is based on the protocol of [86] which works for 2 parties and uses an edit distance approximation [6]. We extend and optimize the protocol for the multiparty setting. In comparison with DN07*, we witness improvements of up to $4\times$ in online run-time and throughput, and savings of 66% in monetary cost.
4. *Biometric matching.* We propose efficient protocols for computing Euclidean distance (ED), which forms the basis for performing biometric matching. Continuing the trend, we witness a $4.6\times$ improvement in online run-time and throughput over DN07*, and savings of up to 85% in monetary cost.

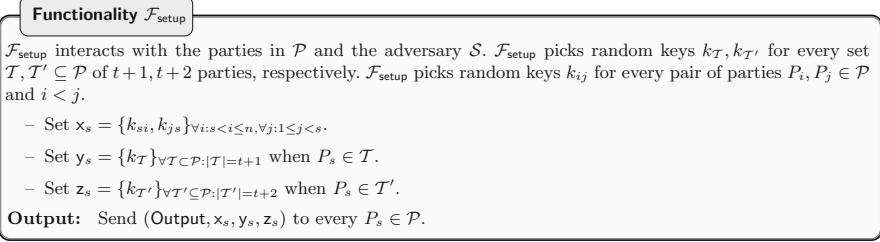


Fig. 3. Ideal functionality for shared-key setup.

2. Preliminaries

Our protocols are designed for rings (\mathbb{Z}_{2^ℓ}) algebraic structure and follows the (function-dependent) preprocessing paradigm to enable a fast online phase. For machine learning applications considered in this work, we use fixed-point arithmetic (FPA) [22, 23, 63, 75, 77] representation to operate over decimal values. Here, a decimal value is represented as an ℓ -bit integer in signed 2's complement representation. The most significant bit (msb) represents the sign bit, and d least significant bits are reserved for the fractional part. The ℓ -bit integer is then treated as an element of \mathbb{Z}_{2^ℓ} , and operations are performed modulo 2^ℓ . We let $\ell = 64$, $d = 13$, with $\ell - d - 1$ bits for the integer part.

Shared-key setup $\mathcal{F}_{\text{setup}}$ [5, 63, 75, 81] enables establishment of common random keys for a pseudo-random function (PRF) F , among parties. This aids in non-interactively generating correlated randomness. Here, $F : \{0, 1\}^k \times \{0, 1\}^k \rightarrow X$ is a secure PRF, with co-domain X being \mathbb{Z}_{2^ℓ} . The semi-honest functionality, $\mathcal{F}_{\text{setup}}$, appears in Fig. 3. The functionality for the malicious case is similar, except that the adversary now has the capability to `abort`.

To sample a random value $r \in \mathbb{Z}_{2^\ell}$ among a set of $t+1$ parties $\mathcal{T} = \{P_1, \dots, P_{t+1}\}$ non-interactively, each $P_i \in \mathcal{T}$ invokes $F_{k_{\mathcal{T}}}(id_{\mathcal{T}})$ and obtains r . Here, $id_{\mathcal{T}}$ denotes a counter maintained by the parties in \mathcal{T} and is updated after every PRF invocation. The appropriate keys used to sample are implicit from the context, from the identities of the parties that sample.

Collision-Resistant Hash Function A family of hash functions [84] $\{\mathcal{H} : \mathcal{K} \times \mathcal{M} \rightarrow \mathcal{Y}\}$ is said to be collision resistant if for all PPT adversaries \mathcal{A} , given the hash function \mathcal{H}_k for $k \in_R \mathcal{K}$, the following holds: $\Pr[(x, x') \leftarrow \mathcal{A}(k) : (x \neq x') \wedge \mathcal{H}_k(x) = \mathcal{H}_k(x')] = \text{negl}(\kappa)$, where $x, x' \in \{0, 1\}^m$, $m = \text{poly}(\kappa)$, and κ is security parameter.

Commitment Scheme Let $\text{Com}(x)$ denote the commitment of a value x [76]. The commitment scheme $\text{Com}(x)$ possesses two properties; *hiding* and *binding*. The former ensures privacy of value x given its commitment $\text{Com}(x)$, while the latter prevents a corrupt party from opening the commitment to a different value $x' \neq x$.

Our System Model This work considers both semi-honest and malicious adversarial models with *static* and at most $t < n/2$ corruptions. For the rest of the paper, we assume maximal corruption in this setting and thus $n = 2t + 1$. The security of our constructions is proved in the standalone, simulation-based security model of MPC, using the

real-world/ideal-world simulation paradigm [70] against a *computationally bounded* adversary, and the details are provided in Sect. 7. Let $\mathcal{P} = \{P_1, P_2, \dots, P_n\}$ denote the set of n parties which are connected by pair-wise private and authentic channels in a *synchronous* network. Additionally, our fair reconstruction protocol in the malicious setting relies on a broadcast channel, which can be instantiated using an existing broadcast protocol such as [41]. Set $\mathcal{E} = \{P_1, P_2, \dots, P_{t+1}\}$, termed as the *evaluator* set, comprises parties that are active during the online phase. Set $\mathcal{D} = \{P_{t+2}, P_{t+3}, \dots, P_n\}$, termed as the *helper* set, comprises parties which help in the preprocessing phase, and in the online verification in the malicious setting. Parties agree on a $P_{\text{king}} \in \mathcal{E}$. Without loss of generality, let $P_{\text{king}} = P_{t+1}$.

Sharing semantics We use the following sharing semantics, based on RSS and additive sharing schemes, which facilitate a fast online phase.

- $\langle \cdot \rangle$ -*sharing*: This denotes the replicated secret sharing (RSS) of a value with threshold t . A value $\mathbf{a} \in \mathbb{Z}_{2^\ell}$ is said to be RSS-shared with threshold t if for every subset $\mathcal{T} \subset \mathcal{P}$ of $n - t$ parties there exists $\langle \mathbf{a} \rangle_{\mathcal{T}} \in \mathbb{Z}_{2^\ell}$ possessed by all $P_i \in \mathcal{T}$ such that $\mathbf{a} = \sum_{\mathcal{T}} \langle \mathbf{a} \rangle_{\mathcal{T}}$.
Alternatively, for every set of t parties, the residual $h = n - t$ parties forming the set \mathcal{T} hold the share $\langle \mathbf{a} \rangle_{\mathcal{T}}$. Let $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_q \subset \mathcal{P}$ be the distinct subsets of size h , where $q = \binom{n}{h}$ represents the total number of shares. Since P_i belongs to $\binom{n-1}{h-1}$ such sets, it holds a tuple of $\binom{n-1}{h-1}$ shares, $\{\langle \mathbf{a} \rangle_{\mathcal{T}}\}$. We denote this tuple of shares that it possesses as $\langle \mathbf{a} \rangle_i$.
- $[\cdot]$ -*sharing*: A value $\mathbf{a} \in \mathbb{Z}_{2^\ell}$ is said to be $[\cdot]$ -shared (additively shared) among parties in \mathcal{P} if $P_i \in \mathcal{P}$ possesses $[\mathbf{a}]_i \in \mathbb{Z}_{2^\ell}$ such that $\mathbf{a} = [\mathbf{a}]_1 + [\mathbf{a}]_2 + \dots + [\mathbf{a}]_n$.
- $^{\mathcal{T}}[\cdot]$ -*sharing*: A value $\mathbf{a} \in \mathbb{Z}_{2^\ell}$ is said to be $^{\mathcal{T}}[\cdot]$ -shared among $t + 1$ parties in \mathcal{T} , if each $P_i \in \mathcal{T}$ holds $^{\mathcal{T}}[\mathbf{a}]_i$ such that $\mathbf{a} = \sum_{P_i \in \mathcal{T}} ^{\mathcal{T}}[\mathbf{a}]_i$. We refer to this sharing scheme as $(t + 1)$ -additive sharing and use $^{\mathcal{E}}[\mathbf{a}]$ to denote such a sharing among parties in \mathcal{E} .
- $\langle\langle \cdot \rangle\rangle$ -*sharing*: A value $\mathbf{a} \in \mathbb{Z}_{2^\ell}$ is said to be $\langle\langle \cdot \rangle\rangle$ -shared in the semi-honest setting if there exist values $\lambda_{\mathbf{a}}, \mathbf{m}_{\mathbf{a}} \in \mathbb{Z}_{2^\ell}$ such that $\mathbf{m}_{\mathbf{a}} = \mathbf{a} + \lambda_{\mathbf{a}}$ where $\lambda_{\mathbf{a}}$ is $\langle \cdot \rangle$ -shared among \mathcal{P} and every $P_i \in \mathcal{E}$ holds $\mathbf{m}_{\mathbf{a}}$. We denote the shares of $P_i \in \mathcal{D}$ by $\langle\langle \mathbf{a} \rangle\rangle_i = \langle \lambda_{\mathbf{a}} \rangle_i$ and that of $P_i \in \mathcal{E}$ as $\langle\langle \mathbf{a} \rangle\rangle_i = (\mathbf{m}_{\mathbf{a}}, \langle \lambda_{\mathbf{a}} \rangle_i)$. In the malicious setting, $\mathbf{m}_{\mathbf{a}}$ is held by all parties, and $\langle\langle \mathbf{a} \rangle\rangle_i = (\mathbf{m}_{\mathbf{a}}, \langle \lambda_{\mathbf{a}} \rangle_i)$ for all $P_i \in \mathcal{P}$.

It is trivial to see that all the sharing schemes mentioned above are linear. This allows parties to compute linear operations such as addition and multiplication with constants locally. The Boolean world operates over \mathbb{Z}_2 , and we denote the corresponding Boolean sharing with a superscript \mathbf{B} . Notations are summarized in Table 1.

Helper primitives We use the primitives described in Table 2 from the literature [15, 17, 30, 81] in our protocols, and their details are presented next. The Boolean variants of corresponding primitives are denoted with a superscript \mathbf{B} .

- (1) $\Pi_{[0]} \rightarrow [0]$ (Fig. 6): To generate $[\cdot]$ -shares of 0, each party non-interactively samples two values, each with one of its neighbouring parties. A party's shares of 0 are defined as the difference between these values.

Table 1. Notations used in this work.

Notation	Description
$n = 2t + 1$	Total number of parties with t corrupt and $h = t + 1$ honest
$\mathcal{T}_1, \dots, \mathcal{T}_q$	$q = \binom{n}{h}$ distinct subsets of \mathcal{P} with $t + 1$ parties each
q	Number of replicated secret shares (RSS) of a value
$g = \binom{n-1}{h-1}$	Number of RSS shares of a value held by a party
\mathcal{E}	Evaluator parties (P_1, \dots, P_{t+1}) that actively carry out the computation
\mathcal{D}	Helper parties (P_{t+2}, \dots, P_n)
\mathbf{a}_i	i^{th} element of vector \vec{a}
$\vec{a} \odot \vec{b}$	dot product of vectors \vec{a} and \vec{b}
$\mathbf{A} \odot \mathbf{B}$	Multiplication of matrices \mathbf{A} and \mathbf{B}
$\mathbf{b}^{\mathbb{R}}$	Arithmetic (Ring) equivalent over \mathbb{Z}_{2^ℓ} of bit $\mathbf{b} \in \mathbb{Z}_2$
$v[i]$	i^{th} bit of ℓ -bit value $v \in \mathbb{Z}_{2^\ell}$
$\mathbf{m}_a = \mathbf{a} + \lambda_a$	Masked value \mathbf{m}_a for $\mathbf{a} \in \mathbb{Z}_{2^\ell}$ with mask $\lambda_a \in \mathbb{Z}_{2^\ell}$
$\mathbf{M}_{a_1 a_2 \dots a_k}$	$\prod_{i=1}^k \mathbf{m}_{a_i}$; Product of masked values $\mathbf{m}_{a_1}, \dots, \mathbf{m}_{a_k}$
$\Lambda_{a_1 a_2 \dots a_k}$	$\prod_{i=1}^k \lambda_{a_i}$; Product of masks $\lambda_{a_1}, \dots, \lambda_{a_k}$

Table 2. Description of helper primitives—all are non-interactive, except Π_{agree} .

Primitive	Input	Output
$\Pi_{[0]}$	–	$[\cdot]$ -sharing of 0
Π_{rand}	–	$\langle \cdot \rangle$ -sharing of a random value $r \in \mathbb{Z}_{2^\ell}$
Π_{pRand}	Identity of a party P_s	$\langle \cdot \rangle$ -sharing of a random value $r \in \mathbb{Z}_{2^\ell}$ s.t. P_s learns all shares
$\Pi_{\cdot \rightarrow \langle \cdot \rangle}$	$\mathbf{a} \in \mathbb{Z}_{2^\ell}$ held by at least $t + 1$ parties	$\langle \mathbf{a} \rangle$ -sharing
$\Pi_{\langle \cdot \rangle \rightarrow \mathcal{T}[\cdot]}$	$\langle \mathbf{a} \rangle$ -sharing, $\mathcal{T} \subset \mathcal{P}$ s.t. $ \mathcal{T} = t + 1$	$\mathcal{T}[\mathbf{a}]$ -sharing
$\Pi_{\langle \cdot \rangle \rightarrow [\cdot]}$	$\langle \mathbf{a} \rangle$ -sharing	$[\mathbf{a}]$ -sharing
$\Pi_{\langle \cdot \rangle \rightarrow \mathcal{T}[\cdot]}$	$\langle \mathbf{a} \rangle$ -sharing, $\mathcal{T} \subset \mathcal{P}$ s.t. $ \mathcal{T} = t + 1$	$\mathcal{T}[\mathbf{a}]$ -sharing
$\Pi_{\langle \cdot \rangle \rightarrow [\cdot]}$	$\langle \mathbf{a} \rangle$ -sharing	$[\mathbf{a}]$ -sharing
$\Pi_{\langle \cdot \rangle \rightarrow \langle \cdot \rangle}$	$\langle \mathbf{a} \rangle$ -sharing	$\langle \mathbf{a} \rangle$ -sharing
$\Pi_{\langle \cdot \rangle, \langle \cdot \rangle \rightarrow [\cdot]}$	$\langle \mathbf{a} \rangle$ -sharing, $\langle \mathbf{b} \rangle$ -sharing	$[\mathbf{ab}]$ -sharing
Π_{agree}	$\mathcal{P}, \vec{v}_1, \dots, \vec{v}_n$	‘continue’ if $\vec{v}_i = \vec{v}_j$ for all $P_i, P_j \in \mathcal{P}$, ‘abort’ otherwise
$\Pi_{\langle \cdot \rangle}$	private input $\mathbf{a} \in \mathbb{Z}_{2^\ell}$ held by party P_s	$\langle \mathbf{a} \rangle$ -sharing

- (2) $\Pi_{\text{rand}} \rightarrow \langle r \rangle$ (Fig. 5): To generate $\langle \cdot \rangle$ -shares of a random $r \in \mathbb{Z}_{2^\ell}$, every set of $t + 1$ parties non-interactively sample a random value using keys established during the setup phase and define r to be the sum of these values.
- (3) $\Pi_{\text{pRand}}(P_s) \rightarrow \langle r \rangle$ (Fig. 6): This protocol generates $\langle \cdot \rangle$ -shares of a random value r such that P_s learns all the shares. Every set of $t + 1$ parties non-interactively samples a random value together with P_s , using the keys established (for every set of $t + 2$ parties) during the setup phase.

- (4) $\Pi_{(\cdot) \rightarrow \langle \cdot \rangle}(\mathbf{a}) \rightarrow \langle \langle \mathbf{a} \rangle \rangle$: This protocol generates $\langle \langle \mathbf{a} \rangle \rangle$ when $\mathbf{a} \in \mathbb{Z}_{2^\ell}$ is held by at least $t + 1$ parties, say parties in \mathcal{E} . For this, $P_i \in \mathcal{E}$ sets $\mathbf{m}_a = \mathbf{a}$ and $\langle \cdot \rangle$ -shares of λ_a as 0. To generate $\langle \langle \mathbf{a} \rangle \rangle$ in the malicious case where all parties hold \mathbf{a} , we let parties set $\mathbf{m}_a = \mathbf{a}$ and shares of λ_a as 0.
- (5) $\Pi_{(\cdot) \rightarrow \mathcal{T}_{[\cdot]}}(\langle \mathbf{a} \rangle) \rightarrow \mathcal{T}[\mathbf{a}]$ (Fig. 7): This protocol enables parties in $\mathcal{T} = \{E_1, E_2, \dots, E_{t+1}\}$ to generate $\mathcal{T}[\mathbf{a}]$ from $[\mathbf{a}]$. To generate $\mathcal{T}[\mathbf{a}]_i$, the idea is to sum up the shares in $\langle \mathbf{a} \rangle_{\mathcal{T}_1}, \dots, \langle \mathbf{a} \rangle_{\mathcal{T}_q}$, while ensuring that every share is accounted for and no share is incorporated more than once. Concretely, for share $\langle \mathbf{a} \rangle_{\mathcal{T}_j}$ held by parties in \mathcal{T}_j for $j \in \{1, \dots, q\}$, $E_i \in \mathcal{T}_j$ incorporates $\langle \mathbf{a} \rangle_{\mathcal{T}_j}$ in its share of $\mathcal{E}[\mathbf{a}]_i$ if E_i has the least index in \mathcal{T}_j .
- (6) $\Pi_{(\cdot) \rightarrow [\cdot]}(\langle \mathbf{a} \rangle) \rightarrow [\mathbf{a}]$: $\langle \cdot \rangle$ -share can be converted to $[\cdot]$ -share following similar procedure as $\Pi_{(\cdot) \rightarrow \mathcal{T}_{[\cdot]}}$ and is denoted as $\Pi_{(\cdot) \rightarrow [\cdot]}(\langle \mathbf{a} \rangle)$. We omit the details due to similarity.
- (7) $\Pi_{\langle \cdot \rangle \rightarrow \mathcal{T}_{[\cdot]}}(\langle \langle \mathbf{a} \rangle \rangle) \rightarrow \mathcal{T}[\mathbf{a}]$: Parties in \mathcal{T} invoke $\Pi_{(\cdot) \rightarrow \mathcal{T}_{[\cdot]}}$ on $-\lambda_a$ to generate $\mathcal{T}[-\lambda_a]$, followed by a designated $P_i \in \mathcal{T}$ that holds \mathbf{m}_a setting $\mathcal{T}[\mathbf{a}]_i = \mathbf{m}_a + \mathcal{E}[-\lambda_a]_i$.
- (8) $\Pi_{\langle \cdot \rangle \rightarrow [\cdot]}(\langle \langle \mathbf{a} \rangle \rangle) \rightarrow [\mathbf{a}]$: $[\mathbf{a}]$ can be generated from $\langle \langle \mathbf{a} \rangle \rangle$ similar to $\Pi_{\langle \cdot \rangle \rightarrow \mathcal{T}_{[\cdot]}}$ and is denoted as $\Pi_{\langle \cdot \rangle \rightarrow [\cdot]}(\langle \langle \mathbf{a} \rangle \rangle)$.
- (9) $\Pi_{(\cdot) \rightarrow \langle \cdot \rangle}(\langle \mathbf{a} \rangle) \rightarrow \langle \langle \mathbf{a} \rangle \rangle$: To convert $\langle \mathbf{a} \rangle$, to $\langle \langle \mathbf{a} \rangle \rangle$, set $\mathbf{m}_a = 0$ and set $\langle \lambda_a \rangle = -\langle \mathbf{a} \rangle$.
- (10) $\Pi_{\langle \cdot \rangle \rightarrow (\cdot)}(\langle \langle \mathbf{a} \rangle \rangle) \rightarrow \langle \mathbf{a} \rangle$: To convert $\langle \langle \mathbf{a} \rangle \rangle$ to $\langle \mathbf{a} \rangle$, set $\langle \mathbf{a} \rangle_{\mathcal{T}_j} = -\langle \lambda_a \rangle_{\mathcal{T}_j}$ for $j \in \{1, \dots, q-1\}$ and $\langle \mathbf{a} \rangle_{\mathcal{T}_q} = \mathbf{m}_a - \langle \lambda_a \rangle_{\mathcal{T}_q}$, where $\mathcal{T}_q = \mathcal{E}$.
- (11) $\Pi_{(\cdot) \cdot (\cdot) \rightarrow [\cdot]}(\langle \mathbf{a} \rangle, \langle \mathbf{b} \rangle) \rightarrow [\mathbf{ab}]$ (Fig. 8): Given $\langle \mathbf{a} \rangle, \langle \mathbf{b} \rangle$, parties non-interactively compute $[\mathbf{ab}]$ as follows. Observe that $[\mathbf{ab}] = \sum_{j=1}^q [\langle \mathbf{a} \rangle_{\mathcal{T}_j} \mathbf{b}]$. To generate $[\langle \mathbf{a} \rangle_{\mathcal{T}_j} \mathbf{b}]$, the idea is to generate $\mathcal{T}_j[\langle \mathbf{a} \rangle_{\mathcal{T}_j} \mathbf{b}]$ and perform a conversion. Parties in \mathcal{T}_j generate $\mathcal{T}_j[\langle \mathbf{a} \rangle_{\mathcal{T}_j} \mathbf{b}]$ as $\mathcal{T}_j[\langle \mathbf{a} \rangle_{\mathcal{T}_j} \mathbf{b}] = (\langle \mathbf{a} \rangle_{\mathcal{T}_j}) \cdot (\mathcal{T}_j[\mathbf{b}])$. To obtain $[\langle \mathbf{a} \rangle_{\mathcal{T}_j} \mathbf{b}]$ from $\mathcal{T}_j[\langle \mathbf{a} \rangle_{\mathcal{T}_j} \mathbf{b}]$, $P_i \in \mathcal{P}$ sets $[\langle \mathbf{a} \rangle_{\mathcal{T}_j} \mathbf{b}]_i = \mathcal{T}_j[\langle \mathbf{a} \rangle_{\mathcal{T}_j} \mathbf{b}]_i$ if $P_i \in \mathcal{T}_j$ and $[\langle \mathbf{a} \rangle_{\mathcal{T}_j} \mathbf{b}]_i = 0$, otherwise.
- (12) $\Pi_{\text{agree}}(\mathcal{P}, \{\vec{v}_1, \dots, \vec{v}_n\}) \rightarrow \text{continue/abort}$: Allows parties to check if they hold the same set of values $\vec{v} = (v_1, \dots, v_m)$, where parties `continue` if the values are same, and `abort` otherwise. We denote the version of \vec{v} held by $P_i \in \mathcal{P}$ as \vec{v}_i . To check for consistency of \vec{v} , parties compute hash, $\mathbf{H} = \mathbf{H}(v_1 || \dots || v_m)$, of the concatenation of all values v_1, \dots, v_m , and exchange \mathbf{H} among themselves. If any party receives inconsistent hashes, it `aborts`; else it `continues`.
- (13) $\Pi_{(\cdot)}(P_s, \mathbf{a}) \rightarrow \langle \mathbf{a} \rangle$: To enable P_s to generate $\langle \mathbf{a} \rangle$, parties generate $\langle \mathbf{a} \rangle_{\mathcal{T}_j}$ for $j \in \{1, \dots, q-1\}$ using Π_{pRand} , with P_s learning $\langle \mathbf{a} \rangle_{\mathcal{T}_j}$ (i.e. $\langle \mathbf{a} \rangle_{\mathcal{T}_j}$ are sampled using common key among $t + 2$ parties). P_s sets $\langle \mathbf{a} \rangle_{\mathcal{T}_q} = \mathbf{a} - \sum_{j=1}^{q-1} \langle \mathbf{a} \rangle_{\mathcal{T}_j}$ and sends $\langle \mathbf{a} \rangle_{\mathcal{T}_q}$ to parties in \mathcal{T}_q . For malicious case, this is followed by invoking $\Pi_{\text{agree}}(\mathcal{P}, \{\langle \mathbf{a} \rangle_{\mathcal{T}_q}\})$ to check consistency of the values sent by P_s .

Protocol $\Pi_{[0]}$

1. P_i, P_{i+1} , for $i \in \{1, \dots, n-1\}$, sample a random value $r_i \in_R \mathbb{Z}_{2^\ell}$, while P_1, P_n sample a random value $r_n \in_R \mathbb{Z}_{2^\ell}$, using their respective common PRF keys.
2. P_i for $i \in \{2, \dots, n\}$ sets $[0]_i = r_i - r_{i-1}$, while P_1 sets $[0]_1 = r_1 - r_n$.

Fig. 4. Generating $[\cdot]$ -shares of 0 .**Protocol** Π_{rand}

1. Every $P_i \in \mathcal{T}_j$ for $j \in \{1, \dots, q\}$, samples $\langle r \rangle_{\mathcal{T}_j} \in_R \mathbb{Z}_{2^\ell}$ using the common PRF key.
2. Define $r = \sum_{j=1}^q \langle r \rangle_{\mathcal{T}_j}$.

Fig. 5. Generating $\langle \cdot \rangle$ -shares of a random value.**Protocol** $\Pi_{\text{pRand}}(P_s)$

1. Every $P_i \in \mathcal{T}_j$ for $j \in \{1, \dots, q\}$, samples $\langle r \rangle_{\mathcal{T}_j} \in_R \mathbb{Z}_{2^\ell}$, together with P_s , using the common PRF key.
2. Define $r = \sum_{j=1}^q \langle r \rangle_{\mathcal{T}_j}$.

Fig. 6. Generating $\langle \cdot \rangle$ -shares of a random value along with P_s .**Protocol** $\Pi_{\langle \cdot \rangle \rightarrow \mathcal{T}[\cdot]}(\langle a \rangle)$

1. Let $\mathcal{T} = \{E_1, \dots, E_{t+1}\}$.
2. $E_i \in \mathcal{T}$ computes $\mathcal{E}[a]_i = \sum_{j=1}^q \langle a \rangle_{\mathcal{T}_j} \cdot e_j^i$, where $e_j^i = 1$ if E_i has the least index in \mathcal{T}_j , and 0, otherwise.

Fig. 7. Conversion from $\langle \cdot \rangle$ -share to $\mathcal{T}[\cdot]$ -share .**Protocol** $\Pi_{\langle \cdot \rangle, \langle \cdot \rangle \rightarrow [\cdot]}(\mathcal{P}, \langle a \rangle, \langle b \rangle)$

1. For $j \in \{1, \dots, q\}$:
 - $P_i \in \mathcal{T}_j$ invokes $\Pi_{\langle \cdot \rangle, \mathcal{T}[\cdot]}$ on $\langle b \rangle$ to generate $\mathcal{T}_j[b]_i$.
 - Set $[\langle a \rangle_{\mathcal{T}_j} b]_i = \langle a \rangle_{\mathcal{T}_j} \cdot \mathcal{T}_j[b]_i$ if $P_i \in \mathcal{T}_j$, and $[\langle a \rangle_{\mathcal{T}_j} b]_i = 0$, otherwise.
2. $P_i \in \mathcal{P}$ computes $[\text{ab}]_i = \sum_{j=1}^q [\langle a \rangle_{\mathcal{T}_j} b]_i$.

Fig. 8. $\langle a \rangle, \langle b \rangle$ to $[\text{ab}]$.

3. MPCLan Protocol

The ideal functionality $\mathcal{F}_{\text{N-PC}}$ for evaluating function f in the n -party setting with semi-honest security appears in Fig. 9. Details of its instantiation over the ring \mathbb{Z}_{2^ℓ} that comprises three phases—input sharing, evaluation (linear operations and multiplication), and output reconstruction—appear next.

Input sharing and output reconstruction To enable $P_s \in \mathcal{P}$ to $\langle \cdot \rangle$ -share a value $v \in \mathbb{Z}_{2^\ell}$, parties first non-interactively sample $\langle \cdot \rangle$ -shares of λ_v , relying on the shared-

Functionality $\mathcal{F}_{n\text{-PC}}$

$\mathcal{F}_{n\text{-PC}}$ interacts with the parties in \mathcal{P} and the adversary \mathcal{S}^{sh} . Let f denote the function to be computed. Let x_s be the input corresponding to the party P_s , and y_s be the corresponding output, i.e. $(\{y_s\}_{s=1}^n) = f(\{x_s\}_{s=1}^n)$.

Step 1: $\mathcal{F}_{n\text{-PC}}$ receives (Input, x_s) from $P_s \in \mathcal{P}$, and computes $(\{y_s\}_{s=1}^n) = f(\{x_s\}_{s=1}^n)$.

Step 2: $\mathcal{F}_{n\text{-PC}}$ sends (Output, y_s) to $P_s \in \mathcal{P}$.

Fig. 9. Semi-honest: ideal functionality for function f .

Protocol $\Pi_{\text{Sh}}(P_s, a)$

Preprocessing: Invoke $\Pi_{\text{pRand}}(P_s)$ to generate $\langle \lambda_a \rangle$, with P_s learning $\lambda_a \in \mathbb{Z}_{2^\ell}$.

Online: P_s computes and sends $m_a = a + \lambda_a$ to all $P_i \in \mathcal{E}$.

Fig. 10. Semi-honest: input sharing protocol.

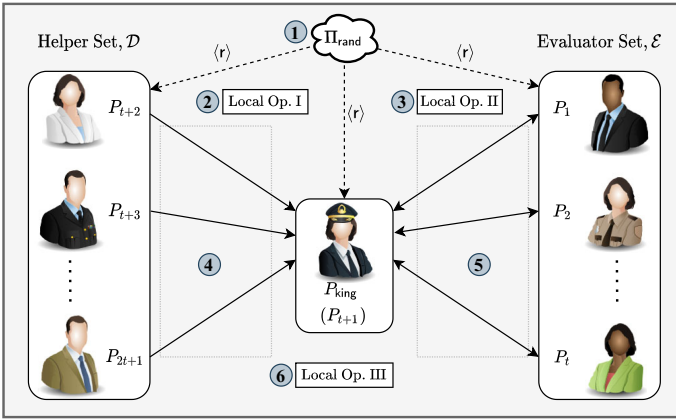


Fig. 11. Steps of semi-honest multiplication protocol.

key setup, such that P_s learns all these shares in clear (via Π_{pRand}). This enables P_s to compute and send $m_v = v + \lambda_v$ to parties in \mathcal{E} , thereby generating $\langle\langle v \rangle\rangle$. The protocol for input sharing appears in Fig. 10.

To reconstruct v towards all parties given $\langle\langle v \rangle\rangle$, observe that parties in \mathcal{E} possess sufficient shares to facilitate the same. Elaborately, parties in \mathcal{E} can non-interactively generate additive shares, $\mathcal{E}[v]$, among themselves (via $\Pi_{\langle\langle \cdot \rangle\rangle \rightarrow \mathcal{E}[-1]}$). These parties can then send their additive shares to P_{king} , who computes and sends v to all parties. Reconstruction towards a single party, say P_s , can proceed similarly except that the protocol terminates after parties in \mathcal{E} send their additive shares of v to $P_{\text{king}} = P_s$, who then computes v .

Evaluation Evaluation comprises linear operations of addition and multiplication with public constant, and nonlinear operations such as multiplication. Parties can non-interactively compute linear operations owing to the linearity of the $\langle\langle \cdot \rangle\rangle$ -sharing. Concretely, given $\langle\langle a \rangle\rangle$, $\langle\langle b \rangle\rangle$ and public constants c_1, c_2 , parties can non-interactively compute $\langle\langle c_1 a + c_2 b \rangle\rangle$ as $c_1 \langle\langle a \rangle\rangle + c_2 \langle\langle b \rangle\rangle$.

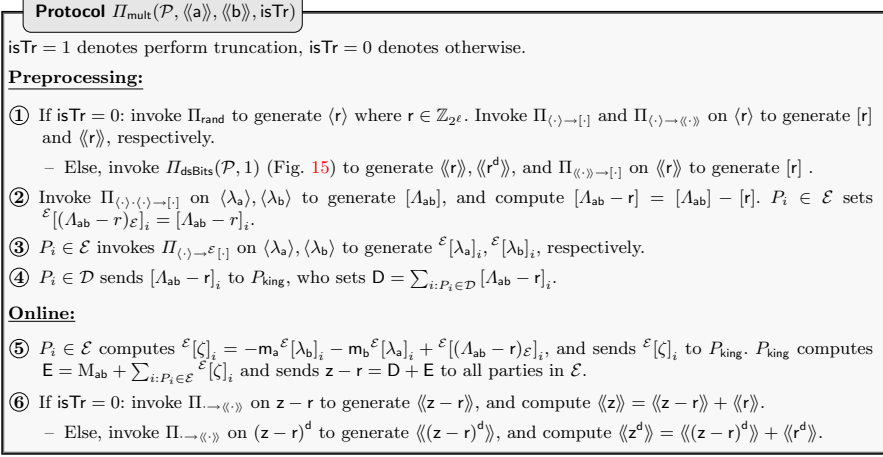


Fig. 12. Semi-honest: multiplication protocol .

To compute $\langle\langle \cdot \rangle\rangle$ -shares for nonlinear operations such as multiplication, say $z = \mathbf{ab}$ given $\langle\mathbf{a}\rangle, \langle\mathbf{b}\rangle$, parties proceed as follows. At a high level, the approach is to enable generation of $\langle\langle z - r \rangle\rangle$ and $\langle\langle r \rangle\rangle$ for a random $r \in \mathbb{Z}_{2^\ell}$, which enables parties to non-interactively compute $\langle\langle z \rangle\rangle = \langle\langle z - r \rangle\rangle + \langle\langle r \rangle\rangle$. Observe that $\langle\langle r \rangle\rangle$ can be generated non-interactively by locally sampling each of its shares. To generate $\langle\langle z - r \rangle\rangle$, we let parties in \mathcal{E} obtain $z - r$, following which $\langle\langle z - r \rangle\rangle$ can be generated non-interactively. (This is achieved via $\Pi_{\langle \cdot \rangle \rightarrow \langle \cdot \rangle}$ where all parties set their shares of $\langle\lambda_{z-r}\rangle$ as 0, and parties in \mathcal{E} set $m_{z-r} = z - r$.) Observe that z remains private while revealing $z - r$ to parties in \mathcal{E} since r is a random mask not known to adversary.

To enable parties in \mathcal{E} to obtain $z - r$, we let $z - r = D + E$, where D is additively shared among parties in \mathcal{D} , while E is additively shared among parties in \mathcal{E} . (D, E are defined in the following paragraphs.) Thus, to reconstruct $z - r$ towards parties in \mathcal{E} , parties send their respective additive shares of D or E towards $P_{\text{king}} \in \mathcal{P}$. P_{king} reconstructs D, E and sends $z - r = D + E$ to parties in \mathcal{E} . Elaborately, as seen in [25, 64], $z - r$ can be computed as

$$\begin{aligned}
 z - r &= \mathbf{ab} - r = (m_a - \lambda_a)(m_b - \lambda_b) - r = M_{ab} - m_a \lambda_b - m_b \lambda_a + A_{ab} - r \\
 &= \underbrace{M_{ab} - m_a \lambda_b - m_b \lambda_a + (A_{ab} - r)\mathcal{E}}_E + \underbrace{(A_{ab} - r)\mathcal{D}}_D
 \end{aligned} \tag{1}$$

where $A_{ab} - r = (A_{ab} - r)\mathcal{D} + (A_{ab} - r)\mathcal{E}$. The multiplication protocol Π_{mult} (Fig. 12) is detailed next, and its schematic representation is provided in Fig. 11.

- *Step* ①: Parties non-interactively generate $\langle r \rangle$ by locally sampling each of its shares (via Π_{rand}). Parties locally compute $[r]$ and $\langle\langle r \rangle\rangle$ from $\langle r \rangle$ using $\Pi_{\langle \cdot \rangle \rightarrow [\cdot]}$ and $\Pi_{\langle \cdot \rangle \rightarrow \langle \cdot \rangle}$, respectively. Looking ahead, $[r]$ aids in generating additive shares of D, E , while $\langle\langle r \rangle\rangle$ aids in computing $\langle\langle z \rangle\rangle$ from $\langle\langle z - r \rangle\rangle$.

- *Step ②*: This step involves computing additive shares of $\Lambda_{ab} - r$ among all parties. For this, parties non-interactively generate $[\Lambda_{ab}]$ from $\langle \lambda_a \rangle, \langle \lambda_b \rangle$ (via $\Pi_{(\cdot, \cdot) \rightarrow [\cdot]}$). $P_i \in \mathcal{P}$ sets its additive share of $\Lambda_{ab} - r$ as $[\Lambda_{ab} - r]_i = [\Lambda_{ab}]_i - [r]_i$. Observe that the shares $[\Lambda_{ab} - r]_i$ of $P_i \in \mathcal{D}$ define the additive shares of $\mathbf{D} = (\Lambda_{ab} - r)_{\mathcal{D}}$ among parties in \mathcal{D} . Similarly, the shares $[\Lambda_{ab} - r]_i$ of $P_i \in \mathcal{E}$ define the additive shares of $(\Lambda_{ab} - r)_{\mathcal{E}}$ among parties in \mathcal{E} (i.e. ${}^{\mathcal{E}}[(\Lambda_{ab} - r)_{\mathcal{E}}]$).
- *Step ③*: Parties in \mathcal{E} generate additive shares of λ_a, λ_b among themselves (${}^{\mathcal{E}}[\cdot]$ -shares, via $\Pi_{(\cdot) \rightarrow {}^{\mathcal{E}}[\cdot]}$). Looking ahead, ${}^{\mathcal{E}}[\lambda_a], {}^{\mathcal{E}}[\lambda_b]$ aid in generating additive shares of \mathbf{E} among \mathcal{E} .
- *Step ④*: Parties in \mathcal{D} send their additive shares of \mathbf{D} (as defined in step ②) to P_{king} , where the latter reconstructs \mathbf{D} .
- *Step ⑤*: $P_i \in \mathcal{E} \setminus \{P_{\text{king}}\}$ non-interactively generates additive share, ${}^{\mathcal{E}}[\mathbf{E}]_i$, of \mathbf{E} among parties in \mathcal{E} as ${}^{\mathcal{E}}[\mathbf{E}]_i = -m_a {}^{\mathcal{E}}[\lambda_b]_i - m_b {}^{\mathcal{E}}[\lambda_a]_i + {}^{\mathcal{E}}[(\Lambda_{ab} - r)_{\mathcal{E}}]_i$. Note that it suffices for only one designated party in \mathcal{E} to add M_{ab} in its share of ${}^{\mathcal{E}}[\mathbf{E}]$, and without loss of generality we let this designated party be P_{king} . For $P_{\text{king}} = P_{t+1}$ in our case, ${}^{\mathcal{E}}[\mathbf{E}]_{t+1} = M_{ab} - m_a {}^{\mathcal{E}}[\lambda_b]_{t+1} - m_b {}^{\mathcal{E}}[\lambda_a]_{t+1} + {}^{\mathcal{E}}[(\Lambda_{ab} - r)_{\mathcal{E}}]_{t+1}$. Parties send their additive shares of \mathbf{E} to P_{king} , who reconstructs \mathbf{E} and sends $\mathbf{z} - r = \mathbf{D} + \mathbf{E}$ to parties in \mathcal{E} .
- *Step ⑥*: Parties non-interactively generate $\langle\langle \mathbf{z} - r \rangle\rangle$ (via $\Pi_{\cdot \rightarrow \langle\langle \cdot \rangle\rangle}$) as explained earlier. Using $\langle\langle r \rangle\rangle$ generated in step ①, parties compute $\langle\langle \mathbf{z} \rangle\rangle = \langle\langle \mathbf{z} - r \rangle\rangle + \langle\langle r \rangle\rangle$, as required.

Lemma 1. *Protocol Π_{mult} (Fig. 12) incurs a communication of t elements in the preprocessing phase and $2t$ elements in 2 rounds in the online phase for multiplication when $\text{isTr} = 0$.*

Analysis: Observe that the communication towards P_{king} in steps ④ and ⑤ can be performed in parallel, resulting in the overall round complexity of the protocol being two. Further, a communication of t elements is required in step ④ and $2t$ elements are required in ⑤ (since $P_{\text{king}} \in \mathcal{E}$), thereby having a total communication complexity of $3t$ ring elements. This complexity resembles that of DN07*. However, our sharing semantics enables us to push some of the steps mentioned above to a preprocessing phase, resulting in a fast online phase, which is non-trivial to achieve in the case of DN07*. Elaborately, observe that since r, λ_a, λ_b are independent of the input (owing to our sharing semantics), computation involving these terms ranging from steps ① to ④ can thus be moved to a preprocessing phase. This improves the online communication complexity by slashing the inward communication towards P_{king} by half. Thus, the online phase requires only $2t$ ring elements of communication while offloading t elements of communication to the preprocessing phase.

Note that a straightforward extension of semi-honest multiplication of DN07* to the preprocessing model, which can be derived from [44], does not provide an efficient solution. Although such a protocol has the same online complexity ($2t$ elements) as our online phase, it has the drawback of inflating the overall communication cost by a factor of $1.6\times$ over DN07*. Elaborately, the online communication cost of $2t$ elements can be attained by appropriately defining the sharing semantics and using the P_{king} approach, similar to our protocol. However, this requires parties to generate the sharing

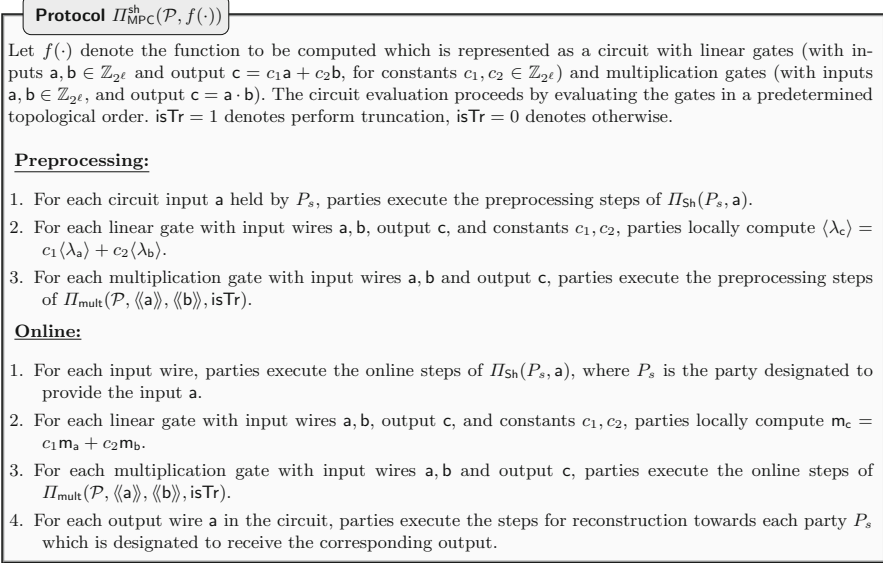


Fig. 13. Semi-honest: the complete MPC protocol.

of $\Lambda_{\mathbf{ab}} = \lambda_a \cdot \lambda_b$ from the shares of λ_a and λ_b during the preprocessing phase and requires a full-fledged multiplication, incurring a cost of $3t$ elements. This yields a protocol with a total cost of $5t$ elements in comparison with the $3t$ cost of the all-online DN07* protocol. Thus, departing from this approach, the novelty of our protocol lies in leveraging the interplay between the sharing semantics and redesigning the communication pattern among the parties to ensure that the total cost of $3t$ does not change.

Furthermore, our protocol design allows parties in \mathcal{D} to remain shut in the online phase, thereby reducing the system's operational load. This is because parties in \mathcal{D} only contribute towards the computation of \mathbf{D} , which can be completed in the preprocessing phase. However, the preprocessing phase becomes function-dependent due to the linear gates, for which the λ value for the output wires cannot be chosen randomly. Concretely, if \mathbf{C} is the output of a linear gate, say addition, with inputs \mathbf{a}, \mathbf{b} , then λ_c cannot be chosen randomly and should be defined as $\lambda_c = \lambda_a + \lambda_b$.

The complete semi-honest secure MPC protocol, Π_{MPC}^{sh} , evaluating a function $f(\cdot)$ appears in Fig. 13.

Online-only mode: Note that in instances where the function description is not known beforehand, our protocol can be run as an all-online protocol with a cost matching that of DN07*. There are two approaches in which this can be achieved. The first approach is as described in steps ①–⑥ discussed above, with the communication towards P_{king} in steps ④ and ⑤ executed simultaneously. The second approach is to begin by performing steps comprising the preprocessing phase, followed by the online phase steps. Although the second approach requires an additional round in the beginning to perform the preprocessing steps, it has the advantage that after completing the preprocessing




Fig. 14. Ideal functionality $\mathcal{F}_{\text{TrGen}}$.

phase, parties in \mathcal{D} can be shut down. This helps in improving the operational cost of the system.

Incorporating truncation To deal with decimal values that arise in several applications, including the ones considered in this work, we operate on the fixed-point arithmetic (FPA) representation [22,23], as described in Sect. 2. In this case, performing multiplication, $z = ab$, results in increasing the number of fractional bits in the result of multiplication, z , from d to $2d$. To retain FPA semantics, it is required to truncate z by d bits, i.e. compute $z^d = z/2^d$. For this, we extend the *probabilistic* truncation technique of [63,64,75] proposed in the small-party domain to the n -party setting. Given (r, r^d) -pair, with $r^d = r/2^d$, the truncated value of z can be obtained as $z^d = (z - r)^d + r^d$. Accuracy and correctness of this method follow from [73,75].

Our multiplication protocol can be modified to additionally perform truncation by incorporating the following two changes—(i) generate $\langle\langle r^d \rangle\rangle$ in step ①, and (ii) compute $\langle\langle z^d \rangle\rangle = \langle\langle (z - r)^d \rangle\rangle + \langle\langle r^d \rangle\rangle$, instead, in step ⑥. For (i), we rely on the ideal functionality, $\mathcal{F}_{\text{TrGen}}$ (Fig. 14), for computing $\langle\langle r \rangle\rangle, \langle\langle r^d \rangle\rangle$. $\mathcal{F}_{\text{TrGen}}$ can be instantiated using the appropriate MPC protocol which will be used as a black-box in our multiplication. Thus, improvements in the MPC protocol that realizes $\mathcal{F}_{\text{TrGen}}$ can be inherited in our multiplication protocol. In our work, we instantiate $\mathcal{F}_{\text{TrGen}}$ using Π_{dsBits} (Fig. 15), which is a slightly modified version of the doubly shared random bit generation protocol of [33], adapted to our n -party setting. Concretely, Π_{dsBits} generates ℓ doubly shared random bits instead of a single bit, as done in the protocol of [33]. Here, a doubly shared random bit is a bit which is arithmetic as well as Boolean shared. With respect to (ii), observe that it is a local operation, and hence performing truncation does not incur any additional overhead in the online phase. The details of Π_{dsBits} , which follows from the protocol of [33], are presented next.

Truncation—Instantiating $\mathcal{F}_{\text{TrGen}}$ We rely on a modified version of the doubly shared random bit (a bit that is arithmetic as well as Boolean shared) generation protocol of [33], extended to our n -party setting, to generate $\langle\langle r \rangle\rangle, \langle\langle r^d \rangle\rangle$ as required to perform truncation. Here, r^d represents the truncated (by d bits) version of $r \in \mathbb{Z}_{2^\ell}$. The resulting protocol is referred to as Π_{dsBits} (Fig. 15).

At a high level, generation of doubly shared bits relies on the property that every nonzero quadratic residue has exactly one root when working over fields. The work of [33], operating over rings, shows that something similar holds over rings as well. Concretely, according to lemma 4.1 of [33]: *if \mathbf{a} is such that $\mathbf{a}^2 \equiv_{\ell} 1$, then \mathbf{a} is congruent mod 2^ℓ to either $1, -1, -1 + 2^{\ell-1}, 1 + 2^{\ell-1}$.* Thus, the doubly shared bit generation protocol of [33] proceeds as follows. Generate \mathbf{a}^2 for $\mathbf{a} \in \mathbb{Z}_{2^{\ell+2}}$ such that $\mathbf{a}^2 \equiv_{\ell+2} 1$, and compute its smallest root $\mathbf{c} \bmod 2^{\ell+2}$. Compute $(\mathbf{c}^{-1}\mathbf{a})$, and by lemma 4.1 of [33] it

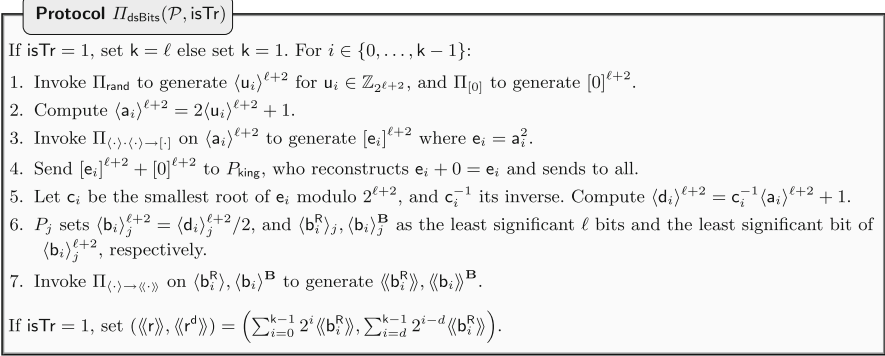


Fig. 15. Semi-honest: doubly shared bits.

follows that $c^{-1}a \in \{\pm 1, \pm 1 + 2^{\ell+1}\}$. That is, $(c^{-1}a)$ is congruent to ± 1 modulo $2^{\ell+1}$. Thus, $d = c^{-1}a + 1$ is congruent to 0 or 2 modulo $2^{\ell+1}$ with equal probability. Hence, setting $b = d/2$ outputs bit $b = 0$ or bit $b = 1$ with equal probability. Observe that the computation has to be performed over $\mathbb{Z}_{2^{\ell+2}}$. Hence, in the protocol description, we use $\ell + 2$ in the superscript to distinguish shares of x over $\mathbb{Z}_{2^{\ell+2}}$ from its shares over \mathbb{Z}_{2^ℓ} .

The main change in Π_{dsBits} from that of the protocol in [33] is that to generate $\langle\langle r \rangle\rangle, \langle\langle r^d \rangle\rangle$ Π_{dsBits} generates ℓ random doubly shared bits $b_0, \dots, b_{\ell-1} \in \mathbb{Z}_2$ instead of a single one, and composes these ℓ bits to generate r , and composes the higher $\ell - d$ bits to generate r^d , as follows.

$$\left(\langle\langle r \rangle\rangle, \langle\langle r^d \rangle\rangle \right) = \left(\sum_{i=0}^{\ell-1} 2^i \langle\langle b_i^R \rangle\rangle, \sum_{i=d}^{\ell-1} 2^{i-d} \langle\langle b_i^R \rangle\rangle \right) \tag{2}$$

Looking ahead, Π_{dsBits} can also be used only to generate a single doubly shared random bit, which finds use in other building blocks such as bit to arithmetic conversion and arithmetic to Boolean conversion. Thus, to distinguish the case when $(\langle\langle r \rangle\rangle, \langle\langle r^d \rangle\rangle)$ has to be generated versus when only a single doubly shared bit is to be generated, Π_{dsBits} takes a bit isTr as input and gives as output a doubly shared bit $\langle\langle b^R \rangle\rangle, \langle\langle b \rangle\rangle^B$ if $\text{isTr} = 0$, and $(\langle\langle r \rangle\rangle, \langle\langle r^d \rangle\rangle)$ otherwise. The protocol appears in Fig. 15.

A final thing to note is that the computation in Π_{dsBits} proceeds over secret-shared data. Thus, to generate shares of the doubly shared bit b , one should be able to divide each share of d by 2, which necessitates d and its shares to be even. This holds true since $\langle d \rangle^{\ell+2} = c^{-1} \langle a \rangle^{\ell+2} + 1 = c^{-1} (2\langle u \rangle^{\ell+2} + 1) + 1 = 2c^{-1} \langle u \rangle^{\ell+2} + c^{-1} + 1$. Here, $2c^{-1} \langle u \rangle^{\ell+2}$ is even due to multiplication by 2, while $c^{-1} + 1$ is even since c^{-1} is odd by definition.

Dot product Given $\langle\langle \cdot \rangle\rangle$ -shares of vectors \vec{x} and \vec{y} of size σ , dot product outputs $\langle\langle z \rangle\rangle$ where $z = \vec{x} \odot \vec{y} = \sum_{k=1}^{\sigma} x_k y_k$ and \odot denotes the dot product operation. The design of our multiplication protocol enables easy extension to support dot product computation

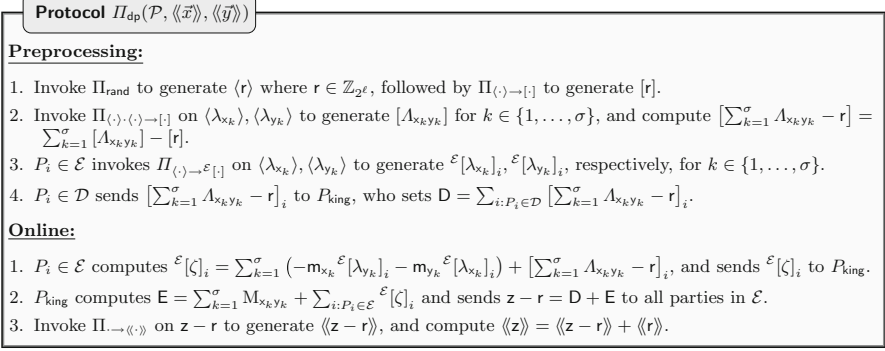


Fig. 16. Semi-honest: dot product protocol.

without incurring any overhead. Concretely, similar to multiplication,

$$\begin{aligned}
 z - r &= (\vec{x} \odot \vec{y}) - r \\
 &= \sum_{k=1}^{\sigma} M_{x_k y_k} - \sum_{k=1}^{\sigma} m_{x_k} \lambda_{y_k} - \sum_{k=1}^{\sigma} m_{y_k} \lambda_{x_k} + \sum_{k=1}^{\sigma} A_{x_k y_k} - r \quad (3)
 \end{aligned}$$

In each of the summands of $z - r$, each of the σ product terms can be generated similar to that in the multiplication protocol, which can then be locally summed up before sending it towards P_{king} . The formal protocol details appear in Fig. 16. Looking ahead, for matrix multiplication, each element of the resultant matrix can be computed via a dot product.

Multi input multiplication 3-input and 4-input multiplication protocols have showcased their wide applicability in improving the online phase complexity [64, 78, 80]. Concretely, computing $z = abc$ (3-input) or $z = abcd$ (4-input) naively requires at least two sequential invocations of 2-input multiplication protocol in the online phase. Instead, 3-input and 4-input multiplication protocol, respectively, enables performing this computation with the same online complexity as that of a *single* 2-input multiplication. Thus, we design 3-input and 4-input multiplication protocols by extending the techniques of [64, 80] to the n -party setting. Designing these protocols require modifications in the preprocessing steps. Consider 3-input multiplication (Fig. 18) where the goal is to generate $\langle\langle \cdot \rangle\rangle$ -sharing of $z = abc$ given $\langle\langle a \rangle\rangle, \langle\langle b \rangle\rangle, \langle\langle c \rangle\rangle$. Note that

Table 3. Semi-honest: communication and round complexity for multi-input multiplications .

Multiplication type	Building Block	Communication		Online Rounds
		Prep.	Online	
$z = abc$	2-input mult.	$2t\ell$	$4t\ell$	4
	3-input mult.	$6t\ell$	$2t\ell$	2
$z = abcd$	2-input mult.	$3t\ell$	$6t\ell$	4
	4-input mult.	$15t\ell$	$2t\ell$	2

$$\begin{aligned} z - r &= abc - r = (m_a - \lambda_a)(m_b - \lambda_b)(m_c - \lambda_c) - r \\ &= M_{abc} - M_{ac}\lambda_b - M_{bc}\lambda_a - M_{ab}\lambda_c + m_a\Lambda_{bc} + m_b\Lambda_{ac} + m_c\Lambda_{ab} - \Lambda_{abc} - r \end{aligned}$$

We follow an approach closely related to 2-input multiplication, with the difference being that parties additionally require to generate the additive sharing of Λ_{bc} , Λ_{ac} and Λ_{abc} during preprocessing. Given these sharings, parties proceed with a similar online phase as in Π_{mult} to compute the 3-input multiplication without inflating the online cost. Specifically, the following steps are performed in the preprocessing phase.

- For generating $\mathcal{E}[\Lambda_{ac}]$, $\mathcal{E}[\Lambda_{bc}]$ parties first compute the respective additive sharings ($[\cdot]$) using $\langle \lambda_a \rangle$, $\langle \lambda_b \rangle$ and $\langle \lambda_c \rangle$ (via two invocations of $\Pi_{(\cdot),(\cdot)} \rightarrow [\cdot]$). Following this parties in \mathcal{D} communicate their share of $[\Lambda_{ac}]$ and $[\Lambda_{bc}]$ to P_{king} , each masked with a random $[\cdot]$ -sharing of 0 (generated using $\Pi_{[0]}$). This establishes $\mathcal{E}[\Lambda_{ac}]$, $\mathcal{E}[\Lambda_{bc}]$ among parties in \mathcal{E} .
- For generating $\mathcal{E}[\Lambda_{ab}]$, a slightly different approach is taken where parties first generate $\langle \Lambda_{ab} \rangle$ using $\langle \lambda_a \rangle$, $\langle \lambda_b \rangle$ (as explained later), followed by non-interactively generating $\mathcal{E}[\Lambda_{ab}]$ (via $\Pi_{(\cdot)} \rightarrow \mathcal{T}_{[\cdot]}$). The reason for generating $\langle \Lambda_{ab} \rangle$ (instead of directly generating $\mathcal{E}[\Lambda_{ab}]$) is to facilitate generation of $\mathcal{E}[\Lambda_{abc} - r]$ from $\langle \Lambda_{ab} \rangle$, $\langle \lambda_c \rangle$ and $[r]$, which closely follows the preprocessing phase of the 2-input multiplication. Specifically, parties can generate $[\Lambda_{abc}]$ using $\Pi_{(\cdot),(\cdot)} \rightarrow [\cdot]$ on $\langle \Lambda_{ab} \rangle$, $\langle \lambda_c \rangle$, followed by parties in \mathcal{D} communicating their $[\Lambda_{abc}]$ shares masked with $[\cdot]$ -sharing of a random r to P_{king} . This generates $\mathcal{E}[\Lambda_{abc} + r]$ -sharing required during online phase.
- Regarding generation of $\langle \Lambda_{ab} \rangle$, all parties generate $\langle \cdot \rangle$ -sharing of a random $\gamma \in \mathbb{Z}_{2\ell}$ non-interactively and convert it to $[\gamma]$. Parties then compute $[\Lambda_{ab} + \gamma]$ by computing $[\Lambda_{ab}]$ from $\langle \lambda_a \rangle$, $\langle \lambda_b \rangle$ followed by summing it up with $[\gamma]$. Parties reconstruct this value towards P_{king} , who then generates $\langle \Lambda_{ab} + \gamma \rangle$, from which parties compute $\langle \Lambda_{ab} \rangle = \langle \Lambda_{ab} + \gamma \rangle - \langle \gamma \rangle$, and thereby $\mathcal{E}[\Lambda_{ab}]$ by invoking $\Pi_{(\cdot)} \rightarrow \mathcal{E}_{[\cdot]}$.

Similarly, for 4-input multiplication, parties need to generate the additive sharing of Λ_{ad} , Λ_{bd} , Λ_{cd} , Λ_{abd} , Λ_{acd} , Λ_{bcd} , Λ_{abcd} in addition to those required in the case of 3-input multiplication. Specifically, generation of $\mathcal{E}[\cdot]$ -shares (additive shares) of Λ_{ac} , Λ_{ad} , Λ_{bc} , Λ_{bd} can proceed similar to generation of $\mathcal{E}[\Lambda_{ac}]$ in $\Pi_{3\text{-mult}}$. Generation of $\mathcal{E}[\cdot]$ -shares of Λ_{ab} , Λ_{cd} is carried out by first generating its $\langle \cdot \rangle$ -shares. This enables generation of $\mathcal{E}[\cdot]$ -shares of Λ_{abc} , Λ_{abd} , Λ_{acd} , Λ_{bcd} following steps similar to




Fig. 17. 4-input multiplication .

generation of $\mathcal{E}[\Lambda_{ac}]$ in $\Pi_{3\text{-mult}}$. Finally, $\mathcal{E}[\Lambda_{abcd} - r]$ is generated similar to generating $\mathcal{E}[\Lambda_{abc} + r]$ in $\Pi_{3\text{-mult}}$. We omit formal details of 4-input multiplication protocol, $\Pi_{4\text{-mult}}$, as it is very close to $\Pi_{3\text{-mult}}$. Table 3 compares the cost of computing $z = abc$ via a 2-input multiplication sequentially vs a 3-input multiplication, and computing $z = abcd$ via a 2-input and 4-input multiplication.

The recent work of [52] provides a method to reduce the round complexity of circuit evaluation. They group the (distinct) consecutive layers in the circuit into pairs and perform a parallel evaluation of all gates in the two layers in a group. Consider a multiplication gate with inputs x, y (obtained as output from a previous layer) and output z . Their approach considers three cases: (i) If x and y are not the outputs of a multiplication gate, (ii) exactly one among x, y is the output of a multiplication gate, and (iii) both x, y are outputs of a multiplication gate. We observe that cases (ii) and (iii) in their approach resemble multi-input multiplication, which allows evaluating the second layer of multiplication ($z = x \cdot y$) non-interactively, thereby saving on rounds. For instance, consider a 2-layer sub-circuit as in Fig. 17, where $x = a \cdot b, y = c \cdot d$ are outputs of a multiplication gate which are fed as input to a multiplication gate in the next level. The approach of [52] allows computation of $z = (a \cdot b) \cdot (c \cdot d)$ in a single shot, which is equivalent to computing z via a 4-input multiplication in our case. Similarly, when only one of the inputs (either x or y) is the output of multiplication, computation of $z = x \cdot y$ resembles a 3-input multiplication. Thus, cases (i), (ii), (iii) correspond to 2-input, 3-input, and 4-input multiplication, respectively, in our work and are sufficient to reduce round complexity of any circuit evaluation by half. Hence, we restrict our focus to 3- and 4-input multiplication, although our technique can be generalized to gates with arbitrarily large fan-in.

4. Extending to Malicious Security

The ideal functionality $\mathcal{F}_{n\text{-PC}}$ for evaluating a function f in the n -party setting while providing malicious security (with fairness) appears in Fig. 19.

The input sharing and output reconstruction protocols for the malicious setting can be obtained efficiently from the semi-honest protocol following standard approaches [44,63,81]. However, the same cannot be said about multiplication. Note that although a maliciously secure multiplication protocol can be achieved by compiling our semi-honest protocol using compiler techniques such as [2,17], the resultant protocol has an expensive online phase. For instance, using the compiler of [2] yields a protocol that requires

Protocol $\Pi_{3\text{-mult}}(\mathcal{P}, \langle\langle \mathbf{a} \rangle\rangle, \langle\langle \mathbf{b} \rangle\rangle, \langle\langle \mathbf{c} \rangle\rangle)$
Preprocessing:

1. Invoke Π_{rand} to generate $\langle r \rangle$ and $\langle \gamma \rangle$ where $r, \gamma \in \mathbb{Z}_{2^t}$. Invoke $\Pi_{(\cdot) \rightarrow [\cdot]}$ to generate $[r], [\gamma]$.
2. Invoke $\Pi_{[0]}$ to generate two different $[\cdot]$ -shares of 0: $[0_1], [0_2]$.
3. Generation of $\mathcal{E}[A_{ac}], \mathcal{E}[A_{bc}]$.
 - Invoke $\Pi_{(\cdot), (\cdot) \rightarrow [\cdot]}$ on $\langle \lambda_a \rangle, \langle \lambda_c \rangle$ to generate $[A_{ac}]_i$, and compute $[A_{ac} + 0_1]_i = [A_{ac}]_i + [0_1]_i$.
 - $P_i \in \mathcal{D}$ sends $[A_{ac} + 0_1]_i$ to $P_{\text{king}} (= P_{t+1})$.
 - Analogous steps are carried out to generate $[A_{bc} + 0_2]$.
 - $P_i \in \mathcal{E} \setminus P_{t+1}$ sets $\mathcal{E}[A_{bc}]_i = [A_{bc} + 0_2]_i$ and $\mathcal{E}[A_{ac}]_i = [A_{ac} + 0_1]_i$.
 - P_{t+1} sets $\mathcal{E}[A_{bc}]_{t+1} = [A_{bc} + 0_2]_{t+1} + \sum_{i: P_i \in \mathcal{D}} [A_{bc} + 0_2]_i$ and $\mathcal{E}[A_{ac}]_{t+1} = [A_{ac} + 0_1]_{t+1} + \sum_{i: P_i \in \mathcal{D}} [A_{ac} + 0_1]_i$.
4. Generation of $\mathcal{E}[A_{ab}]$.
 - Invoke $\Pi_{(\cdot), (\cdot) \rightarrow [\cdot]}$ on $\langle \lambda_a \rangle, \langle \lambda_b \rangle$ to generate $[A_{ab}]_i$, set $[A_{ab} + \gamma]_i = [A_{ab}]_i + [\gamma]_i$, and send $[A_{ab} + \gamma]_i$ to P_{king} .
 - P_{king} reconstructs $A_{ab} + \gamma$, and sends $A_{ab} + \gamma$ to $P_i \in \mathcal{E}$. Parties non-interactively generate $\langle A_{ab} + \gamma \rangle$ via $\Pi_{\rightarrow \langle \cdot \rangle}$ and $\Pi_{\langle \cdot \rangle \rightarrow (\cdot)}$.
 - Compute $\langle A_{ab} \rangle = \langle A_{ab} + \gamma \rangle - \langle \gamma \rangle$ and invoke $\Pi_{(\cdot) \rightarrow \mathcal{E}[\cdot]}$ on $\langle A_{ab} \rangle$ to generate $\mathcal{E}[A_{ab}]$.
5. Generation of $\mathcal{E}[A_{abc} + r]$.
 - Invoke $\Pi_{(\cdot), (\cdot) \rightarrow [\cdot]}$ on $\langle A_{ab} \rangle, \langle \lambda_c \rangle$ to generate $[A_{abc}]_i$, and compute $[A_{abc} + r]_i = [A_{abc}]_i + [r]_i$.
 - $P_i \in \mathcal{D}$ sends $[A_{abc} + r]_i$ to P_{king} .
 - $P_i \in \mathcal{E} \setminus P_{t+1}$ sets $\mathcal{E}[A_{abc} + r]_i = [A_{abc} + r]_i$.
 - P_{t+1} sets $\mathcal{E}[A_{abc} + r]_{t+1} = [A_{abc} + r]_{t+1} + \sum_{i: P_i \in \mathcal{D}} [A_{abc} + r]_i$.
6. $P_i \in \mathcal{E}$ invoke $\Pi_{(\cdot) \rightarrow \mathcal{E}[\cdot]}$ on $\langle \lambda_a \rangle, \langle \lambda_b \rangle$ and $\langle \lambda_c \rangle$ to generate $\mathcal{E}[\lambda_a]_i, \mathcal{E}[\lambda_b]_i, \mathcal{E}[\lambda_c]_i$, respectively.

Online:

1. $P_i \in \mathcal{E}$ computes and sends $\mathcal{E}[\zeta]_i = -M_{ac} \mathcal{E}[\lambda_b]_i - M_{bc} \mathcal{E}[\lambda_a]_i - M_{ab} \mathcal{E}[\lambda_c]_i + m_a \mathcal{E}[A_{bc}]_i + m_b \mathcal{E}[A_{ac}]_i + m_c \mathcal{E}[A_{ab}]_i - \mathcal{E}[A_{abc} + r]_i$ to P_{king} .
2. P_{king} computes and sends $z - r = M_{abc} + \sum_{i: P_i \in \mathcal{E}} \mathcal{E}[\zeta]_i$ to $P_i \in \mathcal{E}$.
3. Invoke $\Pi_{\rightarrow \langle \cdot \rangle}$ on $z - r$ to generate $\langle z - r \rangle$, and compute $\langle z \rangle = \langle (z - r) \rangle + \langle r \rangle$.

Fig. 18. Semi-honest: 3-input multiplication protocol.

computation over *extended* rings and communicating $4t$ extended ring elements in the online phase. This is not favourable compared to working over plain rings, especially in the online phase. Further, compilers such as those in [17] require heavy computational machinery like reliance on zero-knowledge proofs in the online phase, which is also not desirable. Thus, to attain a computation and communication efficient online phase, departing from the aforementioned compiler-based approaches, we design a maliciously secure multiplication protocol that requires communicating $3t$ ring elements in each phase. It is worth noting that we can do this while retaining the benefits of requiring only $t + 1$ parties in the online phase (for most of the computation). The remaining t parties are required to come online only for a short one-time verification phase, that is deferred to the end of the computation. Deferring verification may result in a privacy breach [53]. However, we describe later why the privacy breach does not arise in our protocol. With this background, in what follows next, we begin with describing the input sharing and output reconstruction protocols and then focus on discussing the challenges encountered and their resolutions for obtaining a maliciously secure *multiplication* protocol.

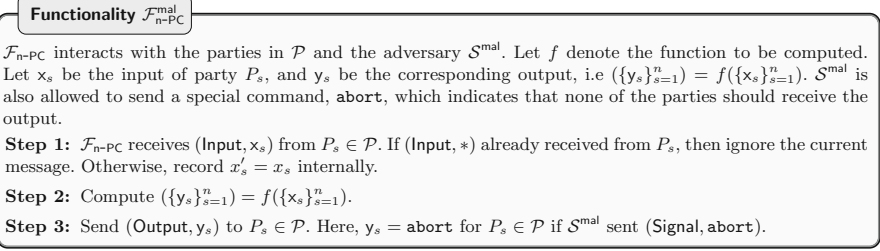


Fig. 19. Malicious: ideal functionality for evaluating function f with fairness.




Fig. 20. Malicious: input sharing protocol.

Input sharing This protocol is similar to the semi-honest one, where to enable P_s to generate $\langle\langle a \rangle\rangle$, parties generate (λ_a) such that P_s learns λ_a , followed by P_s sending the masked value $m_a = a + \lambda_a$ to all. However, note that a corrupt P_s can cause inconsistency among the honest parties by sending different masked values. To ensure the same value is received by all, parties perform a hash-based consistency check, denoted by Π_{agree} (Sect. 2), where each party sends a hash of the received masked value(s) to every other party and **aborts** if it receives inconsistent hashes. Note that this check for all the inputs can be combined, thereby amortizing the cost. The formal protocol appears in Fig. 20.

Reconstruction To reconstruct $\langle\langle \cdot \rangle\rangle$ -shared value a towards $P_s \in \mathcal{P}$, observe that each share that P_s misses is held by $t + 1$ other parties. Each of these parties sends the missing share to P_s . If the received values for a share are consistent, P_s uses this value to perform reconstruction, and **aborts** otherwise. As an optimization, one party can send the missing share while reconstructing several values, and t others can send its hash.

Fairness is a stronger security notion than security with abort, where, during reconstruction, either all parties learn the output or none do. For fair reconstruction, we extend the techniques in [81] to the n -party setting, where commitments are generated on each share of the mask of the output Z (required to reconstruct Z) by $t + 1$ parties in the preprocessing phase.

During the online phase, these commitments are opened towards the respective parties if all the parties are alive (did not **abort**). Since each share of the mask is held by $t + 1$ parties and there is at least one honest party among every set of $t + 1$ parties, it is guaranteed that parties will obtain the correct opening for the commitment of the missing share from the honest party, and all honest parties can reconstruct the output. Else, if the adversary misbehaved at some step during the protocol, none of the honest parties will share the opening information and none will obtain the output. Note that to determine if all parties are alive, each party broadcasts a bit **alive** = 1, where the

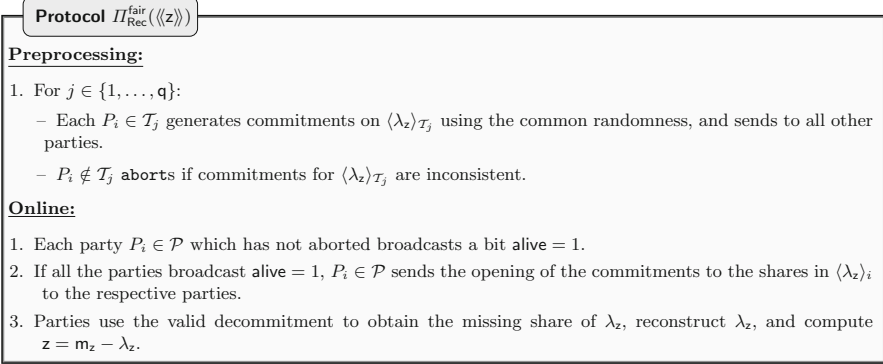


Fig. 21. Malicious: fair reconstruction protocol.

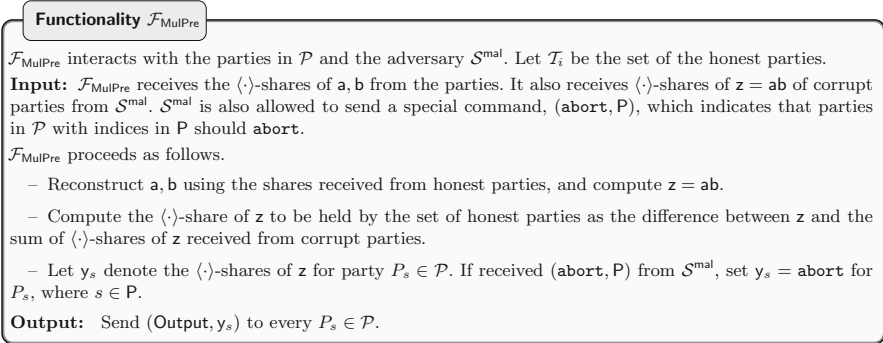


Fig. 22. Ideal functionality $\mathcal{F}_{\text{MulPre}}$.

broadcast can be realized on point-to-point channels using a broadcast protocol such as that of [41]. The formal protocol $\Pi_{\text{Rec}}^{\text{fair}}$ appears in Fig. 21.

Multiplication To enable generation of $\langle\langle \mathbf{z} \rangle\rangle = \langle\langle \mathbf{a}\mathbf{b} \rangle\rangle$ from $\langle\langle \mathbf{a} \rangle\rangle$ and $\langle\langle \mathbf{b} \rangle\rangle$, we retain the high-level ideas from the semi-honest protocol. Our task reduces to (i) generating additive shares of $\Lambda_{\mathbf{a}\mathbf{b}}$ among parties in \mathcal{E} (i.e. $\mathcal{E}[\Lambda_{\mathbf{a}\mathbf{b}}]$) given $\langle \lambda_{\mathbf{a}} \rangle$ and $\langle \lambda_{\mathbf{b}} \rangle$, in the preprocessing phase, and (ii) reconstructing $\mathbf{z} - r$ in the online phase. Given (i), computing $\mathcal{E}[\mathbf{z} - r]$ in the online phase is a local operation. Given (ii), parties can invoke $\Pi_{\rightarrow \langle \cdot \rangle}$ to generate $\langle\langle \mathbf{z} - r \rangle\rangle$ and compute $\langle\langle \mathbf{z} \rangle\rangle = \langle\langle \mathbf{z} - r \rangle\rangle + \langle\langle r \rangle\rangle$, where $\langle\langle r \rangle\rangle$ is generated in the preprocessing phase, as discussed in the semi-honest case.

For task (i), our idea for the semi-honest case, of making parties in \mathcal{D} send their shares to P_{king} , does not work in the presence of a malicious adversary. To address this, we make black-box use of a maliciously secure multiplication protocol, abstracted as a functionality $\mathcal{F}_{\text{MulPre}}$ in Fig. 22, that computes $\langle \Lambda_{\mathbf{a}\mathbf{b}} \rangle$ from $\langle \lambda_{\mathbf{a}} \rangle, \langle \lambda_{\mathbf{b}} \rangle$. In this work, we instantiate $\mathcal{F}_{\text{MulPre}}$ with the state-of-the-art multiplication protocol of [17] that provides abort security and requires $3t$ elements of (amortized) communication. Note that although the protocol of [17] relies on zero-knowledge proofs, this computation is carried out in the preprocessing phase of our multiplication protocol. Moreover, since

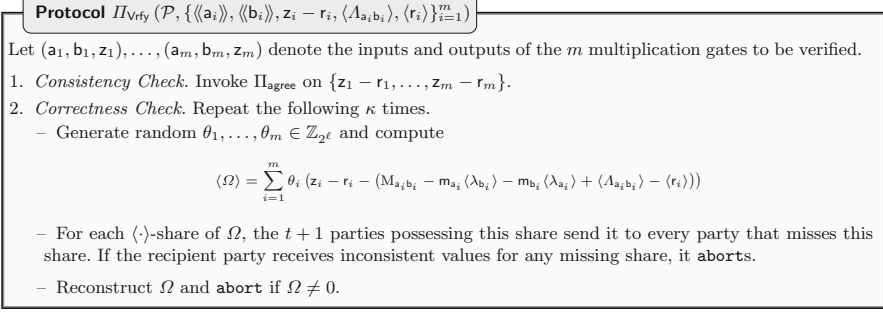


Fig. 23. Malicious: verification protocol for all multiplication gates.

preprocessing is done for many instances in one shot, the zero-knowledge proof can benefit from amortization. The parties then invoke $\Pi_{\langle \cdot \rangle \rightarrow \mathcal{E}[\cdot]}$ to obtain $\mathcal{E}[\Lambda_{\text{ab}}]$ from $\langle \Lambda_{\text{ab}} \rangle$. Looking ahead, $\langle \Lambda_{\text{ab}} \rangle$ also aids in performing the online verification check.

For task (ii), in the online phase, we retain the idea of parties in \mathcal{E} optimistically reconstructing $\mathbf{z} - \mathbf{r}$ from their additive shares ($\mathcal{E}[\cdot]$ -share) to ensure that only the parties in \mathcal{E} remain active for most of the computation. Moreover, this optimistic reconstruction requires only $\mathcal{O}(t)$ -element communication rather than the $\mathcal{O}(t^2)$ required for reconstruction from $\langle \cdot \rangle$ -shares (which is what will be used later for performing verification, albeit to perform only one such reconstruction). Thus, similar to the semi-honest protocol, parties in \mathcal{E} optimistically reconstruct $\mathbf{z} - \mathbf{r}$ towards P_{king} , who further sends the reconstructed value to the parties in \mathcal{E} . In the malicious setting, this approach requires additional care since a malicious party may send a wrong $\mathcal{E}[\cdot]$ -share of $\mathbf{z} - \mathbf{r}$ to P_{king} or a malicious P_{king} may send an incorrectly reconstructed (inconsistent) $\mathbf{z} - \mathbf{r}$ to the parties. To account for these behaviours, the protocol is augmented with a short one-off verification phase to verify the consistency and correctness of $\mathbf{z} - \mathbf{r}$. This phase is executed in the end of the protocol and requires the presence of *all* parties, and hence the possession of $\mathbf{z} - \mathbf{r}$ by all. This is in contrast to the semi-honest protocol where $\mathbf{z} - \mathbf{r}$ is given to only parties in \mathcal{E} . To keep \mathcal{D} disengaged for most of the online phase, sending $\mathbf{z} - \mathbf{r}$ to them is deferred till the end of the protocol. This send is a one-off and can be combined for all multiplication gates. Details of verification protocol Π_{Verify} (Fig. 23) are given next.

Verification comprises two checks—a *consistency* check to first verify that P_{king} has indeed sent the same $\mathbf{z} - \mathbf{r}$ to all the parties, followed by a *correctness* check to verify the correctness of the $\mathbf{z} - \mathbf{r}$. For the former, parties perform a hash-based consistency check of $\mathbf{z} - \mathbf{r}$ and abort in case of any inconsistency. If $\mathbf{z} - \mathbf{r}$ is consistent, parties verify its correctness. The high-level idea for verifying correctness is to *robustly* reconstruct $\mathbf{z} - \mathbf{r}$, but now from its $\langle \cdot \rangle$ -shares (can be computed given $\langle \lambda_{\mathbf{a}} \rangle, \langle \lambda_{\mathbf{b}} \rangle, \langle \Lambda_{\text{ab}} \rangle$ that are generated in the preprocessing phase). Parties can then verify if this reconstructed value equals the value received from P_{king} . Concretely, this is equivalent to robustly reconstructing $\langle \Omega \rangle = \langle \mathbf{z} - \mathbf{r} - (M_{\text{ab}} - m_{\mathbf{a}} \lambda_{\mathbf{b}} - m_{\mathbf{b}} \lambda_{\mathbf{a}} + \Lambda_{\text{ab}} - \mathbf{r}) \rangle$, where $\mathbf{z} - \mathbf{r}$ is the value received from P_{king} , and verifying if $\Omega = 0$. For robust reconstruction of $\langle \Omega \rangle$, every party sends its $\langle \cdot \rangle$ -share to every other party who misses this share, and **aborts** in case of inconsistencies

in the received values. Elaborately, reconstruction of Ω towards $P_s \in \mathcal{P}$ proceeds as follows. For each missing $\langle \cdot \rangle$ -share of Ω at P_s , each of the $t + 1$ parties holding this share sends it to P_s . P_s uses this share for reconstruction if all the $t + 1$ received values are consistent, else it aborts. The presence of at least one honest party among the $t + 1$ guarantees that inconsistency, if any, can be detected. Since each share in $\langle \Omega \rangle$ is held by $t + 1$ parties, comprising at least one honest party, any cheating by up to t corrupt parties is guaranteed to be detected. Since reconstruction should happen towards at least $t + 1$ parties, communicating a missing share towards all these $t + 1$ parties requires $\mathcal{O}(t^2)$ communication in total, and there are $m = \binom{n}{h} - \binom{n-1}{h-1}$ such missing shares. Note that the cost of this reconstruction can be optimized using standard optimization techniques [2, 27], where the correctness of $\mathbf{z} - \mathbf{r}$ for several multiplication gates can be verified with a single reconstruction by reconstructing a linear combination of Ω for several gates and verifying equality with 0. Thus, only one robust reconstruction from $\langle \cdot \rangle$ -shares is required for several multiplication gates, whose cost gets amortized due to verification across multiple gates.

It is worth noting that this random linear combination technique does not trivially work over rings. This is due to the existence of zero divisors which results in the linear combination being 0 with a probability $1/2$ (which denotes the cheating probability of the adversary) [2]. Hence, to obtain the desired security, the verification check is repeated κ times where κ is the security parameter. This bounds the cheating probability of adversary to $1/2^\kappa$. Another approach is to perform the verification over extended rings [15, 16]. Specifically, verification operations are carried out over a ring $\mathbb{Z}_{2^\ell}/f(x)$ which is a ring of all polynomials with coefficients in \mathbb{Z}_{2^ℓ} modulo a degree d polynomial $f(x)$ that is irreducible over \mathbb{Z}_2 . Each element of \mathbb{Z}_{2^ℓ} is lifted to a degree d polynomial in $\mathbb{Z}_{2^\ell}[x]/f(x)$, which increases the communication required to perform verification by a factor of d .

The maliciously secure multiplication protocol (see Fig. 24) can be broken down into the following:

- Preprocessing phase which involves generation of $\langle \Lambda_{ab} \rangle$ by invoking $\mathcal{F}_{\text{MulPre}}$. Malicious behaviour, if any, will be caught by $\mathcal{F}_{\text{MulPre}}$. $\langle \Lambda_{ab} \rangle$ is non-interactively converted into $\mathcal{E}[\cdot]$ -shares of λ_{ab} . $\mathcal{E}[\lambda_a]$, $\mathcal{E}[\lambda_b]$ is also generated non-interactively.
- Generation of $\mathcal{E}[\cdot]$ -shares of λ_a , λ_b , Λ_{ab} during preprocessing enables computation of $\mathcal{E}[\mathbf{z} - \mathbf{r}]$ in the online phase, and thereby reconstruction of $\mathbf{z} - \mathbf{r}$ via P_{King} . The crucial point to note here is that this requires the presence of only parties in \mathcal{E} in the online phase. This is followed by non-interactive generation of $\langle \mathbf{z} - \mathbf{r} \rangle$ from which $\langle \mathbf{z} \rangle$ is computed as $\langle \mathbf{z} \rangle = \langle \mathbf{z} - \mathbf{r} \rangle + \langle \mathbf{r} \rangle$, where $\langle \mathbf{r} \rangle$ is generated during preprocessing.
- Finally, to catch malicious behaviour in the online phase, if any, in the verification phase the correctness of the generated $\langle \mathbf{z} \rangle$ is checked simultaneously, for each \mathbf{z} that is the output of a multiplication gate. This is done by invoking Π_{Verify} . Note that before this verification begins, P_{King} sends $\mathbf{z} - \mathbf{r}$ corresponding to all multiplication gates to parties in \mathcal{D} in a single shot.

As pointed out in [53], deferring the correctness check to later may result in a privacy breach when using a sharing scheme that allows for redundancy (such as RSS or Shamir sharing). We next discuss this breach and explain how it is overcome in our case. We begin

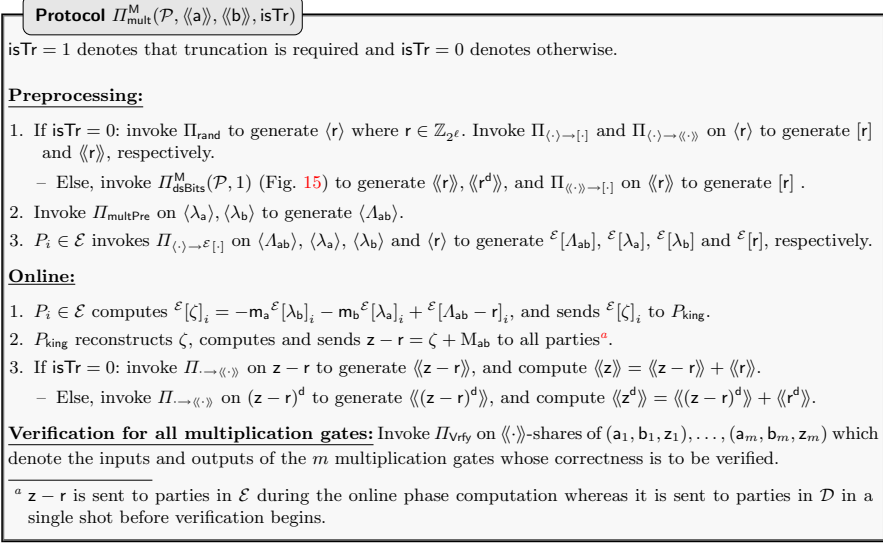


Fig. 24. Malicious: multiplication protocol.

with explaining the attack that a malicious adversary can launch if reconstruction towards P_{king} is performed by relying on RSS (or Shamir sharing), naively. Consider a circuit with two sequential multiplication gates with the output of the first gate, say \mathbf{a} , going as input to the second gate. Let \mathbf{b} denote the other input to the second multiplication gate and \mathbf{z} denote its output. In a P_{king} -based approach for multiplication, t parties send their respective (RSS/Shamir) share of a masked value to P_{king} . In particular, for the first multiplication gate in the circuit mentioned above, t parties send their corresponding share of $\mathbf{a} - r_a$ to P_{king} , who reconstructs it and sends it back to all. Delaying the verification allows a malicious P_{king} to send an inconsistent value of $\mathbf{a} - r_a$ to the parties, using which it can learn the private input \mathbf{b} , as follows. Suppose P_{king} sends the correct $\mathbf{a} - r_a$ to all but one out of the remaining t online parties, to which it sends $\mathbf{a} - r_a + \delta$. Owing to this, for the next multiplication gate P_{king} receives the shares of $\mathbf{z} - r_z$ from the former $t - 1$ parties and a share of $(\mathbf{a} + \delta)\mathbf{b} - r_z = \mathbf{z} + \delta\mathbf{b} - r_z$ from the latter party. Having obtained these and additionally using the shares of $\mathbf{z} - r_z$ and $\mathbf{z} + \delta\mathbf{b} - r_z$ corresponding to the t corrupt parties including itself, a malicious P_{king} can reconstruct $\mathbf{z} - r_z$ as well as $\mathbf{z} + \delta\mathbf{b} - r_z$, thus learning \mathbf{b} in clear. The crux of this attack lies in the fact that a malicious adversary corrupting t parties including P_{king} already possesses t shares each of $\mathbf{z} - r_z$ and $\mathbf{z} + \delta\mathbf{b} - r_z$. Thus, an additional share of these obtained from the online parties allows it to carry out the attack successfully. However, this attack does not hold when working with additive ($\mathcal{E}[\cdot]$) sharing, which is what prevents our protocol from falling prey to this attack.

Elaborately, recall that in our protocol, during reconstruction towards P_{king} , any redundancy due to $\langle\langle \cdot \rangle\rangle$ -sharing is eliminated with parties switching to $\mathcal{E}[\cdot]$ -sharing (additive sharing among parties in \mathcal{E}). Due to this, even if P_{king} sends inconsistent values to the parties, the $\mathcal{E}[\cdot]$ -share of $\mathbf{z} - r_z$ or $\mathbf{z} + \delta\mathbf{b} - r_z$ that it receives corresponds to an additive share defined with respect to parties in \mathcal{E} . Hence, this additionally received additive

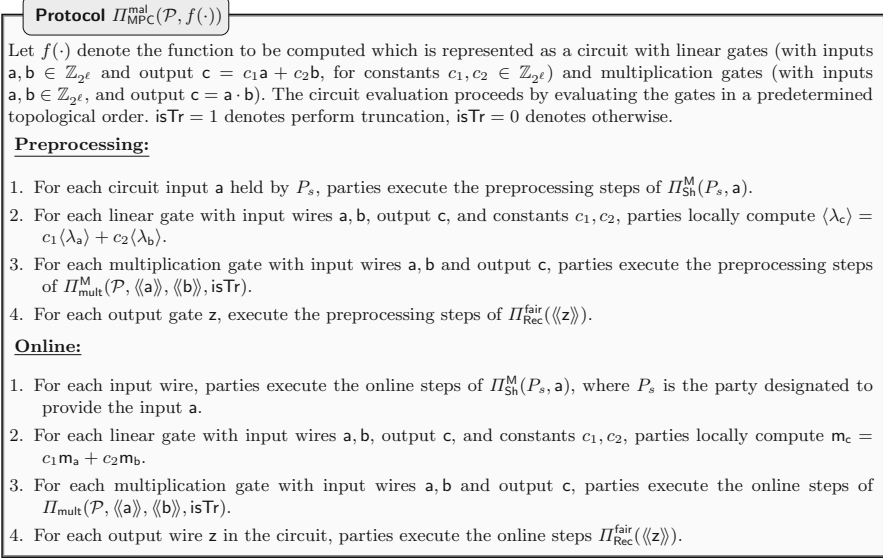


Fig. 25. Malicious: the complete MPC protocol.

share cannot be combined with the shares held by the t corrupt parties to perform the reconstruction. Thus, the earlier strategy of P_{king} of using these additional shares in conjunction with the t corrupt shares to reconstruct $\mathbf{z} - r_z$ and $\mathbf{z} + \delta\mathbf{b} - r_z$ does not hold. The primary reason which prevents the attack is the elimination of redundancy in the sharing scheme by switching to $(t + 1)$ -out-of- $(t + 1)$ additive sharing ($\mathcal{E}[\cdot]$ -sharing) for the set of parties in \mathcal{E} , which is known to withstand this attack [53]. However, this privacy breach persists in the protocol of [44].

Discussion about [44]: The above attack can be circumvented by making P_{king} broadcast the reconstructed value to all the parties, as discussed in [44]. To further optimize the protocol by requiring only $t + 1$ parties to be active in the online phase, they rely on broadcast with abort, which comprises two phases—(i) *send*: where P_{king} sends the value to the recipients, and (ii) *verification*: where the recipients exchange hash of the received value among themselves and abort in case of inconsistency. However, for amortization, they defer the verification (even with respect to broadcast) towards the end of the protocol, thus making their protocol susceptible to the aforementioned attack. We observe that one fix is to perform the verification with respect to broadcast after each level in the circuit. This, however, requires all the parties to be online. An optimization to let only the $t + 1$ parties in the online phase to perform this verification after each level, thereby allowing the remaining t parties to be shut off. Specifically, this involves performing *verification* where the online parties exchange the hash of the received value and abort in case of inconsistency. When the remainder t (offline) parties come online towards the end of the protocol for verifying the correctness of the multiplication gates, this verification should be preceded by first verifying the consistency of the values broadcast by P_{king} to the offline parties (and involves participation of all n parties). Since the

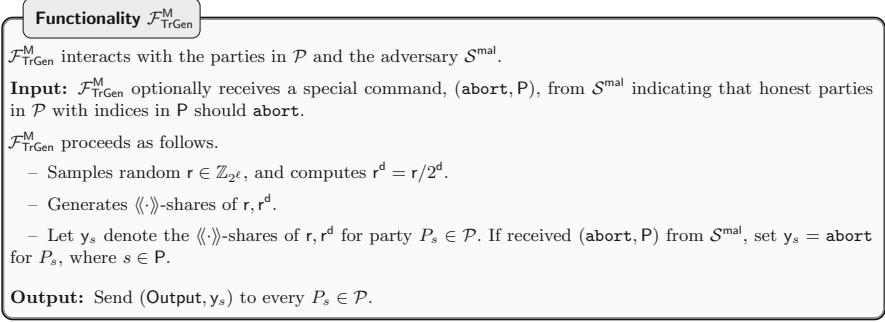


Fig. 26. Ideal functionality $\mathcal{F}_{\text{TrGen}}^M$.

online phase involves broadcasting the reconstructed value to t other online parties, this amounts to an exchange of $\mathcal{O}(t^2)$ hashes after each level, thereby incurring a circuit depth-dependent overhead in the communication cost as well as the rounds. In order for the communication cost to get amortized, it is required that the circuit has $\mathcal{O}(t^2)$ gates at each level. However, the overhead in terms of number of rounds persists.

Lemma 2. *Protocol Π_{mult}^M (Fig. 24) incurs a communication of $3t$ elements in the preprocessing phase and $3t$ elements in 2 rounds in the online phase for multiplication when $\text{isTr} = 0$.*

The complete MPC protocol The maliciously secure MPC protocol, $\Pi_{\text{MPC}}^{\text{mal}}$, evaluating a function $f(\cdot)$ appears in Fig. 25.

Multiplication with truncation Similar to the semi-honest protocol, truncation can be incorporated in the malicious multiplication as well without inflating the online communication. For this, we rely on maliciously secure ideal functionality, $\mathcal{F}_{\text{TrGen}}^M$ (Fig. 26), to generate the $\langle\langle \cdot \rangle\rangle$ -shares of (r, r^d) .

$\mathcal{F}_{\text{TrGen}}^M$ (Fig. 26) can be realized using the maliciously secure variant of Π_{dsBits} (Fig. 15), denoted as Π_{dsBits}^M [33]. This protocol is similar to the semi-honest protocol except with the following differences to account for malicious behaviour. The $\langle \cdot \rangle$ -shares of $\mathbf{e}_i = \mathbf{a}^2$ are generated by invoking Π_{multPre} instead of relying on $\Pi_{\langle \cdot \rangle, \langle \cdot \rangle \rightarrow [\cdot]}$. This ensures generation of correct $\langle \cdot \rangle$ -shares of \mathbf{e}_i , and malicious behaviour, if any, will lead to an abort. Following this, \mathbf{e}_i is either correctly reconstructed towards all or parties abort. This ensures that an adversary cannot lead to reconstruction of an incorrect \mathbf{e}_i . Concretely, for reconstruction, similar to multiplication, every party sends its $\langle \cdot \rangle$ -share to every other party, and aborts in case of inconsistencies in the received values.⁵ The rest of the protocol steps (which are non-interactive) remain unchanged, and hence, a formal protocol is omitted.

⁵This can be optimized similar to the online phase of the multiplication protocol, where the value is first reconstructed towards P_{king} who sends the reconstructed value to all, followed by verifying its correctness via the verification check.

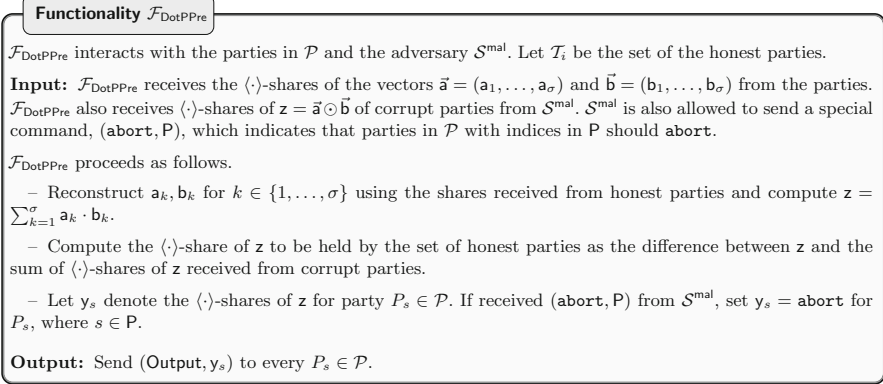


Fig. 27. Ideal functionality for Π_{dotPre} .

Dot product To generate $\langle\langle z \rangle\rangle$ for $z = \vec{x} \odot \vec{y}$ where \vec{x} and \vec{y} are vectors of size σ and are $\langle \cdot \rangle$ -shared, protocol $\Pi_{\text{dp}}^{\text{M}}$ proceeds similar to the semi-honest variant Π_{dp} (Fig. 16). During the preprocessing phase, parties in \mathcal{E} obtain $\mathcal{E}[\cdot]$ -shares of $\Lambda_{\vec{x} \odot \vec{y}} = \sum_{k=1}^{\sigma} \lambda_{x_k} \lambda_{y_k}$ and $\lambda_{x_k}, \lambda_{y_k}$ for $k \in \{1, \dots, \sigma\}$. Although the latter two can be computed by parties locally with an invocation of $\Pi_{\langle \cdot \rangle \rightarrow \mathcal{E}[\cdot]}$ (Fig. 7), computation of the former differs significantly from the semi-honest protocol. For this, we extend the ideas from SWIFT [63] and generate $\langle \Lambda_{\vec{x} \odot \vec{y}} \rangle$, by executing a maliciously secure dot product protocol Π_{dotPre} over $\langle \cdot \rangle$ -shares (abstracted as a functionality $\mathcal{F}_{\text{DotPPre}}$ in Fig. 27). Specifically, parties invoke Π_{dotPre} on $\langle \cdot \rangle$ -shares of $\vec{\lambda}_x = (\lambda_{x_1}, \dots, \lambda_{x_\sigma})$ and $\vec{\lambda}_y = (\lambda_{y_1}, \dots, \lambda_{y_\sigma})$ to compute $\langle \Lambda_{\vec{x} \odot \vec{y}} \rangle$, followed by an invocation of $\Pi_{\langle \cdot \rangle \rightarrow \mathcal{E}[\cdot]}$ to obtain $\mathcal{E}[\Lambda_{\vec{x} \odot \vec{y}}]$. Having computed the necessary preprocessing data, the online phase proceeds similarly to the semi-honest protocol Π_{dp} (Fig. 16), where parties reconstruct $z - r$ via P_{king} as per equation (3). To account for misbehaviour, the protocol is augmented with a verification phase similar to that in malicious multiplication.

Observe that a trivial realization of $\mathcal{F}_{\text{DotPPre}}$ can be reduced to σ instances of multiplication. However, we extend the ideas from [16, 17, 63] and rely on a distributed zero-knowledge proof [17] to eliminate the vector-size dependency in the preprocessing phase. Concretely, we instantiate $\mathcal{F}_{\text{DotPPre}}$ using a semi-honest dot product protocol [54] whose cost matches that of semi-honest multiplication [35] (and thus is independent of the vector-size), followed by a verification phase to verify the correctness of the dot product computation. For the verification, we extend the verification technique for multiplication in [17], to now verify the correctness of the dot product, such that the cost due to verification can be amortized away for multiple dot products, thereby resulting in vector-size independent preprocessing.

Elaborately, the semi-honest dot product protocol in [54] takes as input $\langle \vec{x} \rangle, \langle \vec{y} \rangle$ where \vec{x}, \vec{y} are vectors of size σ , and outputs $\langle z \rangle = \langle \vec{x} \odot \vec{y} \rangle$. For this, parties invoke $\Pi_{\langle \cdot \rangle, \langle \cdot \rangle \rightarrow [\cdot]}$ on each element in \vec{x}, \vec{y} and sum these up to generate $[\rho] = [\vec{x} \odot \vec{y}]$. These shares are randomized by summing with $[r]$ (converted from $\langle r \rangle$) for a random r , and the sum $z + r = (\vec{x} \odot \vec{y}) + r$ is reconstructed towards P_{king} , who sends the reconstructed $z + r$ to

parties in \mathcal{E} . All parties then non-interactively generate $\langle \mathbf{z} + \mathbf{r} \rangle$ by setting one of its share as $\mathbf{z} + \mathbf{r}$ and the others as 0. Given $\langle \mathbf{z} + \mathbf{r} \rangle$, $\langle \mathbf{r} \rangle$, parties can compute $\langle \mathbf{z} \rangle = \langle \mathbf{z} + \mathbf{r} \rangle - \langle \mathbf{r} \rangle$. Observe that communication of $\langle \mathbf{z} + \mathbf{r} \rangle$ to P_{king} requires $2t$ elements, while communicating $\mathbf{z} + \mathbf{r}$ to parties in \mathcal{E} requires t elements, resulting in a matching cost of $3t$ elements as that required for semi-honest multiplication [35]. The correctness of m dot product triples $(\vec{x}_1, \vec{y}_1, \mathbf{z}_1), \dots, (\vec{x}_m, \vec{y}_m, \mathbf{z}_m)$, can be verified by taking a random linear combination,

$$\beta = \sum_{k=1}^m \theta_k \cdot \left(\mathbf{z}_k - \sum_{j=1}^{\sigma} \mathbf{x}_{kj} \cdot \mathbf{y}_{kj} \right)$$

where $\{\theta_k\}_{k=1}^m$ is randomly chosen by all the parties and checking if $\beta = 0$. Given $\langle \cdot \rangle$ -shares of $\vec{x}_k, \vec{y}_k, \mathbf{z}_k$ for $k \in \{1, \dots, m\}$, parties can compute an additive share ($[\cdot]$ -share) of β by invoking $\Pi_{\langle \cdot \rangle, \langle \cdot \rangle \rightarrow [\cdot]}$. However, since $[\cdot]$ -sharing does not allow for robust reconstruction, the approach is to generate $\langle \beta \rangle$ and then robustly reconstruct it and check equality with 0. To generate $\langle \beta \rangle$, parties first $\langle \cdot \rangle$ -share (via $\Pi_{\langle \cdot \rangle}$, Sect. 2) their $[\cdot]$ -share of

$$\psi = \sum_{k=1}^m \theta_k \cdot \sum_{j=1}^{\sigma} \mathbf{x}_{kj} \cdot \mathbf{y}_{kj}.$$

Let ψ^i denote the $[\cdot]$ -share of ψ held by P_i . Given $\langle \psi^i \rangle$ for $i \in \{1, \dots, n\}$, parties can compute

$$\langle \beta \rangle = \sum_{k=1}^m \theta_k \cdot \langle \mathbf{z}_k \rangle - \sum_{i=1}^n \langle \psi^i \rangle$$

and reconstruct β . It is, however, required to ensure that every party P_i $\langle \cdot \rangle$ -shares the correct ψ^i . To check the correctness of ψ^i , parties need to verify if

$$\psi^i - \sum_{k=1}^m \theta_k \left(\sum_{j=1}^{\sigma} \mathbf{x}_{kj}^i \cdot \mathbf{y}_{kj}^i \right) = 0 \quad (4)$$

where $\mathbf{x}_{kj}^i, \mathbf{y}_{kj}^i$ denote the $\langle \cdot \rangle$ -share of $\mathbf{x}_{kj}, \mathbf{y}_{kj}$ held by P_i . Note that following along the lines of $\Pi_{\langle \cdot \rangle \rightarrow \langle \cdot \rangle}$, parties can generate these $\langle \cdot \rangle$ -share of $\mathbf{x}_{kj}^i, \mathbf{y}_{kj}^i$ from $\langle \cdot \rangle$ -shares of $\mathbf{x}_{kj}, \mathbf{y}_{kj}$, non-interactively. Now, setting $\mathbf{a}_{kj} = \theta_k \mathbf{x}_{kj}^i, \mathbf{b}_{kj} = \mathbf{y}_{kj}^i, \mathbf{c} = \psi^i$, for $k \in \{1, \dots, m\}$, Eq. (4), can re-written as

$$\mathbf{c} - \sum_{k=1}^m \sum_{j=1}^{\sigma} \mathbf{a}_{kj} \mathbf{b}_{kj} = 0 \implies \mathbf{c} - \sum_{l=1}^{m\sigma} \tilde{\mathbf{a}}_l \tilde{\mathbf{b}}_l = 0 \quad (5)$$

The correctness of Eq. (5) can be verified by invoking $\mathcal{F}_{\text{proveDeg2Rel}}^{\text{abort}}$ (see section 3 of [17] for the definition and its instantiation), which takes as input $\langle \cdot \rangle$ -shares of $\tilde{\mathbf{a}}_l, \tilde{\mathbf{b}}_l, \mathbf{c}$

for $l \in \{1, \dots, m\sigma\}$, which are known in clear to party P_i , and verifies if Eq. (5) holds. The protocol realizing $\mathcal{F}_{\text{proveDeg2Rel}}^{\text{abort}}$ for all n parties requires communicating $\mathcal{O}(n \log(m\sigma) + n)$ extended ring elements per party. Further, since steps other than $\mathcal{F}_{\text{proveDeg2Rel}}^{\text{abort}}$ require sharing and reconstructing one element, it adds a small constant cost, resulting in the communication cost for verifying m dot products for vector size σ being $\mathcal{O}(n \log(m\sigma) + n)$ extended ring elements per party.

Multi input multiplication This protocol is similar to its semi-honest counterpart with the difference that the preprocessing phase relies on invoking $\mathcal{F}_{\text{MulPre}}$ for generating the required multiplicative terms. At a high level, the malicious variant of multi-input multiplication protocol can be viewed as an amalgamation of the semi-honest multi-input multiplication and the malicious multiplication protocol. For the case of 3-input multiplication, recall that the semi-honest protocol to compute $\langle\langle z \rangle\rangle$ given $\langle\langle a \rangle\rangle$, $\langle\langle b \rangle\rangle$ and $\langle\langle c \rangle\rangle$ where $z = abc$ requires parties to obtain $\mathcal{E}[\Lambda_{ab}]$, $\mathcal{E}[\Lambda_{ac}]$, $\mathcal{E}[\Lambda_{bc}]$ and $\mathcal{E}[\Lambda_{abc}]$ in the preprocessing phase, which is then used to reconstruct m_z in the online phase.

Since parties in \mathcal{E} are required to hold the correct $\mathcal{E}[\cdot]$ -sharings before the online phase begins, as in the case of multiplication, the techniques from the semi-honest protocol fail in this setting. Hence, our protocol uses 4 instances of $\mathcal{F}_{\text{MulPre}}$ in the preprocessing phase, one each to compute $\langle\Lambda_{ab}\rangle$, $\langle\Lambda_{ac}\rangle$, $\langle\Lambda_{bc}\rangle$ and $\langle\Lambda_{abc}\rangle$. Each of the $\langle\cdot\rangle$ -sharing is further converted to $\mathcal{E}[\cdot]$ -sharing using $\Pi_{\langle\cdot\rangle \rightarrow \mathcal{E}[\cdot]}$ to ensure active participation of only $t + 1$ parties in the online phase for reconstruction of $z - r$. Further, to detect malicious behaviour during reconstruction of $z - r$, a verification check similar to the multiplication protocol is performed such that parties abort if the check fails. For 4-input multiplication, parties obtain $\langle\langle \cdot \rangle\rangle$ -sharing of $z = abcd$ using $z - r = (m_a - \lambda_a)(m_b - \lambda_b)(m_c - \lambda_c)(m_d - \lambda_d) - r$. The protocol proceeds in a similar manner as the 3-input case by delegating the computation of product terms to the preprocessing phase.

5. Building Blocks

For completeness, we discuss the building blocks used in our framework. These blocks are known from the literature [64, 80] and we show how these can be extended to the n -party setting.

5.1. Semi-Honest Building Blocks

Bit to arithmetic Given Boolean shares $\langle\langle b \rangle\rangle^{\mathbf{B}}$ of bit b , protocol Π_{bit2A} generates its arithmetic shares, $\langle\langle b^{\mathbf{R}} \rangle\rangle$ over \mathbb{Z}_{2^ℓ} (Fig. 28). Here, $b^{\mathbf{R}}$ denotes the arithmetic value of b over the ring \mathbb{Z}_{2^ℓ} . The approach is to generate a randomized version, $\zeta = b \oplus r$ of b , and then recover arithmetic shares of b by performing the arithmetic equivalent of XOR of $b = \zeta \oplus r$. Specifically, the arithmetic equivalent of $x \oplus y$ is given as $x^{\mathbf{R}} + y^{\mathbf{R}} - 2x^{\mathbf{R}}y^{\mathbf{R}}$.

Bit injection This protocol, denoted as Π_{BitInj} , facilitates generation of $\langle\langle b^{\mathbf{R}} \cdot v \rangle\rangle$ given $\langle\langle b \rangle\rangle^{\mathbf{B}}$, $\langle\langle v \rangle\rangle$ for $b \in \mathbb{Z}_2$ and $v \in \mathbb{Z}_{2^\ell}$. As seen in [64],

$$b^{\mathbf{R}}v = (m_b \oplus \lambda_b)^{\mathbf{R}}(m_v - \lambda_v) = m_b^{\mathbf{R}}m_v - m_b^{\mathbf{R}}\lambda_v + (2m_b^{\mathbf{R}} - 1)(\lambda_b^{\mathbf{R}}\lambda_v - m_v\lambda_b^{\mathbf{R}})$$

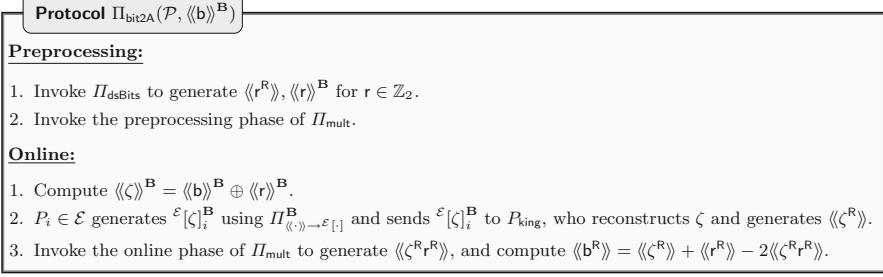


Fig. 28. Semi-honest: bit to arithmetic.

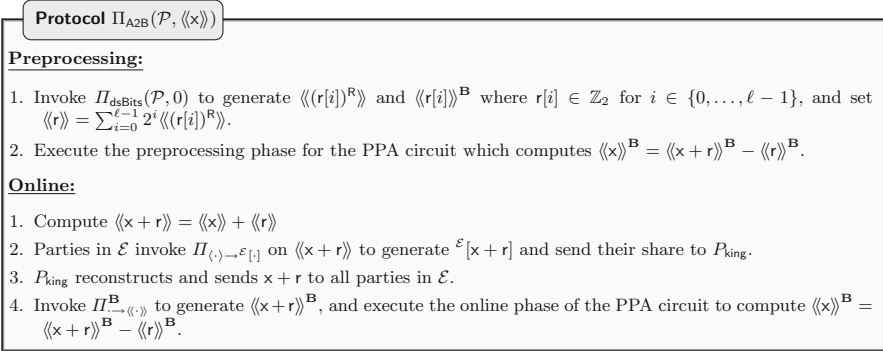


Fig. 29. Semi-honest: arithmetic to Boolean.

Given $\mathcal{E}[\cdot]$ -shares of $\lambda_{\mathbf{v}}, \lambda_{\mathbf{b}}^{\mathbf{R}}, \lambda_{\mathbf{b}}^{\mathbf{R}} \lambda_{\mathbf{v}}, \mathbf{r}$, together with $\langle\langle \mathbf{r} \rangle\rangle$ where $\mathbf{r} \in \mathbb{Z}_{2^\ell}$, and the knowledge that $\mathbf{m}_{\mathbf{v}}, \mathbf{m}_{\mathbf{b}}^{\mathbf{R}}$ is held by all parties in \mathcal{E} , parties can non-interactively compute $\mathcal{E}[\mathbf{b}^{\mathbf{R}} \mathbf{v} + \mathbf{r}]$, reconstruct it via P_{king} and generate $\langle\langle \mathbf{b}^{\mathbf{R}} \mathbf{v} + \mathbf{r} \rangle\rangle$. $\langle\langle \mathbf{b}^{\mathbf{R}} \mathbf{v} \rangle\rangle$ can then be computed as $\langle\langle \mathbf{b}^{\mathbf{R}} \mathbf{v} \rangle\rangle = \langle\langle \mathbf{b}^{\mathbf{R}} \mathbf{v} + \mathbf{r} \rangle\rangle - \langle\langle \mathbf{r} \rangle\rangle$. To facilitate this, in the preprocessing phase parties generate $\mathcal{E}[\cdot]$ -shares of $\mathbf{r}, \lambda_{\mathbf{v}}, \lambda_{\mathbf{b}}^{\mathbf{R}}, \lambda_{\mathbf{b}}^{\mathbf{R}} \lambda_{\mathbf{v}}$, and $\langle\langle \mathbf{r} \rangle\rangle$. Here, $\mathcal{E}[\mathbf{r}]$, $\mathcal{E}[\lambda_{\mathbf{v}}]$ and $\langle\langle \mathbf{r} \rangle\rangle$ are generated as in the preprocessing of multiplication, and $\mathcal{E}[\lambda_{\mathbf{b}}^{\mathbf{R}}]$ is generated via Π_{bit2A} followed by invoking $\Pi_{\langle \cdot \rangle \rightarrow \mathcal{E}_{[i]}}$. Following this, $\mathcal{E}[\lambda_{\mathbf{b}}^{\mathbf{R}} \lambda_{\mathbf{v}}]$ is generated as done in the preprocessing of multiplication.

Arithmetic to Boolean sharing Extending the techniques from [64], protocol Π_{A2B} generates $\langle\langle \mathbf{x} \rangle\rangle^{\mathbf{B}}$ from $\langle\langle \mathbf{x} \rangle\rangle$ for $\mathbf{x} \in \mathbb{Z}_{2^\ell}$. For this, given arithmetic and Boolean shares of $\mathbf{r} \in \mathbb{Z}_{2^\ell}$, Boolean shares of \mathbf{x} are computed as $(\mathbf{x} + \mathbf{r}) - \mathbf{r}$ by evaluating a parallel prefix adder (PPA) circuit [75, 80]. The PPA circuit takes as input two Boolean values $(\mathbf{x} + \mathbf{r}, -\mathbf{r}$ in this case) and outputs their sum. The protocol appears in Fig. 29. Looking ahead, Π_{A2B} is used in the preprocessing phase in the applications considered. Hence, we rely on the PPA circuit from [75] as it provides a good trade-off between rounds and communication as opposed to the circuit from [80] which is optimized to provide a fast online phase at the expense of a higher preprocessing cost (yielding a higher total cost than [75]).

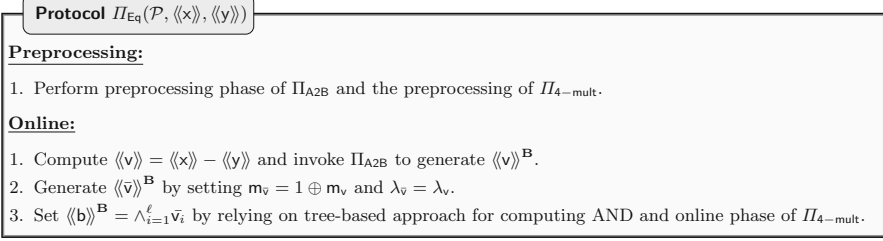


Fig. 30. Semi-honest: equality check protocol.

Boolean to arithmetic sharing This protocol generates $\langle\langle \mathbf{x} \rangle\rangle$ from $\langle\langle \mathbf{x} \rangle\rangle^{\mathbf{B}}$ where $\mathbf{x} \in \mathbb{Z}_{2^\ell}$. Inspired from [63, 64], observe that $\mathbf{x} = \sum_{i=0}^{\ell-1} 2^i (\mathbf{x}[i])^{\mathbf{R}}$. Thus, we invoke Π_{bit2A} on $\mathbf{x}[i]$ for $i \in \{0, \dots, \ell - 1\}$ to generate $\langle\langle \mathbf{x}[i] \rangle\rangle^{\mathbf{R}}$ followed by locally combining it as per the above equation to generate $\langle\langle \mathbf{x} \rangle\rangle$. Optimizations in [64] carry forward to our setting as well.

Comparison To compare $\mathbf{x}, \mathbf{y} \in \mathbb{Z}_{2^\ell}$ in FPA, we extend the technique of [26, 63, 64, 75, 80, 81], where checking $\mathbf{x} < \mathbf{y}$ is equivalent to checking if the most significant bit (msb) of $\mathbf{v} = \mathbf{x} - \mathbf{y}$ is 1. To extract the msb from $\langle\langle \mathbf{v} \rangle\rangle$, we rely on Π_{bitext} which takes as input $\langle\langle \mathbf{v} \rangle\rangle$ and outputs the $\langle\langle \cdot \rangle\rangle^{\mathbf{B}}$ -share of the msb of \mathbf{v} , denoted as $\langle\langle \text{msb}(\mathbf{v}) \rangle\rangle^{\mathbf{B}}$. The optimized bit extraction circuit from [80] is used for computing the msb whose inputs are two $\langle\langle \cdot \rangle\rangle^{\mathbf{B}}$ -shared values and output is the $\langle\langle \cdot \rangle\rangle^{\mathbf{B}}$ -shared msb of the sum of these two inputs. Observe that, given $\langle\langle \mathbf{v} \rangle\rangle$, \mathbf{v} can be written as $\mathbf{v} = m_{\mathbf{v}} - \lambda_{\mathbf{v}}$, and hence $\langle\langle \cdot \rangle\rangle^{\mathbf{B}}$ -shares of $m_{\mathbf{v}}$ and $\lambda_{\mathbf{v}}$ constitute the two inputs to the circuit. While $\langle\langle m_{\mathbf{v}} \rangle\rangle^{\mathbf{B}}$ can be generated non-interactively by invoking $\Pi_{\text{A2B}}^{\mathbf{B}}$ in the online phase, $\langle\langle \lambda_{\mathbf{v}} \rangle\rangle^{\mathbf{B}}$ is generated by performing an arithmetic to Boolean conversion in the preprocessing phase. Evaluation of bit extraction circuit then gives $\langle\langle \text{msb}(\mathbf{v}) \rangle\rangle^{\mathbf{B}}$.

Equality check Given $\langle\langle \cdot \rangle\rangle$ -shared $\mathbf{x}, \mathbf{y} \in \mathbb{Z}_{2^\ell}$, this protocol outputs a $\langle\langle \cdot \rangle\rangle^{\mathbf{B}}$ -shared bit, which is set to 1 if $\mathbf{x} = \mathbf{y}$, and 0 otherwise. The approach is to obtain the bit decomposition of $\mathbf{v} = \mathbf{x} - \mathbf{y}$ by performing Π_{A2B} and check if all bits of \mathbf{v} are 0. For this, parties non-interactively obtain 1's complement of the bits of \mathbf{v} , denoted as $\bar{\mathbf{v}}$, by setting the corresponding $m_{\bar{\mathbf{v}}} = 1 \oplus m_{\mathbf{v}}$ and $\lambda_{\bar{\mathbf{v}}} = \lambda_{\mathbf{v}}$. Parties proceed to compute an AND of all the bits in $\bar{\mathbf{v}}$ following the standard-tree-based approach where we use the 4-input multiplication to save on rounds and communication. If $\mathbf{v} = 0$, then the AND outputs 1 else it outputs a 0. The protocol appears in Fig. 30.

Maxpool/Minpool Maxpool allows parties to compute $\langle\langle \cdot \rangle\rangle$ -share of the maximum value \mathbf{x}_{max} among a vector of values $\vec{\mathbf{x}} = (\mathbf{x}_1, \dots, \mathbf{x}_\sigma)$. For this, we proceed along the lines of [64]. Observe that the maximum among two values $\mathbf{x}_i, \mathbf{x}_j$ can be computed by first using the secure comparison protocol to obtain $\langle\langle \mathbf{b} \rangle\rangle^{\mathbf{B}}$ such that $\mathbf{b} = 0$ if $\mathbf{x}_i \geq \mathbf{x}_j$ and 1 otherwise. Following this, parties can compute $\mathbf{b}(\mathbf{x}_j - \mathbf{x}_i) + \mathbf{x}_i$ using the bit injection protocol, to obtain the maximum value as the output. To compute the maximum among a vector of values, parties follow the standard binary tree-based approach where consecutive pairs of values are compared in a level-by-level manner. We refer to the resulting protocol as Π_{max} . A protocol Π_{min} for minpool can be worked out similarly.

ReLU The ReLU function, $\text{ReLU}(v) = \max(0, v)$, can be written as $\text{ReLU}(v) = \bar{b} \cdot v$, where bit $b = 1$ if $v < 0$ and 0 otherwise. Here \bar{b} denotes the complement of b . Given $\langle\langle v \rangle\rangle$, parties invoke Π_{bitext} on $\langle\langle v \rangle\rangle$ to obtain $\langle\langle b \rangle\rangle^{\mathbf{B}}$. The $\langle\langle \cdot \rangle\rangle^{\mathbf{B}}$ -sharing of \bar{b} is then computed, non-interactively, by setting $m_{\bar{b}} = 1 \oplus m_b$. Given $\langle\langle b \rangle\rangle^{\mathbf{B}}$ and $\langle\langle v \rangle\rangle$, ReLU can be computed using Π_{BitInj} .

5.2. Malicious Blocks

Note that the malicious variants for the building blocks such as bit to arithmetic, Boolean to arithmetic, and arithmetic to Boolean conversion, bit extraction, secure comparison, secure equality check, ReLU, maxpool, and convolutions, follow along similar lines to that of the semi-honest protocols with the difference that the underlying protocols used are replaced with their maliciously secure variants. Moreover, for steps that involve opening values via P_{king} , the reconstructed values are sent to all and are accompanied by a verification check similar to the one in the multiplication protocol.

5.3. Communication Cost

Table 4 summarizes the complexities of the various protocols, in the semi-honest as well as the malicious settings, discussed so far. Specifically, we tabulate the preprocessing communication cost and the online communication cost. We also report the online round complexity of the designed protocols.

6. Applications and Benchmarks

To evaluate the performance of our protocols, we benchmark some of the popular applications such as deep neural networks (NN), graph neural networks (GNN), similar sequence queries (SSQ), and biometric matching where MPC is used to achieve privacy. While these applications have been looked at in the small-party setting [6, 63, 75, 77, 80, 86, 90, 94], we believe the n -party setting is a better fit for reasons described in the introduction. To the best of our knowledge, we are the first to benchmark these in the multiparty honest-majority setting for more than four parties.

Benchmark environment The performance of our protocols is analysed using a prototype implementation building over the ENCRYPTO library [31] in C++17. We chose 64 bit ring ($\mathbb{Z}_{2^{64}}$) for our arithmetic world, and the operations over extended ring were carried out using the NTL library.⁶ Since the correctness and accuracy of the applications considered in the secure computation setting are already established, our benchmark aims to demonstrate our protocols' performance and is not fully functional. Moreover, we believe that incorporating state-of-the-art code optimizations like GPU-assisted computing can enhance the efficiency of our protocols, which is left as future work. Since there is no defined way to capture an adversary's misbehaviour, following standard practice [32, 63, 75], we benchmark honest executions of the protocols, which also include the steps performed for verification in the malicious case. We use multi-threading, wherever

⁶<https://libntl.org>.

Table 4. Communication and round complexity of protocols: semi-honest and malicious .

Building Block	Semi-honest		Malicious		Rounds Online	Rounds Online	Communication Preprocessing	Online	Rounds Online
	Communication Preprocessing	Online	Communication Preprocessing	Online					
Sharing	-	$(t + 1)\ell$	-	$2t\ell$	1	1	-	$2t\ell$	1
Reconstruction ^a	-	$3t\ell$	-	$n(q - g)\ell$	2	1	-	$n(q - g)\ell$	1
Multiplication	$t\ell$	$2t\ell$	$3t\ell$	$3t\ell$	2	2	$3t\ell$	$3t\ell$	2
3-input multiplication	$6t\ell$	$2t\ell$	$12t\ell$	$3t\ell$	2	2	$12t\ell$	$3t\ell$	2
4-input multiplication	$15t\ell$	$2t\ell$	$33t\ell$	$3t\ell$	2	2	$33t\ell$	$3t\ell$	2
Doubly shared bits	$4t(\ell + 2)$	-	$6t(\ell + 2)$	-	-	-	-	-	-
Multiplication with truncation	$4t(\ell + 2)\ell + t\ell$	$2t\ell$	$3t\ell + 6t(\ell + 2)\ell$	$3t\ell$	2	2	$3t\ell + 6t(\ell + 2)\ell$	$3t\ell$	2
Dot product	$t\ell$	$2t\ell$	$3t\ell$	$3t\ell$	2	2	$3t\ell$	$3t\ell$	2
Bit to arithmetic	$4t(\ell + 2) + t\ell$	$4t\ell$	$6t(\ell + 2) + 3t\ell$	$6t\ell$	4	4	$6t(\ell + 2) + 3t\ell$	$6t\ell$	4
Bit injection	$4t(\ell + 2) + 6t\ell$	$2t\ell$	$6t(\ell + 2) + 12t\ell$	$3t\ell$	2	2	$6t(\ell + 2) + 12t\ell$	$3t\ell$	2
Arithmetic to Boolean	$4t(\ell + 2)\ell + t\ell \log_2 \ell$	$2t\ell(1 + \log_2 \ell)$	$6t(\ell + 2)\ell + 3t\ell \log_2 \ell$	$3t\ell(1 + \log_2 \ell)$	2 + 2 $\log_2 \ell$	2 + 2 $\log_2 \ell$	$6t(\ell + 2)\ell + 3t\ell \log_2 \ell$	$3t\ell(1 + \log_2 \ell)$	2 + 2 $\log_2 \ell$
Boolean to arithmetic	$4t(\ell + 2)\ell$	$2t\ell$	$6t(\ell + 2)\ell$	$3t\ell$	2	2	$6t(\ell + 2)\ell$	$3t\ell$	2
Comparison ^b	$u_1 + 4t(\ell + 2)\ell + 3t\ell \log_2 \ell + 2t\ell$	$2tu_2$	$6t(\ell + 2)\ell + 6t\ell \log_2 \ell + 3t\ell + u_1$	$3tu_2$	2 $\log_4 \ell$	2 $\log_4 \ell$	$6t(\ell + 2)\ell + 6t\ell \log_2 \ell + 3t\ell + u_1$	$3tu_2$	2 $\log_4 \ell$

ℓ - size of ring in bits

^aAccounts for reconstruction towards all; $\mathbf{q} = \binom{n}{h}$, $\mathbf{g} = \binom{n-1}{h-1}$.

^b $u_1 = 3t\sigma_2 + 12t\sigma_3 + 33t\sigma_4$, $u_2 = \sigma_2 + \sigma_3 + \sigma_4$, $\sigma_2 = 41$, $\sigma_3 = 27$, $\sigma_4 = 47$ denote the number of AND gates in the bit extraction circuit of ABY2 [80] with 2, 3, 4 inputs, respectively

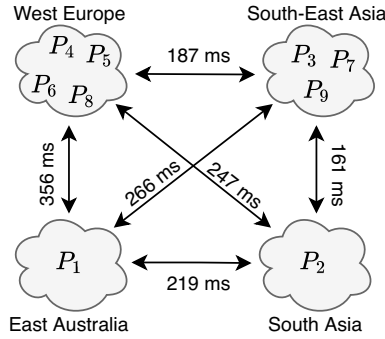


Fig. 31. Round-trip time (rtt) .

possible, to facilitate efficient computation and communication among the parties. The parties in the computation are emulated using Google Cloud (n1-standard-64 instances, 2.0 GHz Intel Xeon Skylake, 64 vCPUs, 240 GB RAM) with machines located in East Australia, South Asia, South East Asia, and West Europe. All our experiments are run for 5, 7, and 9 parties, each. We would like to note that our protocols can be scaled to a larger number of parties. However, recall that reliance on RSS will result in increasing the share size with increasing number of parties. Further, we note that the performance of our semi-honest protocol in the special setting of $t = 1$ is on par with tailor-made protocols such as [25]. For the malicious setting, note that customized protocols with $t = 1$ such as [26, 63, 64] are tailor-made for their setting and hence are more efficient. Specifically, they benefit from a single online round of interaction per multiplication gate, as opposed to two in our case, while having the same communication cost. We estimate this will roughly double the latency of our malicious protocol in this setting of $t = 1$. While a single-round protocol can be designed for the multiparty case, the two-round protocol ensures that communication complexity remains linear in the number of parties, as opposed to quadratic in the former. Since our focus is on attaining protocols that tolerate $t > 1$, we omit providing performance comparisons for these customized protocols in the $t = 1$ setting.

Benchmark parameters We report the run-time and communication of the online phase and total (= preprocessing + online). Note that the reported costs only consider the evaluation phase and do not account for the cost of input sharing and output reconstruction phases (because the latter phases amount to a one-time cost). Hence, for the malicious setting, the reported numbers do not account for the cost of broadcast required for the fair reconstruction. To capture the effect of online round complexity and communication in one go, we also report the throughput (TP [5, 63, 75]) of the online phase. TP denotes the number of operations that can be performed in one minute. Finally, when deployed in the outsourced setting, one pays the price for the communication and uptime of the hired servers. To demonstrate how our protocols fare in this scenario, we additionally report the monetary cost (Cost) [64, 74] for the applications considered. This cost is estimated using Google Cloud Platform [87] pricing, where 1 GB and 1 hour of usage cost USD 0.08 and USD 3.04, respectively.

Table 5. Communication (Preprocessing, Online) in MB for 1 million multiplications.

Ref.	$n = 5$	$n = 7$	$n = 9$
DN07* (semi)	(0, 45.78)	(0, 68.66)	(0, 91.55)
This (semi)	(15.26, 30.52)	(22.88, 45.78)	(30.51, 61.04)
This (mal)	(45.79, 45.78)	(68.67, 68.67)	(91.57, 91.57)

Table 6. Latency in seconds (Preprocessing, Online) for varying depth (d) circuits with 1 million multiplications for $n = 7$.

Ref.	$d = 1$	$d = 100$	$d = 1000$
DN07* (semi)	(0 ^a , 0.65)	(0, 54.97)	(0, 549.69)
This (semi)	(0.47, 0.45)	(0.47, 30.75)	(0.47, 307.48)
This (mal)	(10.52, 1.36)	(10.53, 68.67)	(10.54, 308.39)

^aCost of randomness generation in the preprocessing phase is same for DN07* and our protocols (**This**) and hence is omitted in the reported values

6.1. Comparison with DN07*

In this section, we benchmark our semi-honest and malicious protocols over synthetic circuits comprising one million multiplications with varying depths of 1, 100, and 1000 and compare against the optimized ring variant of DN07* [15]. The gates are distributed equally across each level in the circuit.

Communication The communication cost for 1 million multiplications is tabulated in Table 5 for the 5, 7, and 9 party settings. As can be observed, the online phase of our semi-honest protocol enjoys the benefits of pushing 33% communication to a preprocessing phase compared to DN07*. The observed values corroborate the claimed improvement in the online complexity of our protocol. Our malicious protocol retains the online communication cost of DN07* while incurring a similar overhead in the preprocessing phase.

Note that pushing the communication to the preprocessing phase has several benefits. First, communication with respect to several instances can happen in a single shot and leverage the benefit of serialization. Second, with respect to resource-constrained devices such as mobile phones, the preprocessing communication can occur whenever they have access to a high-bandwidth Wi-Fi network (for instance, when the device is at home overnight). These benefits facilitate a fast online phase, as observed, that may happen over a low-bandwidth network.

Run-time The time taken to evaluate circuits of different depths appears in Table 6. Since the time for the 5, 7, and 9 party settings varies within the range [0, 0.5], we report values only for the 7-party setting in Table 6. With respect to the online run-time, our semi-honest protocol's time is expected to be similar to that of DN07*. However, DN07* demonstrates around $1.5\times$ higher run-time. This difference can be attributed to the asymmetry in the rtt among parties, which vanished when benchmarked over a symmetric rtt setting. Compared to the semi-honest protocol, the malicious variant incurs a minimal overhead of less than one second in the online run-time due to the




Fig. 32. Monetary cost (in USD) for evaluating circuits (1000 instances) of various depths (d) for $n = 9$ parties. The values are reported in \log_2 scale .

one-time verification phase. However, the overhead is higher for the case of the overall run-time. Concretely, it is around 10s and is due to the distributed zero-knowledge proof computation in the preprocessing phase. Note that this overhead is independent of the circuit depth and gets amortized for deeper circuits as evident from Table 6 (depth 1 vs. 1000).

Monetary cost Another key highlight of our protocols is their improved monetary cost, as evident from Fig. 32.⁷ Concretely, for 9 parties (semi-honest), we observe a saving of 17% over DN07* for a depth-1 circuit, and it increases up to 72% for circuits with depth 1000. This is primarily due to the reduction in the number of online parties over DN07*. Comparing our semi-honest and malicious variants, the latter has an overhead of $8\times$ for depth-1 circuit, and it reduces to $1.14\times$ for depth-1000 circuit. This is justified because the verification cost is amortized for deeper circuits, as mentioned earlier. Interestingly, our malicious variant outperforms even the semi-honest DN07* upon reaching circuit depths of 100 and above. A similar analysis holds in the symmetric rtt setting as well, where the saving is up to 56% (for $d = 1000$).

Online throughput (TP): Owing to the asymmetric rtt as described earlier, our semi-honest variant witnesses up to $1.78\times$ improvements in TP (for a single execution) over DN07*, which vanishes in the symmetric rtt setting. However, recall that our protocol requires only $t + 1$ active parties in the online phase, which leaves several channels among the parties underutilized. Hence, we can leverage the load balancing technique where parties' roles are interchanged across various parallel executions. For instance, one approach is to make every party act as P_{king} , i.e. in 5PC, in one execution, $P_{\text{king}} = P_1, \mathcal{E} = \{P_1, P_2, P_3\}, \mathcal{D} = \{P_4, P_5\}$, while in another execution $P_{\text{king}} = P_2, \mathcal{E} = \{P_2, P_3, P_4\}, \mathcal{D} = \{P_5, P_1\}$, and so on. To analyse the effect of load balancing, we performed experiments with similar rtt among the parties and observed a $1.5\times$ improvement in our semi-honest variant over DN07*. This is justified as we communicate over four channels among the parties as opposed to six in DN07*. We note that while enhancing the security from semi-honest to malicious, we observe a significant

⁷Bars in solid colours denote computation over network given in Fig. 31, while the area represented via crosshatch pattern denotes the additional cost incurred in the symmetric rtt setting (356 ms).

drop in TP, which is about $3 \times$ for the depth-1 circuit. This is primarily due to increased run-time owing to the verification in online phase for malicious setting. However, this drop tends to zero for deeper circuits (as verification cost gets amortized), making the online phase of our maliciously secure protocol on par with the semi-honest one.

6.2. Deep Neural Networks (DNN) and Graph Neural Networks (GNN)

We begin by discussing the architectural details of the neural networks and graph neural networks that have been benchmarked in this work.

Neural networks We benchmark three different neural networks (NN) [75,81,94] with increasing number of parameters—(i) NN-1: a 3-layer fully connected network with ReLU activation after each layer, as considered in [63,75,77,81], (ii) NN-2: the LeNet [67] architecture, which contains two convolutional layers and two fully connected layers with ReLU activation after each layer, and maxpool operation after convolutional layers, and (iii) NN-3: VGG16 [91] architecture, that comprises 16 layers in total, which includes fully connected, convolutional, ReLU activation, and maxpool layers. The last 2 NNs were considered in [94]. We benchmark the inference phase of the above NNs, which comprises computing activation matrices, followed by applying an activation function or pooling operation, depending on the network architecture. NN-1 and NN-2 are benchmarked over MNIST dataset [68], while NN-3 is benchmarked using CIFAR-10 dataset [65].

Graph neural networks The goal of spectral-based GNNs [39,62] is to learn a function of signals $\vec{x}_1, \dots, \vec{x}_m$ each of length n , on a graph $G = (V, E, M)$, where V is the set of n vertices of the graph, E is the set of edges and M is the graph description in terms of an $n \times n$ adjacency matrix. The j^{th} component of every signal \vec{x}_i corresponds to j^{th} node of the graph. Training data are used to compute graph description M , which is common for all signals considered.

The approximation of graph filters using a truncated expansion in terms of Chebyshev polynomials was put forth in [39]. Chebyshev polynomials are recursively defined as follows:

$$T_k(x) = \begin{cases} 1 & \text{if } k = 0 \\ x & \text{if } k = 1 \\ 2xT_{k-1}(x) - T_{k-2}(x) & \text{otherwise} \end{cases}$$

The inference phase for a $n \times c$ signal matrix \mathbf{X} with f feature maps, where c represents the dimension of feature vector for each node, with a K -localized filter matrix Θ_k can be performed as $\mathbf{Y} = \sum_{k=0}^{K-1} T_k(\tilde{\mathbf{L}})\mathbf{X}\Theta_k$. Here, $\tilde{\mathbf{L}} = \frac{2}{\lambda_{max}} \cdot \mathbf{L} - \mathbf{I} \cdot \lambda_{max}$, and λ_{max} is the largest eigenvalue of the normalized graph Laplacian \mathbf{L} , \mathbf{Y} is an $n \times f$ -dimensional matrix and the trainable parameter for the k^{th} layer Θ_k is of dimension $c \times f$.

We use the simplified architecture of [39] given in [90]. The GNN architecture in the latter uses one graph convolution layer without pooling operation instead of the original model with two graph convolution layers, each of which is followed by a pooling operation. Further, K is set to 5 instead of 25. This architecture is shown to achieve an




Fig. 33. Comparison for GNN and deep NN between our semi-honest protocol and DN07* (values plotted are logarithmic in base 2) .

accuracy of more than 99% on MNIST classification in [90]. The GNN architecture [90] is as follows.

- Graph convolution layer:
 - *Input:* $T_k(\tilde{\mathbf{L}})$ with dimensions 784×784 , Θ_k with dimensions 1×32 , for $k \in \{0, \dots, K-1\}$, and 28×28 image transformed into a vector $\vec{\mathbf{x}}$ of dimension 784.
 - *Output:* $\sum_{k=0}^{K-1} T_k(\tilde{\mathbf{L}})\vec{\mathbf{x}}\Theta_k$ with dimensions 784×32 .
- ReLU activation: Calculates the ReLU for each input.
- Fully connected layer (FC): with 10 nodes.

Analysis To analyse the improvement of our protocols, we also benchmark (semi-honest) DN07* for the applications by adapting our building blocks to their setting. The semi-honest benchmarks for the different NNs and GNN appear in Table 7, while the malicious ones appear in Table 8. Figure 33 gives a pictorial view of the trends observed while comparing the semi-honest variants and are described next.

We incur a very minimal overhead in the run-time of our protocols when moving from five to nine parties over all the networks considered. Hence, we use $\pm\delta$ to denote this variation in the table. The trends witnessed in synthetic circuit benchmarks (Sect. 6.1) carry forward to neural networks as well due to reasons discussed previously. For instance, the improvement in the online run-time for our semi-honest variant is up to $4.3\times$ over DN07*. The effect of reduced run-time and improved communication results in a significant improvement in online throughput of our protocol over DN07*. Concretely, the gain ranges up to $4.3\times$. Further, the improved run-time coupled with the reduced number of online parties for our case brings in a saving of up to 69% in monetary cost for NN-1. However, the improvement drops to 33% for deep network NN-3. The reduction in savings is due to improved run-time getting nullified by increased communication from NN-1 to NN-3, making communication the dominant factor in determining monetary cost.

Observe that, unlike the case in synthetic circuits (Table 5), the total communication here is an order of magnitude higher. This is primarily due to the higher communication cost incurred for performing the truncation operation—specifically, generation of the doubly shared bits (Π_{dsBits}) in the preprocessing phase. It is worth noting that Π_{dsBits} is used as a black-box, and an improved instantiation for it will lower the communication.

Table 7. Semi-honest: benchmarks for deep NN and GNN .

Ref.	n	NN-1			NN-2			NN-3			GNN			
		Comm	Time	TP ^a	Comm	Time	TP	Comm	Time	TP	Comm	Time	TP	
Online	DN07*	5	0.16	18.55 ±.4	211.69	15.58	46.20	83.12	228.07	152.95	25.11	20.14	7.26	528.66
		7	0.24		202.48	23.39	48.39	79.35	342.24	160.10	23.99	30.22	7.54	509.38
		9	0.33		202.49	31.18	48.40	79.35	456.33	160.14	23.99	40.29	7.56	509.38
This		5	0.02	4.61 ±.02	832.61 ±.04	1.92	11.08 ±.02	346.60 ±.3	29.70	36.92 ±.02	104.01 ±.04	5.34	2.16 ±.02	1777.78 ±.06
		7	0.03		2.88			44.55				8.00		
		9	0.05		3.84			59.40				10.67		
End-to-end	DN07*	5	3.41	21.46	0.06	269.23	56.44	0.21	4288.26	213.77	1.34	956.21	17.00	0.20
		7	5.11	22.29	0.10	403.85	59.60	0.32	6432.39	227.28	1.96	1434.31	17.54	0.29
		9	6.81	22.31	0.13	538.47	59.61	0.42	8576.52	227.33	2.56	1912.41	17.57	0.38
This		5	3.41	11.09 ±.02	0.02	269.50	25.34 ±.03	0.10	4292.06	104.09 ±.03	0.91	956.46	8.97 ±.02	0.17
		7	5.11		0.03	404.25		0.14	6438.09		1.08	1434.69		0.26
		9	6.81		0.04	539.00		0.18	8584.12		1.71	1912.92		0.33

Communication in MB and time in seconds

^aTP denotes throughput

^bmonetary cost in USD




Fig. 34. Similar sequence queries.

Similar trends are observed for GNN as well, where the online run-time of DN07* is up to $3.5\times$ higher than our semi-honest protocol. This is reflected in the throughput where we gain up to $3.4\times$. Further, we observe savings of up to 15% in monetary cost due to the reduced number of active parties and lesser run-time.

Compared to our semi-honest variant for evaluating NNs, the malicious variant incurs a $2\times$ higher online communication cost for NN-1 and NN-2. However, this difference closes in with deeper NNs, with the communication being $1.5\times$ for NN-3. The drop in the difference can be attributed to the one-time cost of verification required in the malicious variant, which gets amortized over deeper circuits. Due to the same reason, in comparison with the semi-honest case, the malicious variant has an overhead of around 1 second in the online run-time, which in turn reflects in the reduced throughput. Similar to the semi-honest evaluation of NNs, the overall communication is an order of magnitude higher than the online communication due to the cost incurred for truncation during preprocessing. Also, analogous to the trend observed for synthetic circuits, the overhead in overall run-time is approximately 11 seconds owing to the distributed zero-knowledge proof verification required in the preprocessing phase. For GNN, the trend follows closely to that of NN-3, where the malicious variant incurs $1.5\times$ higher communication than its semi-honest counterpart.

6.3. Genome Sequence Matching

Given a genome sequence as a query, genome matching aims to identify the most similar sequence from a database of sequences. This task is also known as similar sequence query (SSQ). An SSQ algorithm on two sequences s and q requires the computation of edit distance (ED), which quantifies how different two sequences are by identifying the minimum number of additions, deletions, and substitutions needed to transform one sequence to the other. To compute the ED, we extend the (2-party) protocol from [86] which builds on top of the approximation from [6], to the n -party setting. We proceed to describe the high-level idea of the approximation algorithm for ED computation for a query sequence \mathbf{q} against a database of sequences $\{\mathbf{s}_1, \dots, \mathbf{s}_m\}$.

The ED approximation algorithm has a non-interactive phase, during which the database owner with the sequences $\mathbf{s}_1, \dots, \mathbf{s}_m$, generates a look-up table (LUT) for each sequence. These LUTs are then secret-shared among all the parties. To generate the LUT, the sequences in the database are aligned with respect to a common reference genome sequence (using the Wagner–Fischer algorithm [95]) and divided into blocks of a fixed, predetermined size. Based on the most frequently occurring block sequences in the database, an LUT is constructed consisting of these block values and their distance from each other. Specifically, for a database of m sequences $\{\mathbf{s}_1, \dots, \mathbf{s}_m\}$, each

Table 8. Malicious: benchmarks for deep NN and GNN .

<i>n</i>	NN-1		NN-2		NN-3		GNN						
	Comm	Time	TP ^a	Comm	Time	TP	Comm	Time	TP				
Online	5	0.04	5.44 ± 0.02	706.40 ± 0.04	2.88	11.93 ± 0.03	322.63 ± 2	44.56	37.91 ± 0.02	101.27 ± 0.04	8.01	3.02 ± 0.02	1275.75
	7	0.06			4.32			66.84			12.01		1267.35
	9	0.08			5.77			89.12			16.02		1267.32
End-to-end	5	3.59	22.96 ± 0.02	0.07	286.18	37.71 ± 0.04	0.15	4535.95	124.54	1.04	977.65	22.39 ± 0.03	0.23
	7	5.39		0.10	429.28		0.22	6804.06	126.69	1.53	1466.49		0.34
	9	7.20		0.11	571.98		0.27	9066.43	129.42	1.94	1954.95		0.42

Communication in MB and time in seconds

^aTP denotes throughput

^bmonetary cost in USD




Fig. 35. Edit distance between query \mathbf{q} and sequence \mathbf{s} with respect to a database of m sequences and ω blocks.

of length ω blocks, an LUT_i is constructed for each \mathbf{s}_i . Each LUT has m columns, one corresponding to each \mathbf{s}_i in the database, and ω rows, one corresponding to each block of a sequence, where $\text{LUT}_s[i][j]$ corresponds to the ED between block i of the sequence \mathbf{s} and \mathbf{s}_j . This completes the non-interactive phase of the ED approximation algorithm.

Given the LUTs, when a new query \mathbf{q} has to be processed, its ED must be computed from every sequence \mathbf{s} in the database. For this, similar to the non-interactive phase, the query is first aligned with the reference sequence and broken down into blocks of the same fixed size. Then, the i^{th} block from the query is matched with the i^{th} block of each sequence in the LUT for a sequence \mathbf{s} . If the block values match, then the precomputed distance is taken as the output for that block; otherwise, the output is taken to be 0. Finally, the resultant sum of distances for all the blocks is taken to be the approximated ED between \mathbf{q} and the sequence \mathbf{s} . Computing the ED to all such sequences \mathbf{s} in the database then allows the identification of the most similar sequence for the query using the minpool operation. Algorithms for ED computation between two sequences and SSQ appear in Fig. 35, Fig. 34, respectively, where accuracy and correctness follow from [6]. Since the generation of LUTs happens non-interactively, we only focus on the computation of ED with respect to the new query \mathbf{q} , which requires interaction, and benchmark the same.

The benchmarks for genome sequence matching appear in Table 9. Following [86], we consider three cases with different number of sequences in the database (m) and different block lengths (ω). We witness similar trends here, where our semi-honest protocol has improvements of up to $4\times$ in both online run-time and throughput over DN07*. Our malicious variant incurs a minimal overhead in the range of 5-6% in online run-time and total communication over the semi-honest counterpart.

For the monetary cost (Fig. 36), our semi-honest protocol has up to 66% saving over DN07*, and the malicious variant has around 42%-54% overhead over the semi-honest counterpart.

6.4. Biometric Matching

We extend support for biometric matching, which finds application in many real-world tasks such as face recognition [43] and fingerprint matching [56]. Given a database of m biometric samples $(\vec{s}_1, \dots, \vec{s}_m)$ each of size σ , and a user holding its sample \vec{u} , the goal of biometric matching is to identify the sample from the database that is “closest”

Table 9. Benchmarks for genome sequence matching.

Ref.	n	$m = 1000, \omega = 25$		$m = 2000, \omega = 30$		$m = 4000, \omega = 35$					
		Comm ^a	Time	TP ^b	Comm	Time	TP	Comm	Time	TP	
Online	DN07*	5	10.85	60.58	63.39	25.82	66.33	57.89	59.87	72.08	53.27
		7	16.28	63.60	60.38	38.75	69.63	55.15	89.86	75.65	50.76
		9	21.71	63.62	60.37	51.67	69.66	55.09	119.81	75.67	50.72
	This (semi)	5	6.42	16.12 ± 0.1	236.21 ± 15	15.39	17.61 ± 0.2	217.93 ± 0.2	35.87	19.34 ± 0.2	198.55 ± 3.5
		7	9.63			23.08			53.80		
	This (mal)	9	12.84			30.78			71.74		
		5	9.51	16.8 ± 1	228.71	22.79	18.3 ± 2	209.84	53.11	20.11 ± 0.6	190.95 ± 4
		7	14.14		228.44	33.88		209.49	78.96		
	End-to-End	DN07*	9	18.40		226.82	44.06		207.23	102.68	
5			0.17	74.13	0.25	0.40	82.24	0.31	0.92	92.04	0.43
7			0.25	77.76	0.37	0.60	86.99	0.47	1.39	98.90	0.64
This (semi)		9	0.33	77.79	0.50	0.80	87.39	0.62	1.85	98.93	0.84
		5	0.17	19.08 ± 0.2	0.07	0.40	21.69 ± 0.2	0.11	0.92	25.97 ± 0.3	0.21
This (mal)		7	0.25		0.10	0.60		0.16	1.39		0.31
		9	0.33		0.13	0.80		0.21	1.85		0.39
		5	0.18	31.21 ± 0.8	0.12	0.42	34.52 ± 2	0.17	0.99	41.83 ± 0.6	0.29
7		0.27		0.16	0.64		0.25	1.48		0.42	
9	0.36		0.21	0.85		0.30	1.97		0.52		

^acommunication in MB
^bTP denotes throughput
^ccommunication in GB
^dTime in seconds
^emonetary cost in USD




Fig. 36. Monetary cost for SSQ evaluation for varying number of sequences and block lengths ((1000,25), (2000, 30), (4000,35)) for $n = 9$ parties. Costs for 1000 instances are reported in USD.

to \vec{u} . The notion of “closeness” can be formalized by various distance metrics, of which Euclidean distance (**EuD**) is the most widely used. Following the general trend, we reduce our biometric matching problem to that of finding the sample from the database which has the least **EuD** with the user’s sample \vec{u} .

We follow [77,80] where the **EuD** between two vectors \vec{x}, \vec{y} each of length σ is given as

$$\mathbf{EuD}_{\vec{x}\vec{y}} = \sum_{i=1}^{i=\sigma} (\mathbf{x}_i - \mathbf{y}_i)^2 = \vec{z} \odot \vec{z} \quad (6)$$

where $\vec{z} = ((\mathbf{x}_1 - \mathbf{y}_1), \dots, (\mathbf{x}_\sigma - \mathbf{y}_\sigma))$.

To achieve this goal of performing biometric matching securely, each \vec{s}_i for all $i \in \{1, \dots, m\}$ in the database is $\langle\langle \cdot \rangle\rangle$ -shared among the n parties participating in the computation. Specifically, each component \vec{s}_{i_j} for all $j \in \{1, \dots, \sigma\}$ is $\langle\langle \cdot \rangle\rangle$ -shared among all the parties. Similarly, the user also $\langle\langle \cdot \rangle\rangle$ -shares its sample \vec{u} . The parties compute a $\langle\langle \cdot \rangle\rangle$ -shared distance vector \mathbf{DV} of size m , where the i^{th} component corresponds to the **EuD** between \vec{u} and \vec{s}_i . For this, each party locally obtains $\langle\langle \vec{z}_i \rangle\rangle = \langle\langle \vec{s}_i \rangle\rangle - \langle\langle \vec{u} \rangle\rangle$ and computes $\langle\langle \mathbf{DV}_i \rangle\rangle$ according to Eq. 6 using the dot product operation. The final step is then to identify the minimum of these m components of \mathbf{DV} , which can be performed using the protocol Π_{\min} for minpool operation.

The benchmarks for biometric matching appear in Table 10 for varying number of sequences. As is evident from Table 10, our semi-honest protocol witnesses a $4.6\times$ improvement over DN07* in both online run-time and throughput. Further, in terms of monetary cost, we observe a saving of around 85%. With respect to our maliciously secure protocol, we incur a minimal overhead of around 9.5% in terms of total communication and around 4% in online throughput over our semi-honest variant. We note that our malicious variant outperforms semi-honest DN07* in both online run-time and throughput, thereby achieving our goal of a fast online phase.

Table 10. Benchmarks for biometric matching .

Ref.	n	#seq = 1024		#seq = 4096		#seq = 16384		#seq = 65536						
		Comm	Time	TP ^a	Comm	Time	TP	Comm	Time	TP				
Online	DN07*	5	0.63	55.52	69.17	2.51	66.51	57.73	10.03	77.51	49.54	88.64	43.32	
		7	0.94	58.27	65.90	3.76	69.81	55.00	15.06	81.35	47.20	60.23	93.04	41.27
		9	1.25	58.30	65.88	5.02	69.87	54.97	20.07	81.36	47.14	80.31	93.10	41.16
This (semi)		5	0.09	12.61 ± 0.02	304.62 ± 0.03	0.35	15.07 ± 0.02	254.86 ± 0.03	1.41	17.53 ± 0.02	219.16 ± 0.04	5.62	19.99 ± 0.02	192.09 ± 0.04
		7	0.13			0.53			2.11			8.44		
		9	0.18			0.70			2.81			11.25		
This (mal)		5	0.14	13.43 ± 0.02	285.93 ± 0.2	0.53	15.89 ± 0.02	241.66 ± 0.1	2.11	18.35 ± 0.02	209.24 ± 0.06	8.44	20.86 ± 0.02	183.88 ± 0.07
		7	0.21			0.80			3.17			12.67		
		9	0.28			1.07			4.23			16.89		
End-to-End DN07*		5	6.92	66.55	0.20	27.70	79.85	0.24	110.83	93.47	0.29	443.34	108.53	0.40
		7	10.38	69.32	0.30	41.55	83.24	0.36	166.24	97.70	0.45	665.00	114.45	0.59
		9	13.84	69.35	0.40	55.40	83.30	0.48	221.66	97.71	0.59	886.67	114.55	0.79
This (semi)		5	6.93	14.79 ± 0.02	0.03	27.74	17.35 ± 0.03	0.04	110.99	20.24 ± 0.05	0.06	443.99	24.62 ± 0.1	0.13
		7	10.40		0.05	41.61		0.06	166.49		0.09	665.99		0.18
		9	13.86		0.06	55.49		0.08	221.99		0.11	887.99		0.23
This (mal)		5	7.61	26.67 ± 0.02	0.08	30.43	29.27 ± 0.02	0.09	121.71	32.30 ± 0.03	0.11	486.85	37.33 ± 0.05	0.18
		7	11.42		0.11	45.65		0.13	182.58		0.16	730.28		0.26
		9	15.22		0.14	60.81		0.16	243.19		0.20	972.72		0.33

Communication in MB and time in seconds

^aTP denotes throughput

^bmonetary cost in USD

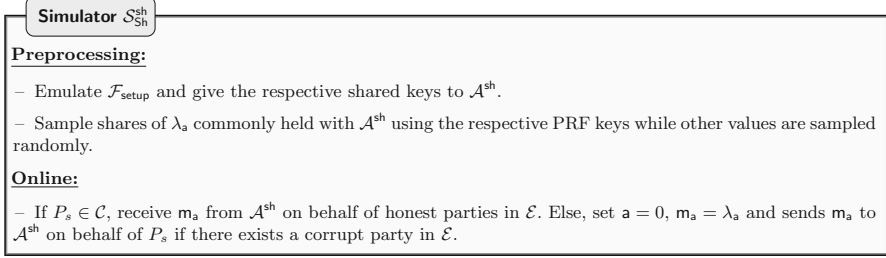


Fig. 37. Semi-honest: simulation for the input sharing protocol Π_{Sh} by P_s .




Fig. 38. Semi-honest: simulation for the reconstruction protocol towards all the parties .

7. Security Proofs

Security proofs are given in the real-world/ideal-world simulation-based paradigm [70]. Let \mathcal{A}^{sh} , \mathcal{A}^{mal} denote the real-world semi-honest, malicious adversary, respectively, corrupting at most t parties in \mathcal{P} , denoted by \mathcal{C} . Let \mathcal{S}^{sh} , \mathcal{S}^{mal} denote the corresponding ideal world semi-honest, malicious adversary, respectively. Security proofs are given in the $\{\mathcal{F}_{\text{setup}}, \mathcal{F}_{\text{TrGen}}\}$ -hybrid (and $\{\mathcal{F}_{\text{Broadcast}}, \mathcal{F}_{\text{TrGen}}^{\text{M}}, \mathcal{F}_{\text{MulPre}}, \mathcal{F}_{\text{DotPPre}}\}$ -hybrid for malicious setting) model. For modularity, we provide simulation steps for each protocol separately.

7.1. Semi-Honest Security

The following is the strategy for simulating the computation of function f (represented by a circuit ckt). The simulator \mathcal{S}^{sh} knows the input and output of the adversary \mathcal{A}^{sh} and sets the inputs of the honest parties to be 0. \mathcal{S}^{sh} emulates $\mathcal{F}_{\text{setup}}$ and gives the respective keys to the \mathcal{A}^{sh} . Knowing all the inputs and randomness, \mathcal{S}^{sh} can compute all the intermediate values for each building block in the clear. Thus, \mathcal{S}^{sh} proceeds to simulate each building block in topological order using the aforementioned values (input and output of \mathcal{A}^{sh} , randomness and intermediate values). We provide the simulation steps for each of the sub-protocols separately for modularity. When carried out in the respective order, these steps result in the simulation steps for the entire computation. To distinguish the simulators for various protocols, we use the corresponding protocol name as the subscript of \mathcal{S}^{sh} .

Sharing and Reconstruction Simulation for input sharing (Fig. 10) and reconstruction appears in Figs. 37, 38, respectively.

Multiplication Simulation steps for multiplication (Fig. 12) are provided in Fig. 39.

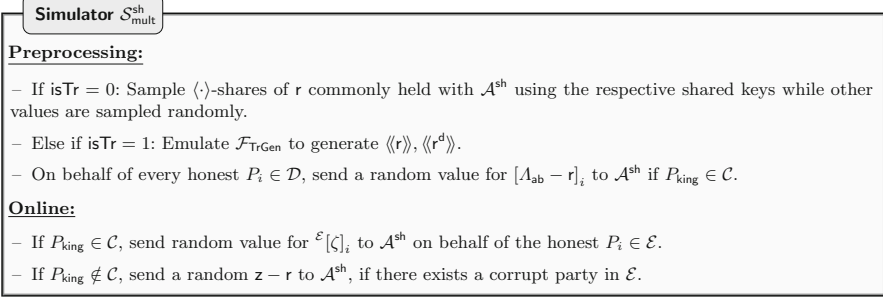


Fig. 39. Semi-honest: simulation for the multiplication protocol Π_{mult} .

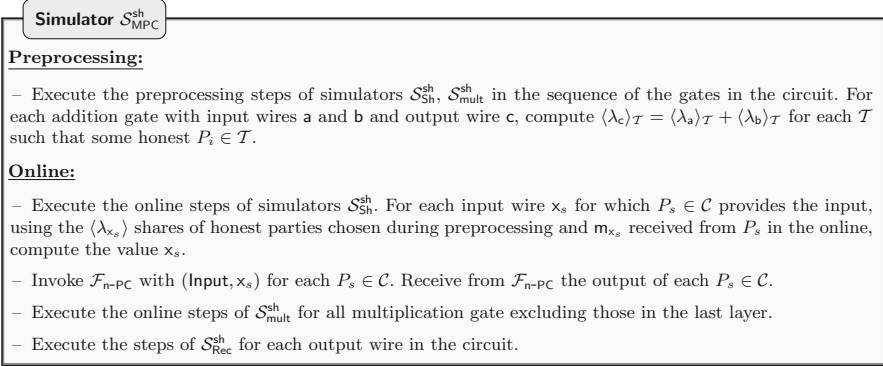


Fig. 40. Semi-honest: simulation for the complete MPC protocol $\Pi_{\text{MPC}}^{\text{sh}}$.

Observe that the adversary's view in the simulation is indistinguishable from its view in the real world since it only receives random values in each step of the protocol.

Other building blocks Simulation steps for the remaining building blocks can be obtained analogously by simulating the steps for the respective underlying protocols in their order of invocations.

Complete MPC protocol Simulation for the complete semi-honest MPC protocol $\Pi_{\text{MPC}}^{\text{sh}}$ (Fig. 13) appears in Fig. 40.

Theorem 1. *Protocol $\Pi_{\text{MPC}}^{\text{sh}}$ (Fig. 13) realizes $\mathcal{F}_{\text{n-PC}}$ (Fig. 9) with computational security against a semi-honest adversary \mathcal{A}^{sh} in the $\{\mathcal{F}_{\text{setup}}, \mathcal{F}_{\text{TrGen}}\}$ -hybrid model.*

Proof. We prove that the adversary's view in the simulation is indistinguishable from its view in the real world via a sequence of hybrids.

Hybrid₀: Execution of protocol $\Pi_{\text{MPC}}^{\text{sh}}$ in the real world.

Hybrid₁: In this hybrid, the execution of Π_{Sh} is replaced by the simulation of $\mathcal{S}_{\text{Sh}}^{\text{sh}}$. The two hybrids differ only in the case of inputs \mathbf{x}_s of each honest P_s . Note that for the input of an honest P_s , the adversary's view consists of $(\beta_{\mathbf{x}_s}, \langle \lambda_{\mathbf{x}_s} \rangle_i)$ for each $P_i \in \mathcal{C}$. Of these, $\langle \lambda_{\mathbf{x}_s} \rangle_i$ consists of random values selected using the shared-key setup among parties and hence is indistinguishable in both the hybrids. Moreover, $\beta_{\mathbf{x}_s}$ remains a random value from the adversary's view in both the hybrids due to the share $\langle \lambda_{\mathbf{x}_s} \rangle_{\mathcal{T}}$ where $\mathcal{T} \subseteq \mathcal{P} \setminus \mathcal{C}$ unknown to the adversary. Hence, the distributions of **Hybrid₀** and **Hybrid₁** are indistinguishable.

Hybrid₂: In this hybrid, the execution of Π_{mult} is replaced with the simulation of $\mathcal{S}_{\text{mult}}^{\text{sh}}$ for all the multiplication gates. The adversary's view here may consist of the reconstructed value $z - r$ if some corrupt party belongs to \mathcal{E} . However, note that it remains a random value from the adversary's view in both the hybrids due to the randomly chosen r . Moreover, r is unknown to the adversary due to the common share held by $n - t$ honest parties. Hence, the distributions of **Hybrid₁** and **Hybrid₂** are indistinguishable.

Hybrid₃: In this hybrid, the reconstruction protocol is replaced with the simulation of $\mathcal{S}_{\text{Rec}}^{\text{sh}}$. Note that this is exactly the execution in the ideal world. The transcript of a corrupt party for an output wire \mathbf{a} consists of shares of $\langle \lambda_{\mathbf{a}} \rangle$ and \mathbf{a} . As described in $\mathcal{S}_{\text{Rec}}^{\text{sh}}$, the simulator obtains the output wire value \mathbf{a} from the functionality and adjusts the shares of $\lambda_{\mathbf{a}}$ held only by the honest parties to ensure a sharing that is consistent with the output \mathbf{a} . Since $\lambda_{\mathbf{a}}$ is random and unknown to the adversary, $\mathbf{m}_{\mathbf{a}}$ is also random. Hence, the adversary's view is indistinguishable in both these executions.

Thus, we conclude that the view of the adversary is indistinguishable in **Hybrid₀**, which is the execution of the protocol in the real world and **Hybrid₃** corresponding to the execution in the ideal world. \square

7.2. Malicious Security

The following is the strategy for simulating the computation of function f (represented by a circuit ckt). The simulator emulates $\mathcal{F}_{\text{setup}}$ and gives the respective keys to the malicious adversary, \mathcal{A}^{mal} . This is followed by the input sharing phase in which \mathcal{S}^{mal} extracts the input of \mathcal{A}^{mal} , using the known keys, and sets the inputs of the honest parties to be 0. Knowing all the inputs, \mathcal{S}^{mal} can compute all the intermediate values for each building block in the clear. \mathcal{S}^{mal} proceeds to simulate each building block in topological order using the aforementioned values (inputs of \mathcal{A}^{mal} , intermediate values). Finally, depending on whether \mathcal{A}^{mal} misbehaved, which \mathcal{S}^{mal} can detect using the aforementioned information, \mathcal{S}^{mal} invokes $\mathcal{F}_{\text{n-PC}}^{\text{mal}}$ to obtain the function output and forwards it to \mathcal{A}^{mal} . As before, we provide the simulation steps for each of the sub-protocols separately for modularity. When carried out in the respective order, these steps result in the simulation steps for the entire computation. To distinguish the simulators for various protocols, the corresponding protocol name appears as the subscript of \mathcal{S}^{mal} .

Sharing Simulation for input sharing appears in Fig. 41.

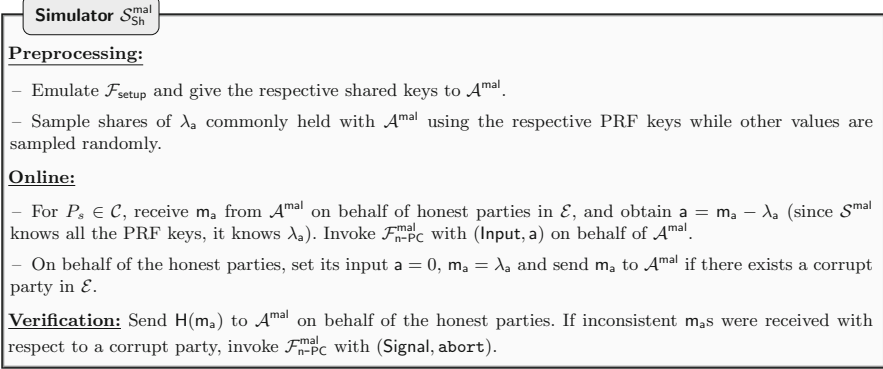


Fig. 41. Malicious: simulation for the input sharing protocol $\Pi_{\text{Sh}}^{\text{M}}$ by P_s .

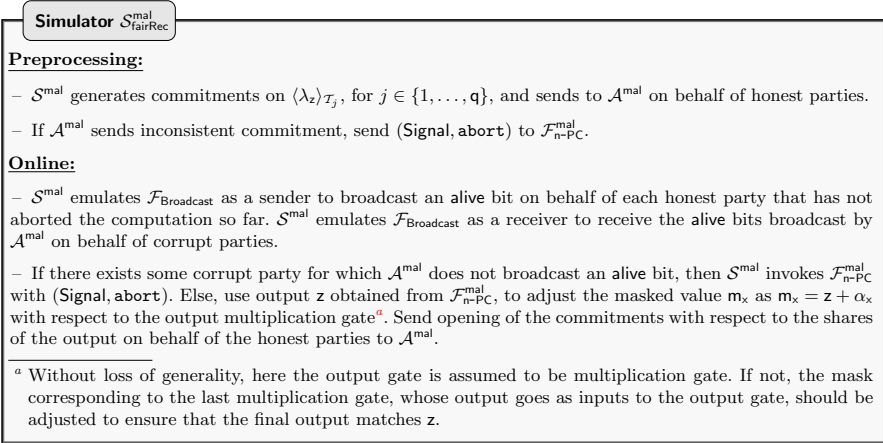


Fig. 42. Malicious: simulation for the fair reconstruction protocol $\Pi_{\text{Rec}}^{\text{fair}}$ towards all the parties.

Reconstruction Simulation for reconstruction (with fairness) appears in Fig. 42.

Multiplication Simulation steps for multiplication (Fig. 24) are provided in Fig. 43.

Observe that since \mathcal{A}^{mal} sees random shares in both the real-world protocol and in the simulation, indistinguishability of the simulation follows.

Other building blocks Simulations for the remaining building blocks can be obtained analogously and using the steps for the underlying protocols.

Complete MPC protocol Simulation for the complete malicious MPC protocol $\Pi_{\text{MPC}}^{\text{mal}}$ (Fig. 25) appears in Fig. 44.

Theorem 2. Protocol $\Pi_{\text{MPC}}^{\text{mal}}$ (Fig. 25) realizes $\mathcal{F}_{n\text{-PC}}^{\text{mal}}$ (Fig. 19) with computational security against a malicious adversary \mathcal{A}^{mal} in the $\{\mathcal{F}_{\text{Broadcast}}, \mathcal{F}_{\text{setup}}, \mathcal{F}_{\text{MulPre}}, \mathcal{F}_{\text{TrGen}}^{\text{M}}\}$ -hybrid model.

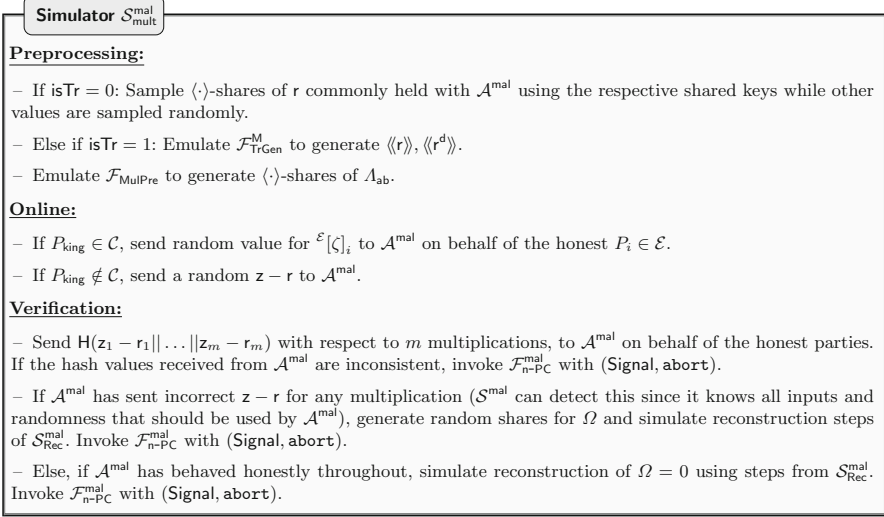


Fig. 43. Malicious: simulation for the multiplication protocol $\Pi_{\text{mult}}^{\text{M}}$.

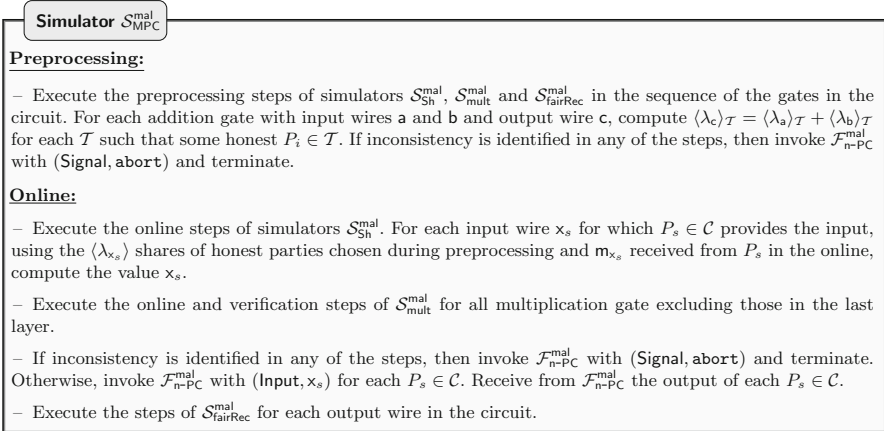


Fig. 44. Malicious: simulation for the complete MPC protocol $\Pi_{\text{MPC}}^{\text{mal}}$.

Proof. We prove that the adversary’s view in the simulation is indistinguishable from its view in the real world via a sequence of hybrids.

Hybrid₀: Execution of protocol $\Pi_{\text{MPC}}^{\text{mal}}$ in the real world.

Hybrid₁: In this hybrid, the execution of $\Pi_{\text{Sh}}^{\text{M}}$ is replaced by the simulation of $\mathcal{S}_{\text{Sh}}^{\text{mal}}$. Similar to the semi-honest protocol, the two hybrids differ only in the case of inputs x_s of each honest P_s . Note that for the input of an honest P_s , the adversary’s view consists of $(\beta_{x_s}, \langle \lambda_{x_s} \rangle_i)$ for each $P_i \in \mathcal{C}$. Of these, $\langle \lambda_{x_s} \rangle_i$ consists of random values selected using the shared-key setup among parties and hence is indistinguishable in both the hybrids.

Moreover, β_{x_s} remains a random value from the adversary's view in both the hybrids due to the share $\langle \lambda_{x_s} \rangle_{\mathcal{T}}$ where $\mathcal{T} \subseteq \mathcal{P} \setminus \mathcal{C}$ unknown to the adversary. Hence, the distributions of **Hybrid₀** and **Hybrid₁** are indistinguishable.

Hybrid₂: In this hybrid, the execution of Π_{mult}^M is replaced with the simulation of $\mathcal{S}_{\text{mult}}^{\text{mal}}$ for all the multiplication gates. The adversary's view here may consist of the reconstructed value $\mathbf{z} - \mathbf{r}$. For multiplication gates in the last layer of the circuit, the simulation differs from the gates in the other layers only for the reconstructed value $\mathbf{z} - \mathbf{r}$. Specifically, the simulator receives the output wire value of the multiplication gate by functionality $\mathcal{F}_{n\text{-PC}}$ and adjusts $\mathbf{z} - \mathbf{r}$ (by adjusting shares held by honest parties) accordingly. However, note that it remains a random value from the adversary's view in both the hybrids due to the randomly chosen \mathbf{r} . Moreover, \mathbf{r} is unknown to the adversary due to the common share held by $n - t$ honest parties in $\mathcal{T} \subseteq \mathcal{P} \setminus \mathcal{C}$. Hence, the distributions of **Hybrid₁** and **Hybrid₂** are indistinguishable.

Hybrid₃: In this hybrid, the fair reconstruction protocol $\Pi_{\text{Rec}}^{\text{fair}}$ is replaced with the simulation of $\mathcal{S}_{\text{fairRec}}^{\text{mal}}$. Note that this is exactly the execution in the ideal world. The transcript of a corrupt party for an output wire \mathbf{a} consists of shares of $\langle \lambda_{\mathbf{a}} \rangle$ and \mathbf{a} . Assume without loss of generality that the output wire is the output of a multiplication gate. As described earlier, the simulator obtains the output wire value \mathbf{a} from the functionality and adjusts the shares of $\mathbf{m}_{\mathbf{a}}$ held only by the honest parties to ensure a sharing that is consistent with the output \mathbf{a} . Since $\lambda_{\mathbf{a}}$ is random and unknown to the adversary, $\mathbf{m}_{\mathbf{a}}$ is also random. Therefore, the adversary's view is indistinguishable in both these executions.

Thus, we conclude that the view of the adversary is indistinguishable in **Hybrid₀**, which is the execution of the protocol in the real world and **Hybrid₃** corresponding to the execution in the ideal world. \square

Conclusion

This work improves the practical efficiency of n -party honest-majority protocols using preprocessing paradigm. While our first construction achieves a fast online phase compared to the semi-honest protocol of DN07*, the second enhances security by tolerating malicious adversaries with minimal overhead in the online phase. The active participation of half of the participants in both of our constructions is a major highlight. This reduction in online parties results in monetary benefits in real-world deployments.

Acknowledgements

The authors would like to acknowledge support from Centre for Networked Intelligence (a Cisco CSR initiative) at the Indian Institute of Science, Bengaluru, SERB MATRICS (Theoretical Sciences) Grant 2020, Google India AI/ML Research Award 2020, DST National Mission on Interdisciplinary Cyber-Physical Systems (NM-ICPS) 2020, National Security Council, India, and the support from Google Cloud to perform the benchmarking. 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). This work was co-funded by the Deutsche Forschungsgemeinschaft (DFG) – SFB 1119 CROSSING/236615297.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- [1] M. Abspoel, R. Cramer, I. Damgård, D. Escudero, C. Yuan, Efficient information-theoretic secure multiparty computation over $\mathbb{Z}/p^k\mathbb{Z}$ via galois rings, in *Theory of Cryptography Conference* (2019)
- [2] M. Abspoel, A.P.K. Dalskov, D. Escudero, A. Nof, An efficient passive-to-active compiler for honest-majority MPC over rings, in *ACNS* (2021)
- [3] A. Aly, E. Orsini, D. Rotaru, N.P. Smart, T. Wood, Zaphod: efficiently combining LSSS and garbled circuits in SCALE, in *ACM WAHC@CCS* (2019)
- [4] T. Araki, A. Barak, J. Furukawa, T. Lichter, Y. Lindell, A. Nof, K. Ohara, A. Watzman, O. Weinstein, Optimized honest-majority MPC for malicious adversaries—breaking the 1 billion-gate per second barrier, in *IEEE S&P* (2017)
- [5] T. Araki, J. Furukawa, Y. Lindell, A. Nof, K. Ohara, High-throughput semi-honest secure three-party computation with an honest majority, in *ACM CCS* (2016)
- [6] G. Asharov, S. Halevi, Y. Lindell, T. Rabin, Privacy-preserving search of similar patients in genomic data, in *PETS* (2018)
- [7] A. Baccarini, M. Blanton, C. Yuan, Multi-party replicated secret sharing over a ring with applications to privacy-preserving machine learning, ePrint Archive (2020). <https://eprint.iacr.org/2020/1577>
- [8] C. Baum, I. Damgård, T. Toft, R.W. Zakarias, Better preprocessing for secure multiparty computation, in *ACNS* (2016)
- [9] A. Ben-Efraim, Y. Lindell, E. Omri, Optimizing semi-honest secure multiparty computation for the internet, in *CCS* (2016)
- [10] A. Ben-Efraim, M. Nielsen, E. Omri, Turbospeedz: double your online spdz! improving SPDZ using function dependent preprocessing, in *ACNS* (2019)
- [11] A. Ben-Efraim, E. Omri, Concrete efficiency improvements for multiparty garbling with an honest majority, in *LATINCRYPT* (2017)
- [12] M. Ben-Or, S. Goldwasser, A. Wigderson, Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract), in *STOC* (1988)
- [13] M. Blanton, A. Kang, C. Yuan, Improved building blocks for secure multi-party computation based on secret sharing with honest majority, in *ACNS* (2020)
- [14] D. Bogdanov, S. Laur, J. Willemson, Sharemind: a framework for fast privacy-preserving computations, in *ESORICS* (2008)
- [15] D. Boneh, E. Boyle, H. Corrigan-Gibbs, N. Gilboa, Y. Ishai, Zero-knowledge proofs on secret-shared data via fully linear PCPs, in *CRYPTO* (2019)
- [16] E. Boyle, N. Gilboa, Y. Ishai, A. Nof, Practical fully secure three-party computation via sublinear distributed zero-knowledge proofs, in *ACM CCS* (2019)
- [17] E. Boyle, N. Gilboa, Y. Ishai, A. Nof, Efficient fully secure computation via distributed zero-knowledge proofs, in *ASIACRYPT* (2020)

- [18] Z. Brakerski, C. Gentry, V. Vaikuntanathan, (Leveled) Fully homomorphic encryption without bootstrapping, *ACM Trans. Comput. Theory* (2014)
- [19] L. Braun, D. Demmler, T. Schneider, O. Tkachenko, Motion—a framework for mixed-protocol multiparty computation, *ACM Trans. Privacy Secur.* (2022)
- [20] M. Byali, H. Chaudhari, A. Patra, A. Suresh, FLASH: fast and robust framework for privacy-preserving machine learning, in *PETS* (2020)
- [21] S. Carpov, K. Deforth, N. Gama, M. Georgieva, D. Jetchev, J. Katz, I. Leontiadis, M. Mohammadi, A. Sae-Tang, M. Vuille, Manticore: efficient framework for scalable secure multiparty computation protocols, ePrint Archive (2021). <https://eprint.iacr.org/2021/200>
- [22] O. Catrina, S.D. Hoogh, Improved primitives for secure multiparty integer computation, in *SCN* (2010)
- [23] O. Catrina, A. Saxena, Secure computation with fixed-point numbers, in *FC* (2010)
- [24] N. Chandran, N. Dasgupta, D. Gupta, S.L.B. Obbattu, S. Sekar, A. Shah, Efficient linear multiparty PSI and extensions to circuit/quorum PSI, in *ACM CCS* (2021)
- [25] H. Chaudhari, A. Choudhury, A. Patra, A. Suresh, ASTRA: high throughput 3PC over rings with application to secure prediction, in *ACM CCSW@CCS* (2019)
- [26] H. Chaudhari, R. Rachuri, A. Suresh, Trident: efficient 4PC framework for privacy preserving machine learning, in *NDSS* (2020)
- [27] K. Chida, D. Genkin, K. Hamada, D. Ikarashi, R. Kikuchi, Y. Lindell, A. Nof, Fast large-scale honest-majority MPC for malicious adversaries, in *CRYPTO* (2018)
- [28] P. Covington, J. Adams, E. Sargin, Deep neural networks for youtube recommendations, in *RecSys* (2016)
- [29] R. Cramer, I. Damgård, D. Escudero, P. Scholl, C. Xing, SPDZ_{2^k} : efficient MPC mod 2^k for dishonest majority, in *CRYPTO* (2018)
- [30] R. Cramer, I. Damgård, Y. Ishai, Share conversion, pseudorandom secret-sharing and applications to secure computation, in *TCC* (2005)
- [31] Cryptography, at TU Darmstadt, P.E.G.: ENCRYPTO Utils (2017). https://github.com/encryptogroup/ENCRYPTO_utils
- [32] A. Dalskov, D. Escudero, M. Keller, Fantastic four: honest-majority four-party secure computation with malicious security, in *USENIX Security* (2021)
- [33] I. Damgård, D. Escudero, T.K. Frederiksen, M. Keller, P. Scholl, N. Volgushev, New primitives for actively-secure MPC over rings with applications to private machine learning, in *IEEE S&P* (2019)
- [34] 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* (2013)
- [35] I. Damgård, J.B. Nielsen, Scalable and unconditionally secure multiparty computation, in *CRYPTO* (2007)
- [36] I. Damgård, C. Orlandi, M. Simkin, Yet another compiler for active security or: efficient MPC over arbitrary rings, in *CRYPTO* (2018)
- [37] I. Damgård, V. Pastro, N.P. Smart, S. Zakarias, Multiparty computation from somewhat homomorphic encryption, in *CRYPTO* (2012)
- [38] E. Dans, How signal cleverly exposed Facebook’s disregard for privacy, in *Forbes* (2021). <https://www.forbes.com/sites/enriquezans/2021/05/07/how-signal-cleverly-exposed-facebooks-disregard-forprivacy>
- [39] M. Defferrard, X. Bresson, P. Vandergheynst, Convolutional neural networks on graphs with fast localized spectral filtering, in *NeurIPS* (2016)
- [40] D. Demmler, T. Schneider, M. Zohner, ABY—a framework for efficient mixed-protocol secure two-party computation, in *NDSS* (2015)
- [41] D. Dolev, H.R. Strong, Authenticated algorithms for Byzantine agreement, *SIAM J. Comput.* (1983)
- [42] C. Dwork, Differential privacy: a survey of results, in *TAMC* (2008)
- [43] Z. Erkin, M. Franz, J. Guajardo, S. Katzenbeisser, I. Lagendijk, T. Toft, Privacy-preserving face recognition, in *PETS* (2009)
- [44] D. Escudero, A. Dalskov, Honest majority MPC with abort with minimal online communication, in *LATINCRYPT* (2021)
- [45] J. Furukawa, Y. Lindell, A. Nof, O. Weinstein, High-throughput secure three-party computation for malicious adversaries and an honest majority, in *EUROCRYPT* (2017)

- [46] D. Genkin, Y. Ishai, M.M. Prabhakaran, A. Sahai, E. Tromer, Circuits resilient to additive attacks with applications to secure computation, in *STOC* (2014)
- [47] C. Gentry, A. Sahai, B. Waters, Homomorphic encryption from learning with errors: conceptually-simpler, asymptotically-faster, attribute-based, in *CRYPTO* (2013)
- [48] O. Goldreich, *Foundations of Cryptography: Volume 2, Basic Applications* (Cambridge University Press, 2009)
- [49] O. Goldreich, S. Micali, A. Wigderson, How to play any mental game or A completeness theorem for protocols with honest majority, in *STOC* (1987)
- [50] S.D. Gordon, S. Ranellucci, X. Wang, Secure computation with low communication from cross-checking, in *ASIACRYPT* (2018)
- [51] S.D. Gordon, D. Starin, A. Yerukhimovich, The more the merrier: reducing the cost of large scale MPC, in *EUROCRYPT* (2021)
- [52] V. Goyal, H. Li, R. Ostrovsky, A. Polychroniadou, Y. Song, ATLAS: efficient and scalable MPC in the honest majority setting, in *CRYPTO* (2021)
- [53] V. Goyal, Y. Liu, Y. Song, Communication-efficient unconditional MPC with guaranteed output delivery, in *CRYPTO* (2019)
- [54] V. Goyal, Y. Song, Malicious security comes free in honest-majority MPC, in *CRYPTO* (2020)
- [55] G. Guido, M.I. Prete, S. Miraglia, I. De Mare, Targeting direct marketing campaigns by neural networks, *J. Market. Manag.* (2011)
- [56] W. Henecka, T. Schneider, Faster secure two-party computation with less memory, in *AsiaCCS* (2013)
- [57] M. Ito, A. Saito, T. Nishizeki, Secret sharing scheme realizing general access structure. *Electron. Commun. Japan (Part III: Fundamental Electronic Science)* (1989)
- [58] J. Katz, V. Kolesnikov, X. Wang, Improved non-interactive zero knowledge with applications to post-quantum signatures, in *CCS* (2018)
- [59] M. Keller, E. Orsini, P. Scholl, MASCOT: faster malicious arithmetic secure computation with oblivious transfer, in *CCS* (2016)
- [60] M. Keller, V. Pastro, D. Rotaru, Overdrive: making SPDZ great again, in *EUROCRYPT* (2018)
- [61] M. Keller, P. Scholl, N.P. Smart, An architecture for practical actively secure MPC with dishonest majority, in *CCS* (2013)
- [62] T.N. Kipf, M. Welling, Semi-supervised classification with graph convolutional networks, in *ICLR* (2017)
- [63] N. Koti, M. Pancholi, A. Patra, A. Suresh, SWIFT: super-fast and robust privacy-preserving machine learning, in *USENIX Security* (2021)
- [64] N. Koti, A. Patra, R. Rachuri, A. Suresh, Tetrad: actively secure 4PC for secure training and inference, in *NDSS* (2022)
- [65] A. Krizhevsky, V. Nair, G. Hinton, The CIFAR-10 dataset (2014). <https://www.cs.toronto.edu/~kriz/cifar.html>
- [66] A. Lapets, N. Volgushev, A. Bestavros, F. Jansen, M. Varia, Secure MPC for analytics as a web application, in *IEEE SecDev* (2016)
- [67] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, Gradient-based learning applied to document recognition, in *Proceedings of the IEEE* (1998)
- [68] Y. LeCun, C. Cortes, MNIST handwritten digit database (2010). <http://yann.lecun.com/exdb/mnist/>
- [69] C. Li, B. Pang, Y. Liu, H. Sun, Z. Liu, X. Xie, T. Yang, Y. Cui, L. Zhang, Q. Zhang, Adsgnn: behavior-graph augmented relevance modeling in sponsored search, in *SIGIR* (2021)
- [70] Y. Lindell, How to simulate it—a tutorial on the simulation proof technique, in *Tutorials on the Foundations of Cryptography* (2017)
- [71] Y. Lindell, B. Pinkas, N.P. Smart, A. Yanai, Efficient constant round multi-party computation combining BMR and SPDZ, in *CRYPTO* (2015)
- [72] M. Malone, *How Does Facebook Know What Ads to Show You? (Example)*. Vici Media (2021). <https://www.vicimediainc.com/how-does-facebook-know-what-ads-to-show-you/>
- [73] S. Mazloom, P.H. Le, S. Ranellucci, S.D. Gordon, Secure parallel computation on national scale volumes of data, in *USENIX Security* (2020)
- [74] P. Miao, S. Patel, M. Raykova, K. Seth, M. Yung, Two-sided malicious security for private intersection-sum with cardinality, in *CRYPTO* (2020)
- [75] P. Mohassel, P. Rindal, ABY³: a mixed protocol framework for machine learning, in *CCS* (2018)

- [76] P. Mohassel, M. Rosulek, Y. Zhang, Fast and secure three-party computation: the garbled circuit approach, in *CCS* (2015)
- [77] P. Mohassel, Y. Zhang, SecureML: a system for scalable privacy-preserving machine learning, in *IEEE S&P* (2017)
- [78] S. Ohata, K. Nuida, Communication-efficient (client-aided) secure two-party protocols and its application, in *FC* (2020)
- [79] K. Park, J. Lee, J. Choi, Deep neural networks for news recommendations, in *CIKM* (2017)
- [80] A. Patra, T. Schneider, A. Suresh, H. Yalame, ABY2.0: improved mixed-protocol secure two-party computation, in *USENIX Security* (2021)
- [81] A. Patra, A. Suresh, BLAZE: blazing fast privacy-preserving machine learning, in *NDSS* (2020)
- [82] R. Poddar, S. Kalra, A. Yanai, R. Deng, R.A. Popa, J.M. Hellerstein, Senate: a maliciously-secure MPC platform for collaborative analytics, in *USENIX Security* (2021)
- [83] 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* (2018)
- [84] P. Rogaway, T. Shrimpton, Cryptographic Hash-function basics: definitions, implications, and separations for preimage resistance, second-preimage resistance, and collision resistance, in *FSE* (2004)
- [85] D. Rotaru, T. Wood, MarBled circuits: mixing arithmetic and Boolean circuits with active security, in *INDOCRYPT* (2019)
- [86] T. Schneider, O. Tkachenko, EPISODE: efficient privacy-preserving similar sequence queries on out-sourced genomic databases, in *AsiaCCS* (2019)
- [87] G.C.C. Services, Google cloud platform, network costs (2008). <https://cloud.google.com/vpc/network-pricing>, Computation costs. <https://cloud.google.com/compute/vm-instance-pricing>
- [88] A. Shamir, *How to Share a Secret Communication* (ACM, 1979), pp. 612–613
- [89] S. Sharma, C. Xing, Y. Liu, Privacy-preserving deep learning with SPDZ (2019)
- [90] L. Shen, X. Chen, J. Shi, Y. Dong, B. Fang, An efficient 3-party framework for privacy-preserving neural network inference, in *ESORICS* (2020)
- [91] K. Simonyan, A. Zisserman, Very deep convolutional networks for large-scale image recognition, in *ICLR* (2015)
- [92] J. So, B. Güler, A.S. Avestimehr, CodedPrivateML: a fast and privacy-preserving framework for distributed machine learning, *IEEE J. Sel. Areas Inf. Theory* (2021)
- [93] A. Suresh, MPCLeague: robust MPC platform for privacy-preserving machine learning, PhD Thesis (2021). <https://arxiv.org/pdf/2112.13338>
- [94] S. Wagh, S. Tople, F. Benhamouda, E. Kushilevitz, P. Mittal, T. Rabin, FALCON: honest-majority maliciously secure framework for private deep learning, in *PETS* (2020)
- [95] R.A. Wagner, M.J. Fischer, The string-to-string correction problem, *J. ACM* (1974)
- [96] X. Wang, S. Ranellucci, J. Katz, Authenticated garbling and efficient maliciously secure two-party computation, in *CCS* (2017)
- [97] J. Yang, Z. Liu, S. Xiao, C. Li, D. Lian, S. Agrawal, A. Singh, G. Sun, X. Xie, Graphformers: Gnn-nested transformers for representation learning on textual graph, in *NeurIPS* (2021)
- [98] A.C. Yao, Protocols for secure computations (extended abstract), in *FOCS* (1982)
- [99] J. Zhu, Y. Cui, Y. Liu, H. Sun, X. Li, M. Pelger, T. Yang, L. Zhang, R. Zhang, H. Zhao, Textgnn: improving text encoder via graph neural network in sponsored search, in *WWW* (2021)