SWIFT: Super-fast and Robust Privacy-Preserving Machine Learning

Abstract

Performing ML computation on private data while maintaining data privacy, aka Privacy-preserving Machine Learning (PPML), is an emergent field of research. Recently, PPML has seen a visible shift towards the adoption of the Secure Outsourced Computation (SOC) paradigm due to the heavy computation that it entails. In the SOC paradigm, computation is outsourced to a set of powerful and specially equipped servers that provide service on a pay-per-use basis. In this work, we propose SWIFT, a *robust* PPML framework for a range of ML algorithms in SOC setting, that guarantees output delivery to the users irrespective of any adversarial behaviour. Robustness, a highly desirable feature, evokes user participation without the fear of denial of service.

At the heart of our framework lies a highly-efficient, maliciously-secure, three-party computation (3PC) over rings that provides guaranteed output delivery (GOD) in the honest-majority setting. To the best of our knowledge, SWIFT is the first robust and efficient PPML framework in the 3PC setting. SWIFT is as fast as (and is strictly better in some cases than) the best-known 3PC framework BLAZE (Patra et al. NDSS'20), which only achieves fairness. We extend our 3PC framework for four parties (4PC). In this regime, SWIFT is as fast as the best known *fair* 4PC framework Trident (Chaudhari et al. NDSS'20) and twice faster than the best-known *robust* 4PC framework FLASH (Byali et al. PETS'20).

We demonstrate our framework's practical relevance by benchmarking popular ML algorithms such as Logistic Regression and deep Neural Networks such as VGG16 and LeNet, both over a 64-bit ring in a WAN setting. For deep NN, our results testify to our claims that we provide improved security guarantee while incurring no additional overhead for 3PC and obtaining $2\times$ improvement for 4PC.

1 Introduction

Privacy Preserving Machine Learning (PPML), a booming field of research, allows Machine Learning (ML) computations over private data of users while ensuring the privacy of the data. PPML finds applications in sectors that deal with sensitive/confidential data, e.g. healthcare, finance, and in cases where organisations are prohibited from sharing client information due to privacy laws such as CCPA and GDPR. However, PPML solutions make the already computationally heavy ML algorithms more compute-intensive. An average end-user who lacks the infrastructure required to run these tasks prefers to outsource the computation to a powerful set of specialized cloud servers and leverage their services on a pay-per-use basis. This is addressed by the Secure Outsourced Computation (SOC) paradigm, and thus is an apt fit for the need of the moment. Many recent works [10, 13, 14, 38, 39, 41, 44, 46, 52] exploit Secure Multiparty Computation (MPC) techniques to realize PPML in the SOC setting where the servers enact the role of the parties. Informally, MPC enables n mutually distrusting parties to compute a function over their private inputs, while ensuring the privacy of the same against an adversary controlling up to t parties. Both the training and prediction phases of PPML can be realized in the SOC setting. The common approach of outsourcing followed in the PPML literature, as well as by our work, requires the users to secret-share[1] their inputs between the set of hired (untrusted) servers, who jointly interact and compute the secret-shared output, and reconstruct it towards the users.

In a bid to improve practical efficiency, many recent works [5, 10, 13, 14, 18, 22–24, 31–33, 44] cast their protocols into the preprocessing model wherein the input-independent (yet function-dependent) phase computationally heavy tasks are computed in advance, resulting in a fast online phase. This paradigm suits scenario analogous to PPML setting, where functions (ML algorithms) typically need to be evaluated a large number of times, and the function description is known beforehand. To further enhance practical efficiency by leveraging CPU optimizations, recent works [6, 19, 21, 23, 25] propose MPC protocols that work over 32 or 64 bit rings. Lastly, solutions for a small number of parties have received a huge momentum due to the many cost-effective

---

[1]The threshold of the secret-sharing is decided based on the number of corrupt servers so that privacy is preserved.

customizations that they permit, for instance, a cheaper realisation of multiplication through custom-made secret sharing schemes [2, 3, 10, 13, 14, 44].

We now motivate the need for robustness aka guaranteed output delivery (GOD) over fairness[2], or even abort security[3], in the domain of PPML. Robustness provides the guarantee of output delivery to all protocol participants, no matter how the adversary misbehaves. Robustness is crucial for real-world deployment and usage of PPML techniques. Consider the following scenario wherein an ML model owner wishes to provide inference service. The model owner shares the model parameters between the servers, while the end-users share their queries. A protocol that provides security with abort or fairness will not suffice as in both the cases a malicious adversary can lead to the protocol aborting, resulting in the user not obtaining the desired output. This leads to denial of service and heavy economic losses for the service provider. For data providers, as more training data leads to more accurate models, collaboratively building a model enables them to provide better ML services, and consequently, attract more clients. A robust framework encourages active involvement from multiple data providers. Hence, for the seamless adoption of PPML solutions in the real world, the robustness of the protocol is of utmost importance. The hall-mark result of [16] suggests that an honest-majority amongst the servers is necessary to achieve robustness. Consequent to the discussion above, we focus on the honest-majority setting with a small set of parties, especially 3 and 4 parties, both of which have drawn enormous attention recently [2, 3, 7, 8, 10, 12–14, 28, 40, 42, 44]. Our protocols work over rings, are cast in the preprocessing paradigm, and achieve GOD.

Related Work We restrict the relevant work to a small number of parties and honest-majority, focusing first on MPC, followed by PPML. MPC protocols for a small population can be cast into orthogonal domains of low latency protocols [11, 12, 43], and high throughput protocols [1–3, 6, 8, 13, 15, 27, 28, 42, 44]. In the 3PC setting, [3, 13] provide efficient semi-honest protocols wherein ASTRA [13] improved upon [3] by casting the protocols in the preprocessing model and provided a fast online phase. ASTRA further provided security with fairness in the malicious setting with an improved online phase compared to [2]. Later, a maliciously-secure 3PC protocol based on distributed zero-knowledge techniques was proposed by Boneh et al. [7] providing abort security. Further, building on [7] and enhancing the security to GOD, Boyle et al. [8] proposed a concretely efficient 3PC protocol with an amortized communication cost of 3 field elements (can be extended to work over rings) per multiplication gate. Concurrently, BLAZE [44] provided a fair protocol in the preprocessing model, which required communicating 3 ring elements in each phase. However, BLAZE eliminated the reliance on the

computationally intensive distributed zero-knowledge system (whose efficiency kicks in for large circuit or many multiplication gates) from the online phase and pushed it to the preprocessing phase. This resulted in a faster online phase compared to [8].

In the regime of 4PC, Gordon et al. [29] presented protocols achieving abort security and GOD. However, [29] relied on expensive public-key primitives and broadcast channels to achieve GOD. Trident [14] improved over the abort protocol of [29], providing a fast online phase achieving security with fairness, and presented a framework for mixed world computations [25]. A robust 4PC protocol was provided in FLASH [10], which requires communicating 6 ring elements, each, in the preprocessing and online phases.

In the PPML domain, MPC has been used for various ML algorithms such as Decision Trees [37], Linear Regression [26, 47], k-means clustering [9, 30], SVM Classification [51, 54], Logistic Regression [49]. In the 3PC SOC setting, the works of ABY3 [39] and SecureNN [52], provide security with abort. This was followed by ASTRA [13], which improves upon ABY3 and achieves security with fairness. ASTRA presents primitives to build protocols for Linear Regression and Logistic Regression inference. Recently, BLAZE improves over the efficiency of ASTRA and additionally tackles training for the above ML tasks, which requires building additional PPML building blocks, such as truncation and bit to arithmetic conversions. In the 4PC setting, the first robust framework for PPML was provided by FLASH [10] which proposed efficient building blocks for ML such as dot product, truncation, MSB extraction, and bit conversion. The works of [10, 13, 14, 39, 41, 44, 52] work over rings to garner practical efficiency. In terms of efficiency, BLAZE and respectively FLASH and Trident are the closest competitors of this work in 3PC and 4PC settings. We now present our contributions and compare them with these works.

1.1 Our Contributions

We propose, **SWIFT**, a robust maliciously-secure framework for PPML in the SOC setting, with a set of 3 and 4 servers having an honest-majority. At the heart of our framework lies highly-efficient, maliciously-secure, 3PC and 4PC over rings (both $\mathbb{Z}_{2^\ell}$ and $\mathbb{Z}_{2^1}$) that provide GOD in the honest-majority setting. We cast our protocols in the preprocessing model, which helps obtain a fast online phase. As mentioned earlier, the input-independent (yet function-dependent) computations will be performed in the preprocessing phase.

To the best of our knowledge, SWIFT is the first robust and efficient PPML framework in the 3PC setting and is as fast as (and is strictly better in some cases than) the best known *fair* 3PC framework BLAZE [44]. We extend our 3PC framework for 4 servers. In this regime, SWIFT is as fast as the best known *fair* 4PC framework Trident [14] and twice faster than best known *robust* 4PC framework FLASH [10]. We detail our contributions next.

Building Blocks	3PC					4PC				
	Ref.	Pre. Comm. (ℓ)	Online Rounds	Online Comm. (ℓ)	Security	Ref.	Pre. Comm. (ℓ)	Online Rounds	Online Comm. (ℓ)	Security
Multiplication	[7]	1	1	2	Abort					
	[8]	-	3	3	GOD	Trident	3	1	3	Fair
	BLAZE	3	1	3	Fair	FLASH	6	1	6	GOD
	SWIFT	**3**	**1**	**3**	GOD	**SWIFT**	**3**	**1**	**3**	GOD
Dot Product						Trident	3	1	3	Fair
	BLAZE	3n	1	3	Fair	FLASH	6	1	6	GOD
	SWIFT	**3**	**1**	**3**	GOD	**SWIFT**	**3**	**1**	**3**	GOD
Dot Product with Truncation						Trident	6	1	3	Fair
	BLAZE	3n + 2	1	3	Fair	FLASH	8	1	6	GOD
	SWIFT	**15**	**1**	**3**	GOD	**SWIFT**	**4**	**1**	**3**	GOD
Bit Extraction						Trident	≈ 8	$\log \ell + 1$	≈ 7	Fair
	BLAZE	9	$1 + \log \ell$	9	Fair	FLASH	14	$\log \ell$	14	GOD
	SWIFT	**9**	$\mathbf{1 + \log \ell}$	**9**	GOD	**SWIFT**	$\approx \mathbf{7}$	$\log \ell$	$\approx \mathbf{7}$	GOD
Bit to Arithmetic						Trident	≈ 3	1	3	Fair
	BLAZE	9	1	4	Fair	FLASH	6	1	8	GOD
	SWIFT	**9**	**1**	**4**	GOD	**SWIFT**	$\approx \mathbf{3}$	**1**	**3**	GOD
Bit Injection						Trident	≈ 6	1	3	Fair
	BLAZE	12	2	7	Fair	FLASH	8	2	10	GOD
	SWIFT	**12**	**2**	**7**	GOD	**SWIFT**	$\approx \mathbf{6}$	**1**	**3**	GOD

– Notations: ℓ - size of ring in bits, n - size of vectors for dot product.

Table 1: 3PC and 4PC: Comparison of SWIFT with its closest competitors in terms of Communication and Round Complexity

Robust 3/4PC frameworks The framework consists of a range of primitives realized in a privacy-preserving way which is ensured via running computation in a secret-shared fashion. We use secret-sharing over both $\mathbb{Z}_{2^\ell}$ and its special instantiation $\mathbb{Z}_{2^1}$ and refer them as *arithmetic* and respectively *boolean* sharing. Our framework consists of realizations for all primitives needed for general MPC and PPML such as multiplication, dot-product, truncation, bit extraction (given arithmetic sharing of a value v, this is used to generate boolean sharing of the most significant bit (msb) of the value), bit to arithmetic sharing conversion (converts the boolean sharing of a single bit value to its arithmetic sharing), bit injection (computes the arithmetic sharing of b · v, given the boolean sharing of a bit b and the arithmetic sharing of a ring element v) and above all, input sharing and output reconstruction in the SOC setting. A highlight of our 3PC framework, which, to the best of our knowledge is achieved for the first time, is a robust dot-product protocol whose (amortized) communication cost is independent of the vector size, which we obtain by extending the techniques of [7, 8]. The performance comparison in terms of concrete cost for communication and rounds, for PPML primitives in both 3PC and 4PC setting, appear in Table 1. As claimed, SWIFT is on par with BLAZE for most of the primitives (while improving security from fair to GOD) and is strictly better than BLAZE in case of dot product and dot product with truncation. For 4PC, SWIFT is on par with Trident in most cases (and is slightly better for dot product with truncation and bit injection), while it is doubly faster than FLASH.

Applications and Benchmarking We demonstrate the practicality of our protocols by benchmarking PPML, particularly, Logistic Regression (training and inference) and popular Neural Networks (inference) such as [41], LeNet [35] and VGG16 [48] having millions of parameters. The NN training requires mixed-world conversions [14,25,39], which we leave as future work. Our PPML blocks can be used to perform training and inference of Linear Regression, Support Vector Machines, and Binarized Neural Networks (as demonstrated in [10, 13, 14, 44]).

Comparisons and Differences with Prior Works To begin with, we introduce a new primitive called Joint Message Passing (jmp) that allows two servers to relay a common message to the third server such that either the relay is successful or an honest server is identified. jmp is extremely efficient as for a message of ℓ elements it only incurs the minimal communication cost of ℓ elements (in an amortized sense). Without any extra cost, it allows us to replace several pivotal private communications, that may lead to abort, either because the malicious sender does not send anything or sends a wrong message. All our primitives, either for a general 3PC or a PPML task, achieve GOD relying on jmp.

Second, instead of using the multiplication of [8] (which has the same overall communication cost as that of our online phase), we build a new protocol. This is because the former involves distributed zero-knowledge protocols. The cost of this heavy machinery gets amortized only for large circuits having millions of gates, which is very unlikely for

inference and moderately heavy training tasks in PPML. As in BLAZE [44], we follow a similar structure for our multiplication protocol but differ considerably in techniques as our goal is to obtain GOD. Our approach is to manipulate and transform some of the protocol steps so that two other servers can locally compute the information required by a server in a round. However, this transformation is not straight forward since BLAZE was constructed with a focus towards providing only fairness (details appear in §3). The multiplication protocol forms a technical basis for our dot product protocol and other PPML building blocks. We emphasise again that the (amortized) cost of our dot product protocol is independent of the vector size.

Third, extending to 4PC brings several performance improvements over 3PC. Most prominent of all is a conceptually simple jmp instantiation, which forgoes the broadcast channel while retaining the same communication cost; and a dot product with cost independent of vector size sans the 3PC amortization technique.

Fourth, we provide robust protocols for input sharing and output reconstruction phase in the SOC setting, wherein a user shares its input with the servers, and the output is reconstructed towards a user. The need for robustness and communication efficiency together makes these tasks slightly non-trivial. As a highlight, we introduce a super-fast online phase for the reconstruction protocol, which gives $4\times$ improvement in terms of rounds (apart from improvement in the communication) compared to BLAZE. Although we aim for GOD, we ensure that an end-user is never part of a broadcast which is relatively expensive than *atomic* point-to-point communication.

1.2 Organisation of the paper

The rest of the paper is organized as follows. In §2 we describe the system model, preliminaries and notations used. §3 and §4 detail our constructs in the 3PC and respectively 4PC setting. These are followed by the applications and benchmarking in §5. The appendix §A elaborates on additional preliminaries. The protocols to complete the PPML framework and detailed cost analysis for all the 3PC and 4PC protocols are provided in appendix §B and §C respectively. The security proofs for our constructions follow in appendix §D.

2 Preliminaries

We consider a set of three servers $\mathcal{P} = \{P_0, P_1, P_2\}$ that are connected by pair-wise private and authentic channels in a synchronous network, and a static, malicious adversary that can corrupt at most one server. We use a broadcast channel for 3PC alone, which is inevitable [17]. For ML training, several data-owners who wish to jointly train a model, secret share (using the sharing semantics that will appear later) their data among the servers. For ML inference, a model-owner

and client secret share the model and the query, respectively, among the servers. Once the inputs are available in the shared format, the servers perform computations and obtain the output in the shared form. In the case of training, the output model is reconstructed towards the data-owners, whereas for inference, the prediction result is reconstructed towards the client. We assume that an arbitrary number of data-owners may collude with a corrupt server for training, whereas for the case of prediction, we assume that either the model-owner or the client can collude with a corrupt server. We prove the security of our protocols using a standard real-world / ideal-world paradigm. We also explore the above model for the four server setting with $\mathcal{P} = \{P_0, P_1, P_2, P_3\}$. The aforementioned setting has been explored extensively [10, 13, 14, 39, 41, 44].

Our constructions achieve the strongest security guarantee of GOD. A protocol is said to be *robust* or achieve GOD if all parties obtain the output of the protocol regardless of how the adversary behaves. In our model, this translates to all the data owners obtaining the trained model for the case of ML training, while the client obtaining the query output for ML inference. All our protocols are cast into: *input-independent* preprocessing phase and *input-dependent* online phase.

For 3/4PC, the function to be computed is expressed as a circuit ckt, whose topology is public, and is evaluated over an arithmetic ring $\mathbb{Z}_{2^\ell}$ or boolean ring $\mathbb{Z}_{2^1}$. For PPML, we consider computation over the same algebraic structure. To deal with floating-point values, we use Fixed-Point Arithmetic (FPA) [10, 13, 14, 39, 41, 44] representation in which 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 x least significant bits are reserved for the fractional part. The ℓ-bit integer is then treated as an element of $\mathbb{Z}_{2^\ell}$, and operations are performed modulo 2^ℓ. We set $\ell = 64, x = 13$, leaving $\ell - x - 1$ bits for the integer part.

The servers use a one-time key setup, modelled as a functionality $\mathcal{F}_{\text{setup}}$ (Fig. 6), to establish pre-shared random keys for pseudo-random functions (PRF) between them. A similar setup is used in [2, 8, 13, 28, 39, 44, 46] for three server case and in [10, 14] for four server setting. The key-setup can be instantiated using any standard MPC protocol in the respective setting. Further, our protocols make use of a *collision-resistant* hash function, denoted by $\mathsf{H}()$, and a commitment scheme, denoted by $\mathsf{Com}()$. The formal details of key setup, hash function, and commitment scheme are deferred to §A.

Notation 2.1. The i^{th} element of a vector $\vec{\mathbf{x}}$ is denoted as x_i. The dot product of two n length vectors, $\vec{\mathbf{x}}$ and $\vec{\mathbf{y}}$, is computed as $\vec{\mathbf{x}} \odot \vec{\mathbf{y}} = \sum_{i=1}^{n} \mathsf{x}_i \mathsf{y}_i$. For two matrices $\mathbf{X}, \mathbf{Y}$, the operation $\mathbf{X} \circ \mathbf{Y}$ denotes the matrix multiplication. The bit in the i^{th} position of an ℓ-bit value v is denoted by $\mathsf{v}[i]$.

Notation 2.2. For a bit $\mathsf{b} \in \{0, 1\}$, we use b^{R} to denote the equivalent value of b over the ring $\mathbb{Z}_{2^\ell}$. b^{R} will have its least significant bit set to b, while all other bits will be set to zero.

3 Robust 3PC and PPML

In this section, we first introduce the sharing semantics for three servers. Then, we introduce our new Joint Message Passing (jmp) primitive, which plays a crucial role in obtaining the strongest security guarantee of GOD, followed by our protocols in the three server setting.

Secret Sharing Semantics We use the following secret-sharing semantics.

○ $[\cdot]$-*sharing:* A value $v \in \mathbb{Z}_{2^\ell}$ is $[\cdot]$-shared among P_1, P_2, if P_s for $s \in \{1, 2\}$ holds $[v]_s \in \mathbb{Z}_{2^\ell}$ such that $v = [v]_1 + [v]_2$.

○ $\langle \cdot \rangle$-*sharing:* A value $v \in \mathbb{Z}_{2^\ell}$ is $\langle \cdot \rangle$-shared among $\mathcal{P}$, if
  – there exists $v_0, v_1, v_2 \in \mathbb{Z}_{2^\ell}$ such that $v = v_0 + v_1 + v_2$.
  – P_s holds $(v_s, v_{(s+1)\%3})$ for $s \in \{0, 1, 2\}$.

○ $\llbracket \cdot \rrbracket$-*sharing:* A value $v \in \mathbb{Z}_{2^\ell}$ is $\llbracket \cdot \rrbracket$-shared among $\mathcal{P}$, if
  – there exists $\alpha_v \in \mathbb{Z}_{2^\ell}$ that is $[\cdot]$-shared among P_1, P_2.
  – there exists $\beta_v, \gamma_v \in \mathbb{Z}_{2^\ell}$ such that $\beta_v = v + \alpha_v$ and P_0 holds $([\alpha_v]_1, [\alpha_v]_2, \beta_v + \gamma_v)$ while P_s for $s \in \{1, 2\}$ holds $([\alpha_v]_s, \beta_v, \gamma_v)$.

Arithmetic and Boolean Sharing *Arithmetic* sharing refers to sharing over $\mathbb{Z}_{2^\ell}$ while *boolean* sharing, denoted as $\llbracket \cdot \rrbracket^{\mathbf{B}}$, refers to sharing over $\mathbb{Z}_{2^1}$.

Linearity of the Secret Sharing Scheme Given $[\cdot]$-shares of v_1, v_2, and public constants c_1, c_2, servers can locally compute $[\cdot]$-share of $c_1 v_1 + c_2 v_2$ as $c_1 [v_1] + c_2 [v_2]$. It is trivial to see that linearity property is satisfied by $\langle \cdot \rangle$ and $\llbracket \cdot \rrbracket$ sharings.

3.1 Joint Message Passing primitive

The jmp primitive allows two servers to relay a common message to the third server such that either the relay is successful or an honest server (or a conflicting pair) is identified. The striking feature of jmp is that it offers a rate-1 communication i.e. for a message of ℓ elements, it only incurs a communication of ℓ elements (in an amortized sense). The task of jmp is captured in an ideal functionality (Fig. 8) and the protocol for the same appears in Fig. 1. Next, we give an overview.

Given two servers P_i, P_j possessing a common value $v \in \mathbb{Z}_{2^\ell}$, protocol Π_{jmp} proceeds as follows. First, P_i sends v to P_k while P_j sends a hash of v to P_k. The communication of the hash is done once and for all from P_j to P_k. In the simplest case, P_k receives a consistent (value, hash) pair, and the protocol terminates. In all other cases, a TTP is identified as follows without having to communicate v again. Importantly, the following part can be run once and for all instances of Π_{jmp} with P_i, P_j, P_k in the same roles, invoked in the final 3PC protocol. Consequently, the cost relevant to this part vanishes in an amortized sense, making the construction rate-1.

Each P_s for $s \in \{i, j, k\}$ maintains a bit b_s initialized to 0, as an indicator for inconsistency. When P_k receives an inconsistent (value, hash) pair, it sets $b_k = 1$ and sends the bit to both P_i, P_j, who cross-check with each other by exchanging

the bit and turn on their inconsistency bit if the bit received from either P_k or its fellow sender is turned on. A server broadcasts a hash of its value when its inconsistency bit is on;[4] P_k's value is the one it receives from P_i. At this stage, there are a bunch of possible cases and a detailed analysis determines an eligible TTP in each case.

Protocol $\Pi_{\mathsf{jmp}}(P_i, P_j, P_k, v)$

– Each server P_s for $s \in \{i, j, k\}$ initializes bit $b_s = 0$.

– P_i sends v to P_k, while P_j sends $\mathsf{H}(v)$ to P_k.

– P_k broadcasts "($\mathtt{accuse}, P_i$)", if P_i is silent and TTP $= P_j$. Analogously for P_j. If P_k accuses both P_i, P_j, then TTP $= P_i$. Otherwise, P_k receives some $\tilde{v}$ and either sets $b_k = 0$ when the value and the hash are consistent or sets $b_k = 1$. P_k then sends b_k to P_i, P_j and terminates if $b_k = 0$.

– If P_i does not receive a bit from P_k, it broadcasts "($\mathtt{accuse}, P_k$)" and TTP $= P_j$. Analogously for P_j. If both P_i, P_j accuse P_k, then TTP $= P_i$. Otherwise, P_s for $s \in \{i, j\}$ sets $b_s = b_k$.

– P_i, P_j exchange their bits to each other. If P_i does not receive b_j from P_j, it broadcasts "($\mathtt{accuse}, P_j$)" and TTP $= P_k$. Analogously for P_j. Otherwise, P_i resets its bit to $b_i \lor b_j$ and likewise P_j resets its bit to $b_j \lor b_i$.

– P_s for $s \in \{i, j, k\}$ broadcasts $\mathsf{H}_s = \mathsf{H}(v^*)$ if $b_s = 1$, where $v^* = v$ for $s \in \{i, j\}$ and $v^* = \tilde{v}$ otherwise. If P_k does not broadcast, terminate. If either P_i or P_j does not broadcast, then TTP $= P_k$. Otherwise,

  • If $\mathsf{H}_i \neq \mathsf{H}_j$: TTP $= P_k$.
  • Else if $\mathsf{H}_i \neq \mathsf{H}_k$: TTP $= P_j$.
  • Else if $\mathsf{H}_i = \mathsf{H}_j = \mathsf{H}_k$: TTP $= P_i$.

Figure 1: 3PC: Joint Message Passing Protocol

When P_k is silent, the protocol is understood to be complete. This is fine irrespective of the status of P_k– an honest P_k never skips this broadcast with inconsistency bit on, and a corrupt P_k implies honest senders. If either P_i or P_j is silent, then P_k is picked as TTP which is surely honest. A corrupt P_k could not make one of $\{P_i, P_j\}$ speak, as the senders (honest in this case) are in agreement on their inconsistency bit (due to their mutual exchange of inconsistency bit). When all of them speak and (i) the senders' hashes do not match, P_k is picked as TTP; (ii) one of the senders conflicts with P_k, the other sender is picked as TTP; and lastly (iii) if there is no conflict, P_i is picked as TTP. The first two cases are self-explanatory. In the last case, either P_j or P_k is corrupt. If not, a corrupt P_i can have honest P_k speak (and hence turn on its inconsistency bit), by sending a v' whose hash is not same as that of v and so inevitably, the hashes of honest P_j and P_k will conflict, contradicting (iii). As a final touch, we ensure that, in each step, a server raises a public alarm (via broadcast) accusing a server which is silent when it is not supposed to be, and the protocol terminates immediately by labelling the server as TTP who is neither the complainer nor the accused.

---
[4] hash can be computed on a combined message across many calls of jmp.

Notation 3.1. We say that P_i, P_j jmp-send v to P_k when they invoke $\Pi_{\mathsf{jmp}}(P_i, P_j, P_k, \mathsf{v})$.

Using jmp *in protocols.* As mentioned in the introduction, the jmp protocol needs to be viewed as consisting of two phases (*send, verify*), where *send* phase consists of P_i sending v to P_k and the rest goes to *verify* phase. Looking ahead, most of our protocols use jmp, and consequently, our final construction, either of general MPC or any PPML task, will have several calls to jmp. To leverage amortization, the *send* phase will be executed in all protocols invoking jmp on the flow, while the *verify* for a fixed ordered pair of senders will be executed once and for all in the end. The *verify* phase will determine if all the sends were correct. If not, a TTP is identified, as explained, and the computation completes with the help of TTP, just as in the ideal-world.

3.2 3PC Protocols

We now describe the protocols for 3 parties/servers and defer the communication analysis of our protocols to §B and security proofs to §D.1.

Sharing Protocol Protocol Π_{sh} (Fig. 9) allows a server P_i to generate $[\![\cdot]\!]$-shares of a value $\mathsf{v} \in \mathbb{Z}_{2^\ell}$. In the preprocessing phase, P_0, P_j for $j \in \{1, 2\}$ along with P_i sample a random $[\alpha_\mathsf{v}]_j \in \mathbb{Z}_{2^\ell}$, while P_1, P_2, P_i sample random $\gamma_\mathsf{v} \in \mathbb{Z}_{2^\ell}$. This allows P_i to know both α_v and γ_v in clear. During the online phase, if $P_i = P_0$, then P_0 sends $\beta_\mathsf{v} = \mathsf{v} + \alpha_\mathsf{v}$ to P_1. P_0, P_1 then jmp-send β_v to P_2 to complete the secret sharing. If $P_i = P_1$, P_1 sends $\beta_\mathsf{v} = \mathsf{v} + \alpha_\mathsf{v}$ to P_2. Then P_1, P_2 jmp-send $\beta_\mathsf{v} + \gamma_\mathsf{v}$ to P_0. The case for $P_i = P_2$ proceeds similar to that of P_1. The correctness of the shares held by each server is assured by the guarantees of Π_{jmp}. We defer the details of Π_{sh} to §B.2.

Joint Sharing Protocol Protocol Π_{jsh} (Fig. 10) allows two servers P_i, P_j to jointly generate a $[\![\cdot]\!]$-sharing of a value $\mathsf{v} \in \mathbb{Z}_{2^\ell}$ that is known to both. Towards this, servers execute the preprocessing of Π_{sh} (Fig. 9) to generate $[\alpha_\mathsf{v}]$ and γ_v. If $(P_i, P_j) = (P_1, P_0)$, then P_1, P_0 jmp-send $\beta_\mathsf{v} = \mathsf{v} + \alpha_\mathsf{v}$ to P_2. The case when $(P_i, P_j) = (P_2, P_0)$ proceeds similarly. The case for $(P_i, P_j) = (P_1, P_2)$ is optimized further as follows: servers locally set $[\alpha_\mathsf{v}]_1 = [\alpha_\mathsf{v}]_2 = 0$. P_1, P_2 together sample random $\gamma_\mathsf{v} \in \mathbb{Z}_{2^\ell}$, set $\beta_\mathsf{v} = \mathsf{v}$ and jmp-send $\beta_\mathsf{v} + \gamma_\mathsf{v}$ to P_0. We defer the formal details of Π_{jsh} to §B.3.

Addition Protocol Given $[\![\cdot]\!]$-shares on input wires x, y, servers can use linearity property of the sharing scheme to locally compute $[\![\cdot]\!]$-shares of the output of addition gate, $\mathsf{z} = \mathsf{x} + \mathsf{y}$ as $[\![\mathsf{z}]\!] = [\![\mathsf{x}]\!] + [\![\mathsf{y}]\!]$.

Multiplication Protocol Protocol $\Pi_{\mathsf{mult}}(\mathcal{P}, [\![\mathsf{x}]\!], [\![\mathsf{y}]\!])$ (Fig. 2) enables the servers in $\mathcal{P}$ to compute $[\![\cdot]\!]$-sharing of $\mathsf{z} = \mathsf{xy}$, given the $[\![\cdot]\!]$-sharing of x and y. We build on the protocol of BLAZE [44] and discuss along the way the differences and resemblances. We begin with a protocol for the semi-honest setting, which is also the starting point of BLAZE. During

the preprocessing phase, P_0, P_j for $j \in \{1, 2\}$ sample random $[\alpha_\mathsf{z}]_j \in \mathbb{Z}_{2^\ell}$, while P_1, P_2 sample random $\gamma_\mathsf{z} \in \mathbb{Z}_{2^\ell}$. In addition, P_0 locally computes $\Gamma_{\mathsf{xy}} = \alpha_\mathsf{x}\alpha_\mathsf{y}$ and generates $[\cdot]$-sharing of the same between P_1, P_2. Since,

$$\beta_\mathsf{z} = \mathsf{z} + \alpha_\mathsf{z} = \mathsf{xy} + \alpha_\mathsf{z} = (\beta_\mathsf{x} - \alpha_\mathsf{x})(\beta_\mathsf{y} - \alpha_\mathsf{y}) + \alpha_\mathsf{z}$$
$$= \beta_\mathsf{x}\beta_\mathsf{y} - \beta_x\alpha_\mathsf{y} - \beta_\mathsf{y}\alpha_\mathsf{x} + \Gamma_{\mathsf{xy}} + \alpha_\mathsf{z} \qquad (1)$$

servers P_1, P_2 locally compute $[\beta_\mathsf{z}]_j = (j-1)\beta_\mathsf{x}\beta_\mathsf{y} - \beta_x[\alpha_\mathsf{y}]_j - \beta_\mathsf{y}[\alpha_\mathsf{x}]_j + [\Gamma_{\mathsf{xy}}]_j + [\alpha_\mathsf{z}]_j$ during the online phase and mutually exchange their shares to reconstruct β_z. P_1 then sends $\beta_\mathsf{z} + \gamma_\mathsf{z}$ to P_0, completing the semi-honest protocol. The correctness that asserts $\mathsf{z} = \mathsf{xy}$ or in other words $\beta_\mathsf{z} - \alpha_\mathsf{z} = \mathsf{xy}$ holds due to Eq. 1.

The following issues arise in the above protocol when a malicious adversary is considered:

1) When P_0 is corrupt, the $[\cdot]$-sharing of Γ_{xy} performed by P_0 might not be correct, i.e. $\Gamma_{\mathsf{xy}} \neq \alpha_\mathsf{x}\alpha_\mathsf{y}$.

2) When P_1 (or P_2) is corrupt, $[\cdot]$-share of β_z handed over to the fellow honest evaluator during the online phase might not be correct, causing reconstruction of an incorrect β_z.

3) When P_1 is corrupt, the value $\beta_\mathsf{z} + \gamma_\mathsf{z}$ that is sent to P_0 during the online phase may not be correct.

All the three issues are common with BLAZE (copied verbatim), but we differ from BLAZE in handling them. We begin with solving the last issue first. We simply make P_1, P_2 jmp-send $\beta_\mathsf{z} + \gamma_\mathsf{z}$ to P_0 (after β_z is computed). This either leads to success or a TTP selection. Due to jmp's rate-1 communication, P_1 alone sending the value to P_0 remains as costly as using jmp in amortized sense. Whereas in BLAZE, the malicious version simply makes P_2 to send a hash of $\beta_\mathsf{z} + \gamma_\mathsf{z}$ to P_0 (in addition to P_1's communication of $\beta_\mathsf{z} + \gamma_\mathsf{z}$ to P_0), who aborts if the received values are inconsistent.

For the remaining two issues, similar to BLAZE, we reduce both to a multiplication (on values unrelated to inputs) in the preprocessing phase. However, our method leads to either success or TTP selection, with no additional cost.

We start with the second issue. To solve it, where a corrupt P_1 (or P_2) sends an incorrect $[\cdot]$-share of β_z, BLAZE makes use of server P_0 to compute a version of β_z for verification, based on β_x and β_y, as follows. Using $\beta_\mathsf{x} + \gamma_\mathsf{x}, \beta_\mathsf{y} + \gamma_\mathsf{y}, \alpha_\mathsf{x}, \alpha_\mathsf{y}, \alpha_\mathsf{z}$ and Γ_{xy}, P_0 computes:

$$\beta_\mathsf{z}^\star = -(\beta_\mathsf{x} + \gamma_\mathsf{x})\alpha_\mathsf{y} - (\beta_\mathsf{y} + \gamma_\mathsf{y})\alpha_\mathsf{x} + 2\Gamma_{\mathsf{xy}} + \alpha_\mathsf{z}$$
$$= (\beta_\mathsf{z} - \beta_\mathsf{x}\beta_\mathsf{y}) - (\gamma_\mathsf{x}\alpha_\mathsf{y} + \gamma_\mathsf{y}\alpha_\mathsf{x} - \Gamma_{\mathsf{xy}}) \qquad \text{[by Eq. 1]}$$
$$= (\beta_\mathsf{z} - \beta_\mathsf{x}\beta_\mathsf{y}) - \chi \qquad \text{[where } \chi = \gamma_\mathsf{x}\alpha_\mathsf{y} + \gamma_\mathsf{y}\alpha_\mathsf{x} - \Gamma_{\mathsf{xy}}]$$

Now if χ can be made available to P_0, it can send $\beta_\mathsf{z}^\star + \chi$ to P_1 and P_2 who using the knowledge of $\beta_\mathsf{x}, \beta_\mathsf{y}$, can verify the correctness of β_z by computing $\beta_\mathsf{z} - \beta_\mathsf{x}\beta_\mathsf{y}$ and checking against the value $\beta_\mathsf{z}^\star + \chi$ received from P_0. However, disclosing χ on clear to P_0 will cause a privacy issue when P_0 is corrupt, because one degree of freedom on the pair $(\gamma_\mathsf{x}, \gamma_\mathsf{y})$ is lost and the same impact percolates down to $(\beta_\mathsf{x}, \beta_\mathsf{y})$ and further to the actual values $(\mathsf{v}_\mathsf{x}, \mathsf{v}_\mathsf{y})$ on the wires x, y. This is resolved through a random value $\psi \in \mathbb{Z}_{2^\ell}$, sampled together

by P_1 and P_2. Now, χ and $\beta_z^\star$ are set to $\gamma_x\alpha_y + \gamma_y\alpha_x - \Gamma_{xy} + \psi$, $(\beta_z - \beta_x\beta_y + \psi) - \chi$, respectively and the check by P_1, P_2 involves computing $\beta_z - \beta_x\beta_y + \psi$. The rest of the logic in BLAZE goes on to discuss how to enforce P_0– (a) to compute a correct χ (when honest), and (b) to share correct Γ_{xy} (when corrupt). Tying the ends together, they identify the precise shared multiplication triple and map its components to χ and Γ_{xy} so that these values are correct by virtue of the correctness of the product relation. This reduces ensuring the correctness of these values to doing a single multiplication of two values in the preprocessing phase.

Protocol $\Pi_{\mathsf{mult}}(\mathcal{P}, [\![x]\!], [\![y]\!])$

Preprocessing:

– P_0, P_j for $j \in \{1,2\}$ together sample random $[\alpha_z]_j \in \mathbb{Z}_{2^\ell}$, while P_1, P_2 sample random $\gamma_z \in \mathbb{Z}_{2^\ell}$.
– Servers in $\mathcal{P}$ locally compute $\langle\cdot\rangle$-sharing of $\mathsf{d} = \gamma_x + \alpha_x$ and $\mathsf{e} = \gamma_y + \alpha_y$ by setting the shares as follows (ref. Table 2):

$$(\mathsf{d}_0 = [\alpha_x]_2, \mathsf{d}_1 = [\alpha_x]_1, \mathsf{d}_2 = \gamma_x), \quad (\mathsf{e}_0 = [\alpha_y]_2, \mathsf{e}_1 = [\alpha_y]_1, \mathsf{e}_2 = \gamma_y)$$

– Servers in $\mathcal{P}$ execute $\Pi_{\mathsf{mulPre}}(\mathcal{P}, \mathsf{d}, \mathsf{e})$ to generate $\langle\mathsf{f}\rangle = \langle\mathsf{de}\rangle$.
– P_0, P_1 locally set $[\chi]_1 = \mathsf{f}_1$, while P_0, P_2 locally set $[\chi]_2 = \mathsf{f}_0$. P_1, P_2 locally compute $\psi = \mathsf{f}_2 - \gamma_x\gamma_y$.

Online:

– P_0, P_j, for $j \in \{1,2\}$, compute $[\beta_z^\star]_j = -(\beta_x + \gamma_x)[\alpha_y]_j - (\beta_y + \gamma_y)[\alpha_x]_j + [\alpha_z]_j + [\chi]_j$.
– P_0, P_1 jmp-send $[\beta_z^\star]_1$ to P_2 and P_0, P_2 jmp-send $[\beta_z^\star]_2$ to P_1.
– P_1, P_2 compute $\beta_z^\star = [\beta_z^\star]_1 + [\beta_z^\star]_2$ and set $\beta_z = \beta_z^\star + \beta_x\beta_y + \psi$.
– P_1, P_2 jmp-send $\beta_z + \gamma_z$ to P_0.

Figure 2: 3PC: Multiplication Protocol ($z = x \cdot y$)

We differ from BLAZE in several ways. First, we do not simply rely on P_0 for the verification information $\beta_z^\star + \chi$, as this may inevitably lead to abort when P_0 is corrupt. Instead, we find (a slightly different) $\beta_z^\star$ that, instead of entirely available to P_0, will be available in $[\cdot]$-shared form between the two teams $\{P_0, P_1\}, \{P_0, P_2\}$, with both servers in $\{P_0, P_i\}$ holding ith share $[\beta_z^\star]_i$. With this edit, the ith team can jmp-send the ith share of $\beta_z^\star$ to the third server which computes $\beta_z^\star$. Due to the presence of one honest server in each team, this $\beta_z^\star$ is correct and P_1, P_2 directly use it to compute β_z, with the knowledge of ψ, β_x, β_y. The outcome of our approach is a win-win situation i.e. either success or TTP selection. Our approach of computing β_z from $\beta_z^\star$ is a departure from BLAZE, where the latter suggests computing β_z from the exchange P_1, P_2's respective share of β_z (as in the semi-honest construction) and use $\beta_z^\star$ for verification. Our new $\beta_z^\star$ and χ are:

$$\chi = \gamma_x\alpha_y + \gamma_y\alpha_x + \Gamma_{xy} - \psi \quad \text{and}$$
$$\beta_z^\star = -(\beta_x + \gamma_x)\alpha_y - (\beta_y + \gamma_y)\alpha_x + \alpha_z + \chi$$
$$= (-\beta_x\alpha_y - \beta_y\alpha_x + \Gamma_{xy} + \alpha_z) - \psi = \beta_z - \beta_x\beta_y - \psi$$

Clearly, both P_0 and P_i can compute $[\beta_z^\star]_i = -(\beta_x + \gamma_x)[\alpha_y]_i - (\beta_y + \gamma_y)[\alpha_x]_i + [\alpha_z]_i + [\chi]_i$ given $[\chi]_i$. The rest of

our discussion explains how (a) ith share of $[\chi]$ can be made available to $\{P_0, P_i\}$ and (b) ψ can be derived by P_1, P_2, from a multiplication triple. Similar to BLAZE, yet for a different triple, we observe that $(\mathsf{d}, \mathsf{e}, \mathsf{f})$ is a multiplication triple, where $\mathsf{d} = (\gamma_x + \alpha_x), \mathsf{e} = (\gamma_y + \alpha_y), \mathsf{f} = (\gamma_x\gamma_y + \psi) + \chi$ if and only if χ and Γ_{xy} are correct. Indeed,

$$\mathsf{de} = (\gamma_x + \alpha_x)(\gamma_y + \alpha_y) = \gamma_x\gamma_y + \gamma_x\alpha_y + \gamma_y\alpha_x + \Gamma_{xy}$$
$$= (\gamma_x\gamma_y + \psi) + (\gamma_x\alpha_y + \gamma_y\alpha_x + \Gamma_{xy} - \psi)$$
$$= (\gamma_x\gamma_y + \psi) + \chi = \mathsf{f}$$

Based on this observation, we compute the above multiplication triple using a multiplication protocol and extract out the values for ψ and χ from the shares of f which are bound to be correct. This can be executed entirely in the preprocessing phase. Specifically, the servers (a) locally obtain $\langle\cdot\rangle$-shares of d, e as in Table 2, (b) compute $\langle\cdot\rangle$-shares of $\mathsf{f}(= \mathsf{de})$, say denoted by $\mathsf{f}_0, \mathsf{f}_1, \mathsf{f}_2$, using an efficient, robust 3-party multiplication protocol, say Π_{mulPre} (abstracted in a functionality Fig. 11) and finally (c) extract out the required preprocessing data *locally* as in Eq. 2. We switch to $\langle\cdot\rangle$-sharing in this part to be able to use the best robust multiplication protocol of [8] that supports this form of secret sharing and requires communication of just 3 elements. Fortunately, the switch does not cost anything, as both the step (a) and (c) (as above) involve local computation and the cost simply reduces to a single run of a multiplication protocol.

$\langle\mathsf{v}\rangle$	P_0 $(\mathsf{v}_0, \mathsf{v}_1)$	P_1 $(\mathsf{v}_1, \mathsf{v}_2)$	P_2 $(\mathsf{v}_2, \mathsf{v}_0)$
$\langle\mathsf{d}\rangle$	$([\alpha_x]_2, [\alpha_x]_1)$	$([\alpha_x]_1, \gamma_x)$	$(\gamma_x, [\alpha_x]_2)$
$\langle\mathsf{e}\rangle$	$([\alpha_y]_2, [\alpha_y]_1)$	$([\alpha_y]_1, \gamma_y)$	$(\gamma_y, [\alpha_y]_2)$

Table 2: The $\langle\cdot\rangle$-sharing of values d and e

$$[\chi]_2 \leftarrow \mathsf{f}_0, \quad [\chi]_1 \leftarrow \mathsf{f}_1, \quad \gamma_x\gamma_y + \psi \leftarrow \mathsf{f}_2. \qquad (2)$$

According to $\langle\cdot\rangle$-sharing, both P_0 and P_1 obtain f_1 and hence obtain $[\chi]_1$. Similarly, P_0, P_2 obtain f_0 and hence $[\chi]_2$. Finally, P_1, P_2 obtain f_2 from which they compute $\psi = \mathsf{f}_2 - \gamma_x\gamma_y$. This completes the informal discussion.

We note that to facilitate a fast online phase for multiplication, our preprocessing phase leverages a robust multiplication protocol [8] in a black-box manner to derive the necessary preprocessing information. A similar black-box approach is also taken for the dot product protocol in the preprocessing phase. This leaves room for further improvements in the communication cost, which can be obtained by instantiating the black-box with an efficient, robust protocol coupled with the fast online phase.

Reconstruction Protocol Protocol Π_{rec} (Fig. 12) allows servers to robustly reconstruct value $\mathsf{v} \in \mathbb{Z}_{2^\ell}$ from its $[\![\cdot]\!]$-shares. Note that each server misses one share of v which is held by the other two servers. Consider the case of P_0 who requires γ_v to compute v. During the preprocessing, P_1, P_2 compute a commitment of γ_v, denoted by $\mathsf{Com}(\gamma_\mathsf{v})$ and jmp-send

the same to P_0. Similar steps are performed for the values $[\alpha_v]_2$ and $[\alpha_v]_1$ that are required by servers P_1 and P_2 respectively. During the online phase, servers open their commitments to the intended server who accepts the opening that is consistent with the agreed upon commitment. We defer the details to §B.5.

The Complete 3PC For the sake of completeness and to demonstrate how GOD is achieved, we show how to compile the above primitives for a general 3PC. A similar approach will be taken for 4PC and each PPML task, and we will avoid repetition. In order to compute an arithmetic circuit over $\mathbb{Z}_{2^\ell}$, we first invoke the key-setup functionality $\mathcal{F}_{\text{setup}}$ (Fig. 6) for key distribution and preprocessing of Π_{sh}, Π_{mult} and Π_{rec}, as per the given circuit. During the online phase, $P_i \in \mathcal{P}$ shares its input x_i by executing online steps of Π_{sh} (Fig. 9). This is followed by the circuit evaluation phase, where severs evaluate the gates in the circuit in the topological order, with addition gates (and multiplication-by-a-constant gates) being computed locally, and multiplication gates being computed via online of Π_{mult} (Fig. 2). Finally, servers run the online steps of Π_{rec} (Fig. 12) on the output wires to reconstruct the function output. To leverage amortization, only *send* phases of all the jmp are run on the flow. At the end of preprocessing, the *verify* phase for all possible ordered pair of senders are run. We carry on computation in the online phase only when the *verify* phases in the preprocessing are successful. Otherwise, the servers simply send their inputs to the elected TTP, who computes the function and returns the result to all the servers. Similarly, depending on the output of the *verify* at the end of the online phase, either the reconstruction is carried out or a TTP is identified. In the latter case, computation completes as mentioned before.

3.3 Building Blocks for PPML using 3PC

This section provides details on robust realizations of the following building blocks for PPML in 3-server setting– i) Dot Product, ii) Truncation, iii) Dot Product with Truncation, iv) Secure Comparison, and v) Non-linear Activation functions– Sigmoid and ReLU. We provide the communication analysis of our protocols in §B and security proofs in §D.1. We begin by providing details of input sharing and reconstruction in the SOC setting.

Input Sharing and Output Reconstruction in the SOC Setting Protocol $\Pi_{\text{sh}}^{\text{SOC}}$ (Fig. 3) extends input sharing to the SOC setting and allows a user U to generate the $[\![\cdot]\!]$-shares of its input v among the three servers. Note that the necessary commitments to facilitate the sharing are generated in the preprocessing phase by the servers which are then communicated to U, along with the opening, in the online phase. U selects the commitment forming the majority (for each share) owing to the presence of an honest majority among the servers, and accepts the corresponding shares. Analogously, protocol $\Pi_{\text{rec}}^{\text{SOC}}$ (Fig. 3) allows the servers to reconstruct a value v

towards user U. In either of the protocols, if at any point, a TTP is identified, then servers signal the TTP's identity to U. U selects the TTP as the one forming a majority and sends its input in the clear to the TTP, who computes the function output and sends it back to U.
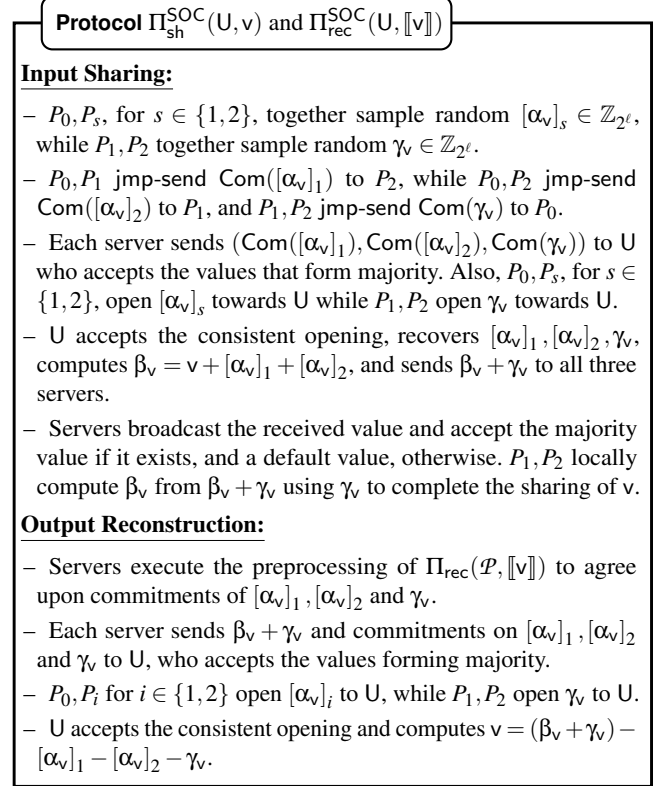

Protocol $\Pi_{\text{sh}}^{\text{SOC}}(U, v)$ and $\Pi_{\text{rec}}^{\text{SOC}}(U, [\![v]\!])$

Input Sharing:

– P_0, P_s, for $s \in \{1,2\}$, together sample random $[\alpha_v]_s \in \mathbb{Z}_{2^\ell}$, while P_1, P_2 together sample random $\gamma_v \in \mathbb{Z}_{2^\ell}$.

– P_0, P_1 jmp-send $\text{Com}([\alpha_v]_1)$ to P_2, while P_0, P_2 jmp-send $\text{Com}([\alpha_v]_2)$ to P_1, and P_1, P_2 jmp-send $\text{Com}(\gamma_v)$ to P_0.

– Each server sends $(\text{Com}([\alpha_v]_1), \text{Com}([\alpha_v]_2), \text{Com}(\gamma_v))$ to U who accepts the values that form majority. Also, P_0, P_s, for $s \in \{1,2\}$, open $[\alpha_v]_s$ towards U while P_1, P_2 open γ_v towards U.

– U accepts the consistent opening, recovers $[\alpha_v]_1, [\alpha_v]_2, \gamma_v$, computes $\beta_v = v + [\alpha_v]_1 + [\alpha_v]_2$, and sends $\beta_v + \gamma_v$ to all three servers.

– Servers broadcast the received value and accept the majority value if it exists, and a default value, otherwise. P_1, P_2 locally compute β_v from $\beta_v + \gamma_v$ using γ_v to complete the sharing of v.

Output Reconstruction:

– Servers execute the preprocessing of $\Pi_{\text{rec}}(\mathcal{P}, [\![v]\!])$ to agree upon commitments of $[\alpha_v]_1, [\alpha_v]_2$ and γ_v.

– Each server sends $\beta_v + \gamma_v$ and commitments on $[\alpha_v]_1, [\alpha_v]_2$ and γ_v to U, who accepts the values forming majority.

– P_0, P_i for $i \in \{1,2\}$ open $[\alpha_v]_i$ to U, while P_1, P_2 open γ_v to U.

– U accepts the consistent opening and computes $v = (\beta_v + \gamma_v) - [\alpha_v]_1 - [\alpha_v]_2 - \gamma_v$.

Figure 3: 3PC: Input Sharing and Output Reconstruction

MSB Extraction, Bit to Arithmetic Conversion and Bit Injection Protocols We provide a high-level overview of three protocols that involve working over arithmetic and boolean rings in a mixed fashion and are used in PPML primitives. The *bit extraction* protocol, Π_{bitext} (§B.6.1) allows servers to compute *boolean* sharing of the most significant bit (msb) of a value v from its arithmetic sharing ($[\![v]\!]$). The *Bit2A* protocol, Π_{bit2A} (§B.6.2), given the boolean sharing of a bit b, denoted as $[\![b]\!]^{\mathbf{B}}$, allows the servers to compute the arithmetic sharing $[\![b^{\mathbf{R}}]\!]$. Here $b^{\mathbf{R}}$ denotes the equivalent value of b over ring $\mathbb{Z}_{2^\ell}$ (see Notation 2.2). Lastly, *Bit Injection* protocol, Π_{BitInj} (§B.6.3), allows servers to compute the arithmetic sharing $[\![bv]\!]$ from boolean sharing of a bit b ($[\![b]\!]^{\mathbf{B}}$) and arithmetic sharing of v ($[\![v]\!]$).

The core techniques used in these protocols follow from BLAZE [44], where multiplication calls are replaced with our new Π_{mult}, and several private communications are replaced with jmp-send to ensure a successful run or TTP selection. These PPML building-blocks can be understood without details of the constructs and hence they are moved to §B.6.

Dot Product Given the $[\![\cdot]\!]$-sharing of vectors $\vec{x}$ and $\vec{y}$, protocol Π_{dotp} (Fig. 14) allows servers to generate $[\![\cdot]\!]$-sharing of

$z = \vec{x} \odot \vec{y}$ robustly. $\llbracket \cdot \rrbracket$-sharing of a vector $\vec{x}$ of size n, means that each element $x_i \in \mathbb{Z}_{2^\ell}$ of $\vec{x}$, for $i \in [n]$, is $\llbracket \cdot \rrbracket$-shared. We borrow ideas from BLAZE for obtaining an online communication cost *independent* of n and use jmp primitive to ensure either success or TTP selection. Analogous to our multiplication protocol, our dot product offloads one call to a robust dot product protocol to the preprocessing. By extending techniques of [7, 8], we give an instantiation for the dot product protocol used in our preprocessing whose (amortized) communication cost is constant, thereby making our preprocessing cost also *independent* of n.

To begin with, $z = \vec{x} \odot \vec{y}$ can be viewed as n parallel multiplication instances of the form $z_i = x_i y_i$ for $i \in [n]$, followed by adding up the results. Let $\beta_z^\star = \sum_{i=1}^n \beta_{z_i}^\star$. Then,

$$\beta_z^\star = -\sum_{i=1}^n (\beta_{x_i} + \gamma_{x_i})\alpha_{y_i} - \sum_{i=1}^n (\beta_{y_i} + \gamma_{y_i})\alpha_{x_i} + \alpha_z + \chi \quad (3)$$

where $\chi = \sum_{i=1}^n (\gamma_{x_i}\alpha_{y_i} + \gamma_{y_i}\alpha_{x_i} + \Gamma_{x_i y_i} - \psi_i)$.

Apart from the aforementioned modification, the online phase for dot product proceeds similar to that of multiplication protocol. P_0, P_1 locally compute $[\beta_z^\star]_1$ as per Eq. 3 and jmp-send $[\beta_z^\star]_1$ to P_2. P_1 obtains $[\beta_z^\star]_2$ in a similar fashion. P_1, P_2 reconstruct $\beta_z^\star = [\beta_z^\star]_1 + [\beta_z^\star]_2$ and compute $\beta_z = \beta_z^\star + \sum_{i=1}^n \beta_{x_i}\beta_{y_i} + \psi$. Here, the value ψ has to be correctly generated in the preprocessing phase satisfying Eq. 3. Finally, P_1, P_2 jmp-send $\beta_z + \gamma_z$ to P_0.

We now provide the details for preprocessing phase that enable servers to obtain the required values (χ, ψ) with the invocation of a dot product protocol in a black-box way. Towards this, let $\vec{d} = [d_1, \ldots, d_n]$ and $\vec{e} = [e_1, \ldots, e_n]$, where $d_i = \gamma_{x_i} + \alpha_{x_i}$ and $e_i = \gamma_{y_i} + \alpha_{y_i}$ for $i \in [n]$, as in the case of multiplication. Then for $f = \vec{d} \odot \vec{e}$,

$$f = \vec{d} \odot \vec{e} = \sum_{i=1}^n d_i e_i = \sum_{i=1}^n (\gamma_{x_i} + \alpha_{x_i})(\gamma_{y_i} + \alpha_{y_i})$$
$$= \sum_{i=1}^n (\gamma_{x_i}\gamma_{y_i} + \psi_i) + \sum_{i=1}^n \chi_i = \sum_{i=1}^n (\gamma_{x_i}\gamma_{y_i} + \psi_i) + \chi$$
$$= \sum_{i=1}^n (\gamma_{x_i}\gamma_{y_i} + \psi_i) + [\chi]_1 + [\chi]_2 = f_2 + f_1 + f_0.$$

where $f_2 = \sum_{i=1}^n (\gamma_{x_i}\gamma_{y_i} + \psi_i), f_1 = [\chi]_1$ and $f_0 = [\chi]_2$.

Using the above relation, the preprocessing phase proceeds as follows: P_0, P_j for $j \in \{1, 2\}$ sample a random $[\alpha_z]_j \in \mathbb{Z}_{2^\ell}$, while P_1, P_2 sample random γ_z. Servers locally prepare $\langle \vec{d} \rangle, \langle \vec{e} \rangle$ similar to that of multiplication protocol. Servers then execute a robust 3PC dot product protocol, denoted by Π_{dotpPre}, that takes $\langle \vec{d} \rangle, \langle \vec{e} \rangle$ as input and compute $\langle f \rangle$ with $f = \vec{d} \odot \vec{e}$. Given $\langle f \rangle$, the ψ and $[\chi]$ values are extracted as follows (ref. Eq. 4):

$$\psi = f_2 - \sum_{i=1}^n \gamma_{x_i}\gamma_{y_i}, \quad [\chi]_1 = f_1, \quad [\chi]_2 = f_0, \quad (4)$$

It is easy to see from the semantics of $\langle \cdot \rangle$-sharing that both P_1, P_2 obtain f_2 and hence ψ. Similarly, both P_0, P_1 obtain f_1 and hence $[\chi]_1$, while P_0, P_2 obtain $[\chi]_2$.

A trivial way to instantiate Π_{dotpPre} is to treat a dot product operation as n multiplications. However, this results in a communication cost that is linearly dependent on the feature size. Instead, we instantiate Π_{dotpPre} by a semi-honest dot product protocol followed by a verification phase to check the correctness. For the verification phase, we extend the techniques of [7, 8] to provide support for verification of dot product tuples. Setting the verification phase parameters appropriately gives a Π_{dotpPre} whose (amortized) communication cost is independent of the feature size. Details appear in §B.7.

Truncation Working over fixed-point values, repeated multiplications using FPA arithmetic can lead to an overflow resulting in loss of significant bits of information. This put forth the need for truncation [10, 13, 39, 41, 44] that re-adjusts the shares after multiplication so that FPA semantics are maintained. As shown in SecureML [41], the method of truncation would result in loss of information on the least significant bits and affect the accuracy by a very minimal amount.

For truncation, servers execute Π_{trgen} (Fig. 16) to generate $([r], \llbracket r^d \rrbracket)$-pair, where r is a random ring element, and r^d is the truncated value of r, i.e the value r right-shifted by d bit positions. Recall that d denotes the number of bits allocated for the fractional part in the FPA representation. Given (r, r^d), the truncated value of v, denoted as v^d, is computed as $v^d = (v - r)^d + r^d$. The correctness and accuracy of this method was shown in ABY3 [39].

Protocol Π_{trgen} is inspired from [14, 39] and proceeds as follows to generate $([r], \llbracket r^d \rrbracket)$. Analogous to the approach of ABY3 [39], servers generate a boolean sharing of an ℓ-bit value $r = r_1 \oplus r_2$, non-interactively. Each server truncates its share of r locally to obtain a boolean sharing of r^d by removing the lower d bits. To obtain the arithmetic shares of (r, r^d) from their boolean sharing, we do not, however, rely on the approach of ABY3 as it requires more rounds. Instead, we implicitly perform a *boolean to arithmetic conversion*, as was proposed in Trident [14], to obtain the arithmetic shares of (r, r^d). This entails performing two dot product operations and constitutes the cost for Π_{trgen}. We defer details to §B.8.

Dot Product with Truncation Given the $\llbracket \cdot \rrbracket$-sharing of vectors $\vec{x}$ and $\vec{y}$, protocol Π_{dotpt} (Fig. 17) allows servers to generate $\llbracket z^d \rrbracket$, where z^d denotes the truncated value of $z = \vec{x} \odot \vec{y}$. A naive way is to compute the dot product using Π_{dotp}, followed by performing truncation using the (r, r^d) pair. Instead, we follow the optimization of BLAZE where the online phase of Π_{dotp} is modified to integrate the truncation using (r, r^d) at no additional cost.

The preprocessing phase now consists of the execution of one instance of Π_{trgen} (Fig. 16) and the preprocessing corresponding to Π_{dotp} (Fig. 14). In the online phase, servers enable P_1, P_2 to obtain $z^\star - r$ instead of $\beta_z^\star$, where $z^\star = \beta_z^\star - \alpha_z$. Using $z^\star - r$, both P_1, P_2 then compute $(z - r)$

locally, truncate it to obtain $(z-r)^d$ and execute Π_{jsh} to generate $[\![(z-r)^d]\!]$. Finally, servers locally compute the result as $[\![z^d]\!] = [\![(z-r)^d]\!] + [\![r^d]\!]$. We defer the details to §B.9.

Secure Comparison Secure comparison allows servers to check whether $x < y$, given their $[\![\cdot]\!]$-shares. In FPA representation, checking $x < y$ is equivalent to checking the msb of $v = x - y$. Towards this, servers locally compute $[\![v]\!] = [\![x]\!] - [\![y]\!]$ and extract the msb of v using Π_{bitext} (§B.6.1). In case an arithmetic sharing is desired, servers can apply Π_{bit2A} (Fig. 13) protocol on the outcome of Π_{bitext} protocol.

Activation Functions We now elaborate on two of the most prominently used activation functions: i) Rectified Linear Unit (ReLU) and (ii) Sigmoid (Sig).

  – *ReLU:* The ReLU function, $\text{relu}(v) = \max(0, v)$, can be viewed 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 $[\![v]\!]$, servers execute Π_{bitext} on $[\![v]\!]$ to generate $[\![b]\!]^{\mathbf{B}}$. $[\![\cdot]\!]^{\mathbf{B}}$-sharing of $\bar{b}$ is locally computed by setting $\beta_{\bar{b}} = 1 \oplus \beta_b$. Servers execute Π_{BitInj} protocol on $[\![\bar{b}]\!]^{\mathbf{B}}$ and $[\![v]\!]$ to obtain the desired result.

  – *Sig:* In this work, we use the MPC-friendly variant of the Sigmoid function [13, 39, 41] (ref. §B.10). Note that $\text{sig}(v) = \bar{b_1}b_2(v + 1/2) + \bar{b_2}$, where $b_1 = 1$ if $v + 1/2 < 0$ and $b_2 = 1$ if $v - 1/2 < 0$. To compute $[\![\text{sig}(v)]\!]$, servers proceed in a similar fashion as in ReLU, and hence, we skip the details.

Maxpool, Convolution and Matrix Multiplication The goal of maxpool is to find the maximum value in a vector $\vec{x}$ of m values. Maximum between two elements x_i, x_j can be computed by applying secure comparison, which returns a binary sharing of a bit b such that $b = 0$ if $x_i > x_j$, or 1, otherwise, followed by computing $(b)^{\mathbf{B}}(x_j - x_i) + x_i$, which can be performed using bit injection (§B.6.3). To find the maximum value in vector $\vec{x}$, the servers first group the values in $\vec{x}$ into pairs and securely compare each pair to obtain the maximum of the two. This results in a vector of size $m/2$. This process is repeated for $\mathrm{O}(\log m)$ rounds to obtain the maximum value in the entire vector.

Convolutions, which form another important building block in PPML tasks, can be cast into matrix multiplication. Our protocol to compute a matrix of dimension $p \times r$ after multiplication requires only $p \times r$ multiplications, and is explained in §B.11.

4   Robust 4PC and PPML

In this section, we extend our 3PC results to the 4-party case and observe substantial efficiency gain. First, the use of broadcast is eliminated. Second, the preprocessing of multiplication becomes substantially computationally light, eliminating the multiplication protocol (used in the preprocessing) altogether. Third, we achieve a dot product protocol with communication cost independent of the size of the vector, completely eliminating the complex machinery required as in the 3PC case. At the heart of our 4PC constructions lies an efficient 4-party

jmp primitive, denoted as jmp4, that allows two servers to send a common value to a third server robustly. We start with the secret-sharing semantics for 4 servers. We only use an extended version of $[\![\cdot]\!]$-sharing defined below.

Secret Sharing Semantics For a value v, the shares for P_0, P_1 and P_2 remain the same as that for 3PC case. That is, P_0 holds $([\alpha_v]_1, [\alpha_v]_2, \beta_v + \gamma_v)$ while P_i for $i \in \{1, 2\}$ holds $([\alpha_v]_i, \beta_v, \gamma_v)$. The shares for the fourth server P_3 is defined as $([\alpha_v]_1, [\alpha_v]_2, \gamma_v)$. Clearly, the secret is defined as $v = \beta_v - [\alpha_v]_1 - [\alpha_v]_2$.

4PC Joint Message Passing Primitive The jmp4 primitive enables two servers P_i, P_j to send a common value $v \in \mathbb{Z}_{2^\ell}$ to a third server P_k, or identify a TTP in case of any inconsistency. This primitive is analogous to jmp (Fig. 1) in spirit but is significantly optimized and free from broadcast calls. Similar to the 3PC counterpart, each server maintains a bit and P_i sends the value, and P_j the hash of it to P_k. P_k sets its inconsistency bit to 1 when the (value, hash) pair is inconsistent. This is followed by relaying the bit to all the servers, who exchange it among themselves and agree on the bit that forms majority (1 indicates the presence of inconsistency, and 0 indicates consistency). The presence of an honest majority among P_i, P_j, P_l, guarantees agreement on the presence/absence of an inconsistency as conveyed by P_k. Observe that inconsistency can only be caused either due to a corrupt sender sending an incorrect value (or hash), or a corrupt receiver falsely announcing the presence of inconsistency. Hence, the fourth server, P_l, can safely be employed as TTP. The protocol appears in Fig. 4.

Protocol $\Pi_{\text{jmp4}}(P_i, P_j, P_k, v, P_l)$

– $P_s \in \mathcal{P}$ initializes an inconsistency bit $b_s = 0$. If P_s remains silent instead of sending b_s in any of the following rounds, the recipient sets b_s to 1.

– P_i sends v to P_k and P_j sends $\mathsf{H}(v)$ to P_k. P_k sets $b_k = 1$ if the received values are inconsistent or if the value is not received.

– P_k sends b_k to all servers. P_s for $s \in \{i, j, l\}$ sets $b_s = b_k$.

– P_s for $s \in \{i, j, l\}$ mutually exchange their bits. P_s resets $b_s = b'$ where b' denotes the bit which appears in majority among b_i, b_j, b_l.

– All servers set $\mathsf{TTP} = P_l$ if $b' = 1$, terminate otherwise.

Figure 4: 4PC: Joint Message Passing Primitive

Notation 4.1. We say that P_i, P_j jmp4-send v to P_k when they invoke $\Pi_{\text{jmp4}}(P_i, P_j, P_k, v, P_l)$.

We note that the end goal of jmp4 primitive relates closely to the bi-convey primitive of FLASH [10]. Bi-convey allows two servers S_1, S_2 to convey a value to a server R, and in case of an inconsistency, a pair of honest servers mutually identify each other, followed by exchanging their internal randomness to recover the clear inputs, computing the circuit, and sending the output to all. Note, however, that jmp4 primitive is more efficient and differs significantly in techniques from the bi-convey primitive. Unlike in bi-convey, in case of an

inconsistency, jmp4 enables servers to learn the TTP's identity unanimously. Moreover, bi-convey demands that honest servers, identified during an inconsistency, exchange their internal randomness (which comprises of the shared keys established during the key-setup phase) to proceed with the computation. This enforces the need for a fresh key-setup every time inconsistency is detected. On the efficiency front, jmp4 simply halves the communication cost of bi-convey, giving a $2\times$ improvement.

4.1 4PC Protocols

In this section, we revisit the protocols from 3PC (§3) and suggest optimizations leveraging the presence of an additional honest party in the system. While we provide details for the protocols that vary significantly from their 3PC counterpart in this section, the details for other protocols along with the communication analysis are deferred to §C. We provide security proofs in §D.2.

Sharing Protocol To enable P_i to share a value v, protocol Π_{sh4} (Fig. 19) proceeds similar to that of 3PC case with the addition that P_3 also samples the values $[\alpha_v]_1, [\alpha_v]_2, \gamma_v$ using the shared randomness with the respective servers. On a high level, P_i computes $\beta_v = v + [\alpha_v]_1 + [\alpha_v]_2$ and sends β_v (or $\beta_v + \gamma_v$) to another server and they together jmp4-send this information to the intended servers. We defer the formal details of Π_{sh4} to §C.2.

Protocol $\Pi_{\text{mult4}}(\mathcal{P}, [\![x]\!], [\![y]\!])$

Preprocessing:

– P_0, P_3, P_j, for $j \in \{1, 2\}$, sample random $[\alpha_z]_j \in \mathbb{Z}_{2^\ell}$, and P_0, P_1, P_3 sample random $[\Gamma_{xy}]_1 \in \mathbb{Z}_{2^\ell}$.

– P_1, P_2, P_3 sample random $\gamma_z, \psi, r \in \mathbb{Z}_{2^\ell}$ and set $[\psi]_1 = r, [\psi]_2 = \psi - r$.

– P_0, P_3 set $[\Gamma_{xy}]_2 = \Gamma_{xy} - [\Gamma_{xy}]_1$, where $\Gamma_{xy} = \alpha_x \alpha_y$. P_0, P_3 jmp4-send $[\Gamma_{xy}]_2$ to P_2.

– P_3, P_j, for $j \in \{1, 2\}$, set $[\chi]_j = \gamma_x [\alpha_y]_j + \gamma_y [\alpha_x]_j + [\Gamma_{xy}]_j - [\psi]_j$. P_1, P_3 jmp4-send $[\chi]_1$ to P_0, while P_2, P_3 jmp4-send $[\chi]_2$ to P_0.

Online:

– P_0, P_j, for $j \in \{1, 2\}$, compute $[\beta_z^\star]_j = -(\beta_x + \gamma_x)[\alpha_y]_j - (\beta_y + \gamma_y)[\alpha_x]_j + [\alpha_z]_j + [\chi]_j$.

– P_1, P_0 jmp4-send $[\beta_z^\star]_1$ to P_2, and P_2, P_0 jmp4-send $[\beta_z^\star]_2$ to P_1.

– P_j, for $j \in \{1, 2\}$, computes $\beta_z^\star = [\beta_z^\star]_1 + [\beta_z^\star]_2$ and sets $\beta_z = \beta_z^\star + \beta_x \beta_y + \psi$.

– P_1, P_2 jmp4-send $\beta_z + \gamma_z$ to P_0.

Figure 5: 4PC: Multiplication Protocol ($z = x \cdot y$)

Multiplication Protocol Given the $[\![\cdot]\!]$-shares of x and y, protocol Π_{mult4} (Fig. 5) allows servers to compute $[\![z]\!]$ with $z = xy$. When compared with the state-of-the-art 4PC GOD protocol of FLASH [10], our solution improves communica-

tion in both, the preprocessing and online phase, from 6 to 3 ring elements. Moreover, our communication cost matches with the state-of-the-art 4PC protocol of Trident [14] that only provides security with fairness.

Recall that the goal of preprocessing in 3PC multiplication was to enable P_1, P_2 obtain ψ, and P_0, P_i for $i \in \{1, 2\}$ obtain $[\chi]_i$ where $\chi = \gamma_x \alpha_y + \gamma_y \alpha_x + \Gamma_{xy} - \psi$. Here ψ is a random value known to both P_1, P_2. With the help of P_3, we let the servers obtain the respective preprocessing data as follows: P_0, P_3, P_1 together samples random $[\Gamma_{xy}]_1 \in \mathbb{Z}_{2^\ell}$. P_0, P_3 locally compute $\Gamma_{xy} = \alpha_x \alpha_y$, set $[\Gamma_{xy}]_2 = \Gamma_{xy} - [\Gamma_{xy}]_1$ and jmp4-send $[\Gamma_{xy}]_2$ to P_2. P_1, P_2, P_3 locally sample ψ, r and generate $[\cdot]$-shares of ψ by setting $[\psi]_1 = r$ and $[\psi]_2 = \psi - r$. Then P_j, P_3 for $j \in \{1, 2\}$ compute $[\chi]_j = \gamma_x [\alpha_y]_j + \gamma_y [\alpha_x]_j + [\Gamma_{xy}]_j - [\psi]_j$ and jmp4-send $[\chi]_j$ to P_0. The online phase is similar to that of 3PC, apart from Π_{jmp4} being used instead of Π_{jmp} for communication. Since P_3 is not involved in the online computation phase, we can safely assume P_3 to serve as the TTP for the Π_{jmp4} executions in the online phase.

Reconstruction Protocol Given $[\![v]\!]$, protocol Π_{rec4} (Fig. 22) enables servers to robustly reconstruct the value v among the servers. Note that every server lacks one share for reconstruction and the same is available with three other servers. Hence, they communicate the missing share among themselves, and the majority value is accepted. As an optimization, two among the three servers can send the missing share while the third one can send a hash of the same for verification. We defer the formal details of Π_{rec4} to §C.6. Notice that, as opposed to the 3PC case, this protocol does not require commitments.

Input Sharing and Output Reconstruction in SOC Setting We extend input sharing and reconstruction in the SOC setting as follows. To generate $[\![\cdot]\!]$-shares for its input v, U receives each of the shares $[\alpha_v]_1, [\alpha_v]_2$, and γ_v from three out of the four servers as well as a random value $r \in \mathbb{Z}_{2^\ell}$ sampled together by P_0, P_1, P_2 and accepts the values that form the majority. U locally computes $u = v + [\alpha_v]_1 + [\alpha_v]_2 + \gamma_v + r$ and sends u to all the servers. Servers then execute a two round byzantine agreement (BA) [45] to agree on u (or $\bot$). We refer the readers to [45] for the formal details of the agreement protocol. On successful completion of BA, P_0 computes $\beta_v + \gamma_v$ from u while P_1, P_2 compute β_v from u locally. For the reconstruction of a value v, servers send their $[\![\cdot]\!]$-shares of v to U, who selects the majority value for each share and reconstructs the output. At any point, if a TTP is identified, the servers proceed as follows. All servers send their $[\![\cdot]\!]$-share of the input to the TTP. TTP picks the majority value for each share and computes the function output. It then sends this output to U. U also receives the identity of the TTP from all servers and accepts the output received from the TTP forming majority.

Dot Product Given $[\![\cdot]\!]$-shares of two n-sized vectors $\vec{x}, \vec{y}$, protocol Π_{dotp4} (Fig. 25) enables servers to compute $[\![z]\!]$ with $z = \vec{x} \odot \vec{y}$. The protocol is essentially similar to n instances of

multiplications of the form $z_i = x_i y_i$ for $i \in [n]$. But instead of communicating values corresponding to each of the n instances, servers locally sum up the shares and communicate a single value. This helps to obtain a communication cost independent of the size of the vectors. Details appear in §C.8.

5 Applications and Benchmarking

In this section, we empirically show the practicality of our protocols for PPML. We consider training and inference for Logistic Regression, and inference for 3 different Neural Networks (NN). NN training requires additional tools to allow mixed world computations, which we leave as future work. We refer readers to SecureML [41], ABY3 [39], BLAZE [44], FALCON [53] for a detailed description of the training and inference steps for the aforementioned ML algorithms. All our benchmarking is done over the publicly available MNIST [36] and CIFAR-10 [34] dataset. For training, we use a batch size of $B = 128$ and define 1 KB = 8192 bits.

In 3PC, we compare our results against the best-known framework BLAZE that provides fairness in the same setting. We observe that the technique of making the dot product cost independent of feature size can also be applied to BLAZE to obtain better costs. Hence, for a fair comparison, we additionally report these improved values for BLAZE. Further, we only consider the PPA circuit based variant of bit extraction for BLAZE since we aim for high throughput; the GC based variant results in huge communication and is not efficient for deep NNs. Our results imply that we get GOD at no additional cost compared to BLAZE. For 4PC, we compare our results with two best-known works FLASH [10] (which is robust) and Trident [14] (which is fair). Our results halve the cost of FLASH and are on par with Trident.

Benchmarking Environment We use a 64-bit ring ($\mathbb{Z}_{2^{64}}$). The benchmarking is performed over a WAN that was instantiated using n1-standard-8 instances of Google Cloud[5], with machines located in East Australia (P_0), South Asia (P_1), South East Asia (P_2), and West Europe (P_3). The machines are equipped with 2.3 GHz Intel Xeon E5 v3 (Haswell) processors supporting hyper-threading, with 8 vCPUs, and 30 GB of RAM Memory and with a bandwidth of 40 Mbps. The average round-trip time (rtt) was taken as the time for communicating 1 KB of data between a pair of parties, and the rtt values were as follows.

P_0-P_1	P_0-P_2	P_0-P_3	P_1-P_2	P_1-P_3	P_2-P_3
151.40ms	59.95ms	275.02ms	92.94ms	173.93ms	219.37ms

Software Details We implement our protocols[6] using the publicly available ENCRYPTO library [20] in C++17. We

obtained the code of BLAZE and FLASH from the respective authors and executed them in our environment. The collision-resistant hash function was instantiated using SHA-256. We have used multi-threading, and our machines were capable of handling a total of 32 threads. Each experiment is run for 20 times, and the average values are reported.

Datasets We use the following datasets:
- MNIST [36] is a collection of 28×28 pixel, handwritten digit images along with a label between 0 and 9 for each image. It has $60,000$ and respectively, $10,000$ images in the training and test set. We evaluate logistic regression, and NN-1, NN-2 (cf. §5.2) on this dataset.
- CIFAR-10 [34] consists of 32×32 pixel images of 10 different classes such as dogs, horses, etc. There are $50,000$ images for training and $10,000$ for testing, with 6000 images in each class. We evaluate NN-3 (cf. §5.2) on this dataset.

Benchmarking Parameters We use *throughput* (TP) as the benchmarking parameter following BLAZE and ABY3 [39] as it would help to analyse the effect of improved communication and round complexity in a single shot. Here, TP denotes the number of operations ("iterations" for the case of training and "queries" for the case of inference) that can be performed in unit time. We consider minute as the unit time since most of our protocols over WAN requires more than a second to complete. An *iteration* in ML training consists of a *forward propagation* phase followed by a *backward propagation* phase. In the former phase, servers compute the output from the inputs. At the same time, in the latter, the model parameters are adjusted according to the difference in the computed output and the actual output. The inference can be viewed as one forward propagation of the algorithm alone. In addition to TP, we provide the online and overall communication and latency for all the benchmarked ML algorithms.

We observe that due to our protocols' asymmetric nature, the communication load is unevenly distributed among all the servers, which leaves several communication channels underutilized. Thus, to improve the performance, we perform *load balancing*, where we run several parallel execution threads, each with roles of the servers changed. This helps in utilizing all channels and improving the performance.

5.1 Logistic Regression

In Logistic Regression, one iteration comprises updating the weight vector $\vec{\mathbf{w}}$ using the gradient descent algorithm (GD). It is updated according to the function given below: $\vec{\mathbf{w}} = \vec{\mathbf{w}} - \frac{\alpha}{B} \mathbf{X}_i^T \circ (\text{sig}(\mathbf{X}_i \circ \vec{\mathbf{w}}) - \mathbf{Y}_i)$. where α and $\mathbf{X}_i$ denote the learning rate, and a subset, of batch size B, randomly selected from the entire dataset in the ith iteration, respectively. The forward propagation comprises of computing the value $\mathbf{X}_i \circ \vec{\mathbf{w}}$ followed by an application of a sigmoid function on it. The weight vector is updated in the backward propagation,

which internally requires the computation of a series of matrix multiplications, and can be achieved using a dot product. The update function can be computed using $[\![\cdot]\!]$ shares as: $[\![\vec{\mathbf{w}}]\!] = [\![\vec{\mathbf{w}}]\!] - \frac{\alpha}{B}[\![\mathbf{X}_j^T]\!] \circ (\mathsf{sig}([\![\mathbf{X}_j]\!] \circ [\![\vec{\mathbf{w}}]\!]) - [\![\mathbf{Y}_j]\!])$. We summarize our results in Table 3.

Setting	Ref.	Online (TP in $\times 10^3$)			Total	
		Latency (s)	Com [KB]	TP	Latency (s)	Com [KB]
3PC Training	BLAZE	0.74	50.26	4872.38	0.93	203.35
	SWIFT	1.05	50.32	4872.38	1.54	203.47
3PC Inference	BLAZE	0.66	0.28	7852.05	0.84	0.74
	SWIFT	0.97	0.34	6076.46	1.46	0.86
4PC Training	FLASH	0.83	88.93	5194.18	1.11	166.75
	SWIFT	0.83	41.32	11969.48	1.11	92.91
4PC Inference	FLASH	0.76	0.50	7678.40	1.04	0.96
	SWIFT	0.75	0.27	15586.96	1.03	0.57

Table 3: Logistic Regression training and inference. TP is given in (#it/min) for training and (#queries/min) for inference.

We observe that the online TP for the case of 3PC inference is slightly lower compared to that of BLAZE. This is because the total number of rounds for the inference phase is slightly higher in our case due to the additional rounds introduced by the verification mechanism (aka *verify* phase which also needs broadcast). This gap becomes less evident for protocols with more number of rounds, as is demonstrated in the case of NN (presented next), where verification for several iterations is clubbed together, making the overhead for verification insignificant.

For the case of 4PC, our solution outperforms FLASH in terms of communication as well as throughput. Concretely, we observe a $2\times$ improvement in TP for inference and a $2.3\times$ improvement for training. For Trident [14], we observe a drop of 15.86% in TP for inference due to the extra rounds required for verification to achieve GOD. This loss is, however, traded-off with the stronger security guarantee. For training, we are on par with Trident as the effect of extra rounds becomes less significant for more number of rounds, as will also be evident from the comparisons for NN inference.

As a final remark, note that our 4PC sees roughly $2.5\times$ improvement compared to our 3PC for logistic regression.

5.2 NN Inference

We consider the following popular neural networks for benchmarking. These are chosen based on the different range of model parameters and types of layers used in the network. We refer readers to [53] for a detailed architecture of the neural networks.

NN-1: This is a 3-layered fully connected network with ReLU activation after each layer. This network has around 118K parameters and is chosen from [39, 44].

NN-2: This network, called LeNet [35], contains 2 convolutional layers and 2 fully connected layers with ReLU activation after each layer, additionally followed by maxpool for convolutional layers. This network has approximately 431K parameters.

NN-3: This network, called VGG16 [48], was the runner-up of ILSVRC-2014 competition. This network has 16 layers in total and comprises of fully-connected, convolutional, ReLU activation and maxpool layers. This network has about 138 million parameters.

Network	Ref.	Online			Total	
		Latency (s)	Com [MB]	TP	Latency (s)	Com [MB]
NN-1	BLAZE	1.92	0.04	49275.19	2.35	0.11
	SWIFT	2.22	0.04	49275.19	2.97	0.11
NN-2	BLAZE	4.77	3.54	536.52	5.61	9.59
	SWIFT	5.08	3.54	536.52	6.22	9.59
NN-3	BLAZE	15.58	52.58	36.03	18.81	148.02
	SWIFT	15.89	52.58	36.03	19.29	148.02

Table 4: 3PC NN Inference. TP is given in (#queries/min).

Table 4 summarise our benchmarking results for 3PC NN inference. As illustrated, the performance of our 3PC framework is on par with BLAZE while providing better security guarantee.

Network	Ref.	Online			Total	
		Latency (s)	Com [MB]	TP	Latency (s)	Com [MB]
NN-1	FLASH	1.70	0.06	59130.23	2.17	0.12
	SWIFT	1.70	0.03	147825.56	2.17	0.06
NN-2	FLASH	3.93	5.51	653.67	4.71	10.50
	SWIFT	3.93	2.33	1672.55	4.71	5.40
NN-3	FLASH	12.65	82.54	43.61	15.31	157.11
	SWIFT	12.50	35.21	110.47	15.14	81.46

Table 5: 4PC NN Inference. TP is given in (#queries/min).

Table 5 summarises NN inference for 4PC setting. Here, we outperform FLASH in every aspect, with the improvement in TP being at least $2.5\times$ for each NN architecture. Further, we are on par with Trident [14] because the extra rounds required for verification get amortized with an increase in the number of rounds required for computing NN inference. This establishes the practical relevance of our work.

As a final remark, note that our 4PC sees roughly $3\times$ improvement compared to our 3PC for NN inference. This reflects the improvements brought in by the additional honest server in the system.

6 Conclusion

In this work, we presented an efficient framework for PPML that achieves the strongest security of GOD or robustness. Our 3PC protocol builds upon the recent work of BLAZE [44] and achieves almost similar (in some cases, better) performance albeit improving the security guarantee. For the case of 4PC, we outperform the best-known– (a) robust protocol of FLASH [10] by $2\times$ performance-wise and (b) fair protocol of Trident [14] by uplifting its security.

We leave the problem of extending our framework to support mixed-world conversions as well as to design protocols to support algorithms like Decision Trees, k-means Clustering etc. as open problem.

References

[1] M. Abspoel, A. P. K. Dalskov, D. Escudero, and A. Nof. An efficient passive-to-active compiler for honest-majority MPC over rings. Cryptology ePrint Archive, Report 2019/1298, 2019. https://eprint.iacr.org/2019/1298.

[2] T. Araki, A. Barak, J. Furukawa, T. Lichter, Y. Lindell, A. Nof, K. Ohara, A. Watzman, and O. Weinstein. Optimized honest-majority MPC for malicious adversaries - breaking the 1 billion-gate per second barrier. In *IEEE S&P*, pages 843–862, 2017.

[3] T. Araki, J. Furukawa, Y. Lindell, A. Nof, and K. Ohara. High-throughput semi-honest secure three-party computation with an honest majority. In *ACM CCS*, pages 805–817, 2016.

[4] G. Asharov and Y. Lindell. A full proof of the BGW protocol for perfectly secure multiparty computation. *J. Cryptology*, pages 58–151, 2017.

[5] C. Baum, I. Damgård, T. Toft, and R. W. Zakarias. Better preprocessing for secure multiparty computation. In *ACNS*, pages 327–345, 2016.

[6] D. Bogdanov, S. Laur, and J. Willemson. Sharemind: A framework for fast privacy-preserving computations. In *ESORICS*, pages 192–206, 2008.

[7] D. Boneh, E. Boyle, H. Corrigan-Gibbs, N. Gilboa, and Y. Ishai. Zero-knowledge proofs on secret-shared data via fully linear pcps. In *CRYPTO*, pages 67–97, 2019.

[8] E. Boyle, N. Gilboa, Y. Ishai, and A. Nof. Practical fully secure three-party computation via sublinear distributed zero-knowledge proofs. In *ACM CCS*, pages 869–886, 2019.

[9] P. Bunn and R. Ostrovsky. Secure two-party k-means clustering. In *ACM CCS*, pages 486–497, 2007.

[10] M. Byali, H. Chaudhari, A. Patra, and A. Suresh. FLASH: fast and robust framework for privacy-preserving machine learning. *PETS*, 2020. https://eprint.iacr.org/2019/1365.

[11] M. Byali, C. Hazay, A. Patra, and S. Singla. Fast actively secure five-party computation with security beyond abort. In *ACM CCS*, pages 1573–1590, 2019.

[12] M. Byali, A. Joseph, A. Patra, and D. Ravi. Fast secure computation for small population over the internet. In *ACM CCS*, pages 677–694, 2018.

[13] H. Chaudhari, A. Choudhury, A. Patra, and A. Suresh. ASTRA: High Throughput 3PC over Rings with Application to Secure Prediction. In *ACM CCSW@CCS*, 2019.

[14] H. Chaudhari, R. Rachuri, and A. Suresh. Trident: Efficient 4PC Framework for Privacy Preserving Machine Learning. *NDSS*, 2020.

[15] K. Chida, D. Genkin, K. Hamada, D. Ikarashi, R. Kikuchi, Y. Lindell, and A. Nof. Fast large-scale honest-majority MPC for malicious adversaries. In *CRYPTO*, pages 34–64, 2018.

[16] R. Cleve. Limits on the security of coin flips when half the processors are faulty (extended abstract). In *ACM STOC*, pages 364–369, 1986.

[17] R. Cohen, I. Haitner, E. Omri, and L. Rotem. Characterization of secure multiparty computation without broadcast. *J. Cryptology*, pages 587–609, 2018.

[18] R. Cramer, I. Damgård, D. Escudero, P. Scholl, and C. Xing. SpdF_{2^k}: Efficient MPC mod 2^k for dishonest majority. In *CRYPTO*, pages 769–798, 2018.

[19] R. Cramer, S. Fehr, Y. Ishai, and E. Kushilevitz. Efficient multi-party computation over rings. In *EUROCRYPT*, pages 596–613, 2003.

[20] Cryptography and P. E. G. at TU Darmstadt. ENCRYPTO Utils. https://github.com/encryptogroup/ENCRYPTO_utils.

[21] I. Damgård, D. Escudero, T. K. Frederiksen, M. Keller, P. Scholl, and N. Volgushev. New primitives for actively-secure MPC over rings with applications to private machine learning. *IEEE S&P*, 2019.

[22] I. Damgård, M. Keller, E. Larraia, V. Pastro, P. Scholl, and N. P. Smart. Practical covertly secure MPC for dishonest majority - or: Breaking the SPDZ limits. In *ESORICS*, pages 1–18, 2013.

[23] I. Damgård, C. Orlandi, and M. Simkin. Yet another compiler for active security or: Efficient MPC over arbitrary rings. In *CRYPTO*, pages 799–829, 2018.

[24] I. Damgård, V. Pastro, N. P. Smart, and S. Zakarias. Multiparty computation from somewhat homomorphic encryption. In *CRYPTO*, pages 643–662, 2012.

[25] D. Demmler, T. Schneider, and M. Zohner. ABY - A framework for efficient mixed-protocol secure two-party computation. In *NDSS*, 2015.

[26] W. Du and M. J. Atallah. Privacy-preserving cooperative scientific computations. In *IEEE CSFW-14*, pages 273–294, 2001.

[27] H. Eerikson, M. Keller, C. Orlandi, P. Pullonen, J. Puura, and M. Simkin. Use your brain! arithmetic 3pc for any modulus with active security. In *1st Conference on Information-Theoretic Cryptography (ITC*

2020). Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.

[28] J. Furukawa, Y. Lindell, A. Nof, and O. Weinstein. High-throughput secure three-party computation for malicious adversaries and an honest majority. In *EUROCRYPT*, pages 225–255, 2017.

[29] S. D. Gordon, S. Ranellucci, and X. Wang. Secure computation with low communication from cross-checking. In *ASIACRYPT*, pages 59–85, 2018.

[30] G. Jagannathan and R. N. Wright. Privacy-preserving distributed k-means clustering over arbitrarily partitioned data. In *ACM SIGKDD*, pages 593–599, 2005.

[31] M. Keller, E. Orsini, and P. Scholl. MASCOT: faster malicious arithmetic secure computation with oblivious transfer. In *ACM CCS*, pages 830–842, 2016.

[32] M. Keller, V. Pastro, and D. Rotaru. Overdrive: Making SPDZ great again. In *EUROCRYPT*, pages 158–189, 2018.

[33] M. Keller, P. Scholl, and N. P. Smart. An architecture for practical actively secure MPC with dishonest majority. In *ACM CCS*, pages 549–560, 2013.

[34] A. Krizhevsky, V. Nair, and G. Hinton. The CIFAR-10 dataset. 2014. https://www.cs.toronto.edu/~kriz/cifar.html.

[35] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, pages 2278–2324, 1998.

[36] Y. LeCun and C. Cortes. MNIST handwritten digit database. 2010. http://yann.lecun.com/exdb/mnist/.

[37] Y. Lindell and B. Pinkas. Privacy preserving data mining. *J. Cryptology*, pages 177–206, 2002.

[38] E. Makri, D. Rotaru, N. P. Smart, and F. Vercauteren. EPIC: efficient private image classification (or: Learning from the masters). In *CT-RSA*, pages 473–492, 2019.

[39] P. Mohassel and P. Rindal. ABY^3: A mixed protocol framework for machine learning. In *ACM CCS*, pages 35–52, 2018.

[40] P. Mohassel, M. Rosulek, and Y. Zhang. Fast and secure three-party computation: The garbled circuit approach. In *ACM CCS*, pages 591–602, 2015.

[41] P. Mohassel and Y. Zhang. Secureml: A system for scalable privacy-preserving machine learning. In *IEEE S&P*, pages 19–38, 2017.

[42] P. S. Nordholt and M. Veeningen. Minimising communication in honest-majority MPC by batchwise multiplication verification. In *ACNS*, pages 321–339, 2018.

[43] A. Patra and D. Ravi. On the exact round complexity of secure three-party computation. In *CRYPTO*, pages 425–458, 2018.

[44] A. Patra and A. Suresh. BLAZE: Blazing Fast Privacy-Preserving Machine Learning. *NDSS*, 2020. https://eprint.iacr.org/2020/042.

[45] M. C. Pease, R. E. Shostak, and L. Lamport. Reaching agreement in the presence of faults. *J. ACM*, pages 228–234, 1980.

[46] M. S. Riazi, C. Weinert, O. Tkachenko, E. M. Songhori, T. Schneider, and F. Koushanfar. Chameleon: A hybrid secure computation framework for machine learning applications. In *AsiaCCS*, pages 707–721, 2018.

[47] A. P. Sanil, A. F. Karr, X. Lin, and J. P. Reiter. Privacy preserving regression modelling via distributed computation. In *ACM SIGKDD*, pages 677–682, 2004.

[48] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

[49] A. B. Slavkovic, Y. Nardi, and M. M. Tibbits. Secure logistic regression of horizontally and vertically partitioned distributed databases. In *ICDM*, pages 723–728, 2007.

[50] Stanford. CS231n: Convolutional Neural Networks for Visual Recognition.

[51] J. Vaidya, H. Yu, and X. Jiang. Privacy-preserving SVM classification. *Knowl. Inf. Syst.*, pages 161–178, 2008.

[52] S. Wagh, D. Gupta, and N. Chandran. Securenn: 3-party secure computation for neural network training. *PoPETs*, pages 26–49, 2019.

[53] S. Wagh, S. Tople, F. Benhamouda, E. Kushilevitz, P. Mittal, and T. Rabin. Falcon: Honest-majority maliciously secure framework for private deep learning. *arXiv preprint*, 2020. https://arxiv.org/abs/2004.02229v1.

[54] H. Yu, J. Vaidya, and X. Jiang. Privacy-preserving SVM classification on vertically partitioned data. In *PAKDD*, pages 647–656, 2006.

A Preliminaries

A.1 Shared Key Setup

Let $F : \{0,1\}^{\kappa} \times \{0,1\}^{\kappa} \to X$ be a secure pseudo-random function (PRF), with co-domain X being $\mathbb{Z}_{2^{\ell}}$. The set of keys established between the servers for 3PC is as follows:

– One key shared between every pair– k_{01}, k_{02}, k_{12} for the servers $(P_0, P_1), (P_0, P_2)$ and (P_1, P_2), respectively.
– One shared key known to all the servers– $k_{\mathcal{P}}$.

Suppose P_0, P_1 wish to sample a random value $r \in \mathbb{Z}_{2^{\ell}}$ non-interactively. To do so they invoke $F_{k_{01}}(id_{01})$ and obtain r. Here, id_{01} denotes a counter maintained by the servers, and is updated after every PRF invocation. The appropriate keys used to sample is implicit from the context, from the identities of the pair that sample or from the fact that it is sampled by all, and, hence, is omitted.

The key setup is modelled via a functionality $\mathcal{F}_{\text{setup}}$ (Fig. 6) that can be realised using any secure MPC protocol. Analogously, key setup functionality for 4PC is given in Fig. 7.

Functionality $\mathcal{F}_{\text{setup}}$

$\mathcal{F}_{\text{setup}}$ interacts with the servers in $\mathcal{P}$ and the adversary $\mathcal{S}$. $\mathcal{F}_{\text{setup}}$ picks random keys k_{ij} for $i, j \in \{0,1,2\}$ and $k_{\mathcal{P}}$. Let y_s denote the keys corresponding to server P_s. Then

– $y_s = (k_{01}, k_{02}$ and $k_{\mathcal{P}})$ when $P_s = P_0$.
– $y_s = (k_{01}, k_{12}$ and $k_{\mathcal{P}})$ when $P_s = P_1$.
– $y_s = (k_{02}, k_{12}$ and $k_{\mathcal{P}})$ when $P_s = P_2$.

Output: Send (Output, y_s) to every $P_s \in \mathcal{P}$.
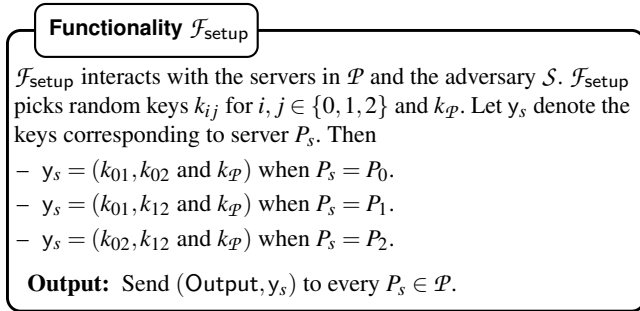
Figure 6: 3PC: Ideal functionality for shared-key setup

Functionality $\mathcal{F}_{\text{setup4}}$

$\mathcal{F}_{\text{setup4}}$ interacts with the servers in $\mathcal{P}$ and the adversary $\mathcal{S}$. $\mathcal{F}_{\text{setup4}}$ picks random keys k_{ij} and k_{ijk} for $i, j, k \in \{0,1,2\}$ and $k_{\mathcal{P}}$. Let y_s denote the keys corresponding to server P_s. Then

– $y_s = (k_{01}, k_{02}, k_{03}, k_{012}, k_{013}, k_{023}$ and $k_{\mathcal{P}})$ when $P_s = P_0$.
– $y_s = (k_{01}, k_{12}, k_{13}, k_{012}, k_{013}, k_{123}$ and $k_{\mathcal{P}})$ when $P_s = P_1$.
– $y_s = (k_{02}, k_{12}, k_{23}, k_{012}, k_{023}, k_{123}$ and $k_{\mathcal{P}})$ when $P_s = P_2$.
– $y_s = (k_{03}, k_{13}, k_{23}, k_{013}, k_{023}, k_{123}$ and $k_{\mathcal{P}})$ when $P_s = P_3$.

Output: Send (Output, y_s) to every $P_s \in \mathcal{P}$.

Figure 7: 4PC: Ideal functionality for shared-key setup

To generate a *3-out-of-3* additive sharing of 0 i.e. ζ_s for $s \in \{0,1,2\}$ such that P_s holds ζ_s, and $\zeta_0 + \zeta_1 + \zeta_2 = 0$, servers proceed as follows. Every pair of servers, $P_s, P_{(s+1)\%3}$, non-interactively generate r_s, as described earlier, and each P_s sets $\zeta_s = r_s - r_{(s-1)\%3}$.

A.2 Collision Resistant Hash Function

Consider a hash function family $\mathsf{H} = \mathcal{K} \times \mathcal{L} \to \mathcal{Y}$. The hash function H is said to be collision resistant if, for all probabilistic polynomial-time adversaries $\mathcal{A}$, given the description of H_k where $k \in_R \mathcal{K}$, there exists a negligible function $\text{negl}()$ such that $\Pr[(x_1, x_2) \leftarrow \mathcal{A}(k) : (x_1 \neq x_2) \wedge \mathsf{H}_k(x_1) = \mathsf{H}_k(x_2)] \leq \text{negl}(\kappa)$, where $m = \text{poly}(\kappa)$ and $x_1, x_2 \in_R \{0,1\}^m$.

A.3 Commitment Scheme

Let $\mathsf{Com}(x)$ denote the commitment of a value x. The commitment scheme $\mathsf{Com}(x)$ possesses two properties; *hiding* and *binding*. The former ensures privacy of the value v given just its commitment $\mathsf{Com}(v)$, while the latter prevents a corrupt server from opening the commitment to a different value $x' \neq x$. The practical realization of a commitment scheme is via a hash function $\mathcal{H}()$ given below, whose security can be proved in the random-oracle model (ROM)– for $(c, o) = (\mathcal{H}(x||r), x||r) = \mathsf{Com}(x; r)$.

B 3PC Protocols

In this section, we provide a detailed communication cost analysis for our protocols in the 3PC setting. Also detailed information regarding some of the protocols are provided.

B.1 Joint Message Passing

The ideal functionality for jmp appears in Fig. 8.

Functionality $\mathcal{F}_{\text{jmp}}$

$\mathcal{F}_{\text{jmp}}$ interacts with the servers in $\mathcal{P}$ and the adversary $\mathcal{S}$.

Step 1: $\mathcal{F}_{\text{jmp}}$ receives (Input, v_s) from P_s for $s \in \{i, j\}$, while it receives $(\text{Select}, \text{ttp})$ from $\mathcal{S}$. Here ttp denotes the server that $\mathcal{S}$ wants to choose as the TTP. Let $P^{\star} \in \mathcal{P}$ denote the server corrupted by $\mathcal{S}$.

Step 2: If $v_i = v_j$ and $\text{ttp} = \perp$, then set $\text{msg}_i = \text{msg}_j = \perp, \text{msg}_k = v_i$ and go to **Step 5**.

Step 3: If $\text{ttp} \in \mathcal{P} \setminus \{P^{\star}\}$, then set $\text{msg}_i = \text{msg}_j = \text{msg}_k = \text{ttp}$.

Step 4: Else, TTP is set to be the honest server with smallest index. Set $\text{msg}_i = \text{msg}_j = \text{msg}_k = \text{TTP}$

Step 5: Send $(\text{Output}, \text{msg}_s)$ to P_s for $s \in \{0,1,2\}$.

Figure 8: 3PC: Ideal functionality for jmp primitive

Lemma B.1 (Communication). *Protocol* Π_{jmp} *(Fig. 1) requires* 1 *round and an amortized communication of* ℓ *bits overall.*

Proof. Server P_i sends value v to P_k while P_j sends hash of the same to P_k. This accounts for one round and communication of ℓ bits. P_k then sends back its inconsistency bit to P_i, P_j, who then exchange it; this takes another two rounds. This is followed by servers broadcasting hashes on their values and selecting a TTP based on it, which takes one more round. All except the first round can be combined for several instances of Π_{jmp} protocol and hence the cost gets amortized. $\square$

B.2 Sharing Protocol

Lemma B.2 (Communication). *Protocol Π_{sh} (Fig. 9) is non-interactive in the preprocessing phase and requires 2 rounds and an amortized communication of 2ℓ bits in the online phase.*

Proof. During the preprocessing phase, servers non-interactively sample the $[\cdot]$-shares of α_v and γ_v values using the shared key setup. In the online phase, when $P_i = P_0$, it computes β_v and sends it to P_1, resulting in one round and ℓ bits communicated. They then jmp-send β_v to P_2, which requires additional one round in an amortized sense, and ℓ bits to be communicated. For the case when $P_i = P_1$, it sends β_v to P_2, resulting in one round and a communication of ℓ bits. Then, P_1, P_2 jmp-send $\beta_v + \gamma_v$ to P_0. This again requires an additional one round and ℓ bits. The analysis is similar in the case of $P_i = P_2$. $\qquad\square$

Protocol $\Pi_{sh}(P_i, v)$

Preprocessing:

- If $P_i = P_0$: P_0, P_j, for $j \in \{1, 2\}$, together sample random $[\alpha_v]_j \in \mathbb{Z}_{2^\ell}$, while $\mathcal{P}$ together sample random $\gamma_v \in \mathbb{Z}_{2^\ell}$.
- If $P_i = P_1$: P_0, P_1 together sample random $[\alpha_v]_1 \in \mathbb{Z}_{2^\ell}$, while $\mathcal{P}$ together sample a random $[\alpha_v]_2 \in \mathbb{Z}_{2^\ell}$. Also, P_1, P_2 together sample random $\gamma_v \in \mathbb{Z}_{2^\ell}$.
- If $P_i = P_2$: Symmetric to the case when $P_i = P_1$.

Online:

- If $P_i = P_0$: P_0 computes $\beta_v = v + \alpha_v$ and sends β_v to P_1. P_1, P_0 jmp-send β_v to P_2.
- If $P_i = P_j$, for $j \in \{1, 2\}$: P_j computes $\beta_v = v + \alpha_v$, sends β_v to P_{3-j}. P_1, P_2 jmp-send $\beta_v + \gamma_v$ to P_0.

Figure 9: 3PC: Generating $[\![v]\!]$-shares by server P_i

B.3 Joint Sharing Protocol

Protocol $\Pi_{jsh}(P_i, P_j, v)$

Preprocessing:

- If $(P_i, P_j) = (P_1, P_0)$: Servers execute the preprocessing of $\Pi_{sh}(P_1, v)$ and then locally set $\gamma_v = 0$.
- If $(P_i, P_j) = (P_2, P_0)$: Similar to the case above.
- If $(P_i, P_j) = (P_1, P_2)$: P_1, P_2 together sample random $\gamma_v \in \mathbb{Z}_{2^\ell}$. Servers locally set $[\alpha_v]_1 = [\alpha_v]_2 = 0$.

Online:

- If $(P_i, P_j) = (P_1, P_0)$: P_0, P_1 compute $\beta_v = v + [\alpha_v]_1 + [\alpha_v]_2$. P_0, P_1 jmp-send β_v to P_2.
- If $(P_i, P_j) = (P_2, P_0)$: Similar to the case above.
- If $(P_i, P_j) = (P_1, P_2)$: P_1, P_2 locally set $\beta_v = v$. P_1, P_2 jmp-send $\beta_v + \gamma_v$ to P_0.

Figure 10: 3PC: $[\![\cdot]\!]$-sharing of a value $v \in \mathbb{Z}_{2^\ell}$ jointly by P_i, P_j

---

When the value v is available to both P_i, P_j in the preprocessing phase, protocol Π_{jsh} can be made non-interactive in the following way: $\mathcal{P}$ sample a random $r \in \mathbb{Z}_{2^\ell}$ and locally set their share according to Table 6.

	(P_1, P_2)		(P_1, P_0)		(P_2, P_0)	
	$[\alpha_v]_1 = 0, [\alpha_v]_2 = 0$		$[\alpha_v]_1 = -v, [\alpha_v]_2 = 0$		$[\alpha_v]_1 = 0, [\alpha_v]_2 = -v$	
	$\beta_v = v, \gamma_v = r - v$		$\beta_v = 0, \gamma_v = r$		$\beta_v = 0, \gamma_v = r$	
P_0	$(0, 0, r\ )$		$(-v, 0, r)$		$(0, -v, r)$	
P_1	$(0, v, r-v)$		$(-v, 0, r)$		$(0, 0, r)$	
P_2	$(0, v, r-v)$		$(0, 0, r)$		$(0, -v, r)$	

Table 6: The columns depict the three distinct possibility of input contributing pairs. The first row shows the assignment to various components of the sharing. The last row, along with three sub-rows, specify the shares held by the three servers.

Lemma B.3 (Communication). *Protocol Π_{jsh} (Fig. 10) is non-interactive in the preprocessing phase and requires 1 round and an amortized communication of ℓ bits in the online phase.*

Proof. In this protocol, servers execute Π_{jmp} protocol once. Hence the overall cost follows from that of an instance of the Π_{jmp} protocol (Lemma B.1). $\qquad\square$

B.4 Multiplication Protocol

The ideal functionality for Π_{mulPre} appears in Fig. 11.

Functionality $\mathcal{F}_{MulPre}$

$\mathcal{F}_{MulPre}$ interacts with the servers in $\mathcal{P}$ and the adversary $\mathcal{S}$. $\mathcal{F}_{MulPre}$ receives $\langle\cdot\rangle$-shares of d, e from the servers where P_s, for $s \in \{0, 1, 2\}$, holds $\langle d \rangle_s = (d_s, d_{(s+1)\%3})$ and $\langle e \rangle_s = (e_s, e_{(s+1)\%3})$ such that $d = d_0 + d_1 + d_2$ and $e = e_0 + e_1 + e_2$. Let P_i denotes the server corrupted by $\mathcal{S}$. $\mathcal{F}_{MulPre}$ receives $\langle f \rangle_i = (f_i, f_{(i+1)\%3})$ from $\mathcal{S}$ where $f = de$. $\mathcal{F}_{MulPre}$ proceeds as follows:

- Reconstructs d, e using the shares received from honest servers and compute $f = de$.
- Compute $f_{(i+2)\%3} = f - f_i - f_{(i+1)\%3}$ and set the output shares as $\langle f \rangle_0 = (f_0, f_1), \langle f \rangle_1 = (f_1, f_2), \langle f \rangle_2 = (f_2, f_0)$.
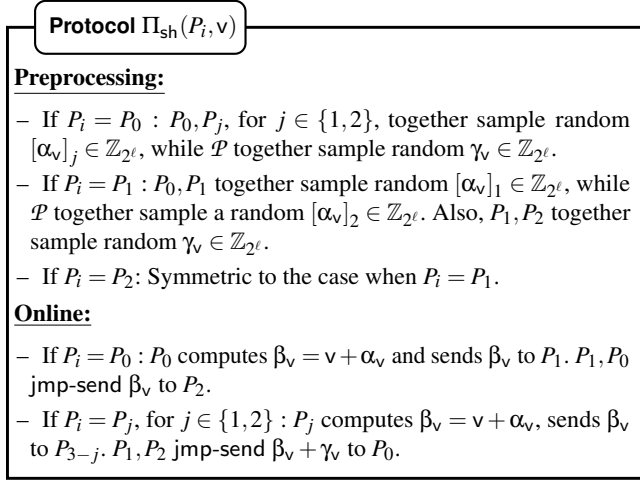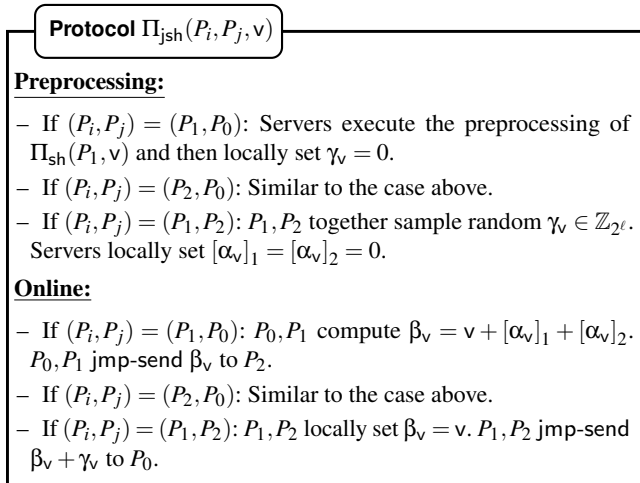- Send $(\mathsf{Output}, \langle f \rangle_s)$ to server $P_s \in \mathcal{P}$.

Figure 11: 3PC: Ideal functionality for Π_{mulPre} protocol

Lemma B.4 (Communication). *Protocol Π_{mult} (Fig. 2) requires an amortized cost of 3ℓ bits in the preprocessing phase, and 1 round and amortized cost of 3ℓ bits in the online phase.*

Proof. In the preprocessing phase, generation of α_z and γ_z are non-interactive. This is followed by one execution of Π_{mulPre}, which requires an amortized communication cost of 3ℓ bits. During the online phase, P_0, P_1 jmp-send $[\beta_z^\star]_1$ to P_2, while P_0, P_2 jmp-send $[\beta_z^\star]_2$ to P_1. This requires one round and a communication of 2ℓ bits. Following this, P_1, P_2 jmp-send $\beta_z + \gamma_z$ to P_0, which requires one round and a communication

of ℓ bits. However, jmp-send of $\beta_z + \gamma_z$ can be delayed till the end of the protocol, and will require only one round for the entire circuit and can be amortized. $\square$

B.5 Reconstruction Protocol

Protocol Π_{rec} appears in Fig. 12.

Protocol $\Pi_{\text{rec}}(\mathcal{P}, [\![v]\!])$

Preprocessing:

– P_0, P_j, for $j \in \{1, 2\}$, compute $\text{Com}([\alpha_v]_j)$, while P_1, P_2 compute $\text{Com}(\gamma_v)$.

– P_1, P_2 jmp-send $\text{Com}(\gamma_v)$ to P_0, while P_0, P_1 jmp-send $\text{Com}([\alpha_v]_1)$ to P_2, and P_0, P_2 jmp-send $\text{Com}([\alpha_v]_2)$ to P_1,
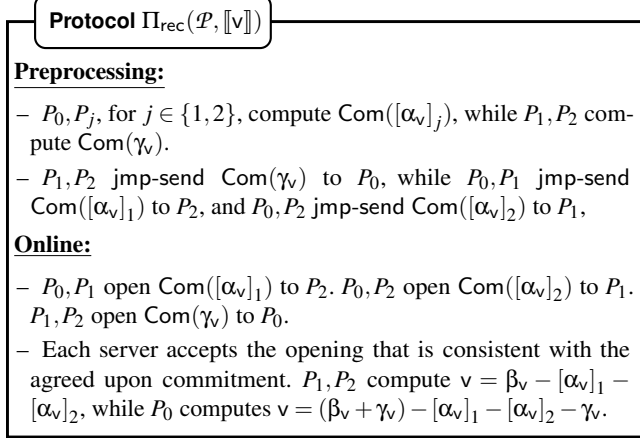
Online:

– P_0, P_1 open $\text{Com}([\alpha_v]_1)$ to P_2. P_0, P_2 open $\text{Com}([\alpha_v]_2)$ to P_1. P_1, P_2 open $\text{Com}(\gamma_v)$ to P_0.

– Each server accepts the opening that is consistent with the agreed upon commitment. P_1, P_2 compute $v = \beta_v - [\alpha_v]_1 - [\alpha_v]_2$, while P_0 computes $v = (\beta_v + \gamma_v) - [\alpha_v]_1 - [\alpha_v]_2 - \gamma_v$.

Figure 12: 3PC: Reconstruction of v among the servers

Lemma B.5 (Communication). *Protocol Π_{rec} (Fig. 12) requires* 1 *round and a communication of* 6ℓ *bits in the online phase.*

Proof. The preprocessing phase consists of communication of commitment values using the Π_{jmp} protocol. The hash-based commitment scheme allows generation of a single commitment for several values and hence the cost gets amortised away for multiple instances. During the online phase, each server receives an opening for the commitment from other two servers, which requires one round and an overall communication of 6ℓ bits. $\square$

B.6 Special protocols

Here we provide details regarding the special protocols - i) Bit Extraction, ii) Bit2A, and iii) Bit Injection.

B.6.1 Bit Extraction protocol

Protocol Π_{bitext} allows servers to compute the *boolean* sharing of the most significant bit (msb) of a value v given its arithmetic sharing $[\![v]\!]$. To compute the msb, we use the optimized 2-input Parallel Prefix Adder (PPA) boolean circuit proposed by ABY3 [39]. The PPA circuit consists of $2\ell - 2$ AND gates and has a multiplicative depth of $\log \ell$.

Let $v_0 = \beta_v, v_1 = -[\alpha_v]_1$ and $v_2 = -[\alpha_v]_2$. Then $v = v_0 + v_1 + v_2$. Servers first locally compute the boolean shares corresponding to each bit of the values v_0, v_1 and v_2 according to Table 7. It has been shown in ABY3 that $v = v_0 + v_1 + v_2$ can also be expressed as $v = 2c + s$ where $\text{FA}(v_0[i], v_1[i], v_2[i]) \rightarrow (c[i], s[i])$ for $i \in \{0, \ldots, \ell - 1\}$. Here FA denotes a Full Adder

	P_0	P_1	P_2
$[\![v_0[i]]\!]^{\mathbf{B}}$	$(0, 0, 0)$	$(0, v_0[i], v_0[i])$	$(0, v_0[i], v_0[i])$
$[\![v_1[i]]\!]^{\mathbf{B}}$	$(v_1[i], 0, 0)$	$(v_1[i], 0, 0)$	$(0, 0, 0)$
$[\![v_2[i]]\!]^{\mathbf{B}}$	$(0, v_2[i], 0)$	$(0, 0, 0)$	$(0, v_2[i], 0)$

Table 7: The $[\![\cdot]\!]^{\mathbf{B}}$-sharing corresponding to i^{th} bit of $v_0 = \beta_v, v_1 = -[\alpha_v]_1$ and $v_2 = -[\alpha_v]_2$. Here $i \in \{0, \ldots, \ell - 1\}$.

circuit while s and c denote the sum and carry bits respectively. To summarize, servers execute ℓ instances of FA in parallel to compute $[\![c]\!]^{\mathbf{B}}$ and $[\![s]\!]^{\mathbf{B}}$. The FA's are executed independently and require one round of communication. The final result is then computed as $\text{msb}(2[\![c]\!]^{\mathbf{B}} + [\![s]\!]^{\mathbf{B}})$ using the optimized PPA circuit.

Lemma B.6 (Communication). *Protocol Π_{bitext} requires a communication cost of $9\ell - 6$ bits in the preprocessing phase and require $\log \ell + 1$ rounds and an amortized communication of $9\ell - 6$ bits in the online phase.*

Proof. In Π_{bitext}, first round comprises of ℓ Full Adder (FA) circuits executing in parallel, each comprising of single AND gate. This is followed by the execution of the optimized PPA circuit of ABY3 [39], which comprises of $2\ell - 2$ AND gates and has a multiplicative depth of $\log \ell$. Hence the communication cost follows from the multiplication for $3\ell - 2$ AND gates. $\square$

B.6.2 Bit2A Conversion protocol

Given the boolean sharing of a bit b, denoted as $[\![b]\!]^{\mathbf{B}}$, protocol Π_{bit2A} (Fig. 13) allows servers to compute the arithmetic sharing $[\![b^{\mathsf{R}}]\!]$. Here b^{R} denotes the equivalent value of b over ring $\mathbb{Z}_{2^\ell}$ (see Notation 2.2). As pointed out in BLAZE, $b^{\mathsf{R}} = (\beta_b \oplus \alpha_b)^{\mathsf{R}} = \beta_b^{\mathsf{R}} + \alpha_b^{\mathsf{R}} - 2\beta_b^{\mathsf{R}}\alpha_b^{\mathsf{R}}$. Also $\alpha_b^{\mathsf{R}} = ([\alpha_b]_1 \oplus [\alpha_b]_2)^{\mathsf{R}} = [\alpha_b]_1^{\mathsf{R}} + [\alpha_b]_2^{\mathsf{R}} - 2[\alpha_b]_1^{\mathsf{R}}[\alpha_b]_2^{\mathsf{R}}$. During the preprocessing phase, P_0, P_j for $j \in \{1, 2\}$ execute Π_{jsh} on $[\alpha_b]_j^{\mathsf{R}}$ to generate $[\![[\alpha_b]_j^{\mathsf{R}}]\!]$. Servers then execute Π_{mult} on $[\![[\alpha_b]_1^{\mathsf{R}}]\!]$ and $[\![[\alpha_b]_2^{\mathsf{R}}]\!]$ to generate $[\![[\alpha_b]_1^{\mathsf{R}}[\alpha_b]_2^{\mathsf{R}}]\!]$ followed by locally computing $[\![\alpha_b^{\mathsf{R}}]\!]$. During the online phase, P_1, P_2 execute Π_{jsh} on β_b^{R} to jointly generate $[\![\beta_b^{\mathsf{R}}]\!]$. Servers then execute Π_{mult} protocol on $[\![\beta_b^{\mathsf{R}}]\!]$ and $[\![\alpha_b^{\mathsf{R}}]\!]$ to compute $[\![\beta_b^{\mathsf{R}}\alpha_b^{\mathsf{R}}]\!]$ followed by locally computing $[\![b^{\mathsf{R}}]\!]$. The formal details for Π_{bit2A} protocol appears in Fig. 13.

Lemma B.7 (Communication). *Protocol Π_{bit2A} (Fig. 13) requires an amortized communication cost of 9ℓ bits in the preprocessing phase and requires* 1 *round and an amortized communication of 4ℓ bits in the online phase.*

Proof. In the preprocessing phase, servers run two instances of Π_{jsh}, which can be done non-interactively (ref. § B.3). This is followed by an execution of entire multiplication protocol, which requires 6ℓ bits to be communicated (Lemma B.4). Parallelly, the servers execute the preprocessing phase of

Π_{mult}, resulting in an additional 3ℓ bits of communication (Lemma B.4). During the online phase, P_1, P_2 execute Π_{jsh} once, which requires one round and ℓ bits to be communicated. In Π_{jsh}, the communication towards P_0 can be deferred till the end, thereby requiring a single round for multiple instances. This is followed by an execution of the online phase of Π_{mult}, which requires one round and a communication of 3ℓ bits. $\qquad\square$

Protocol $\Pi_{\text{bit2A}}(\mathcal{P}, [\![b]\!]^{\mathbf{B}})$

Preprocessing:

– P_0, P_j for $j \in \{1,2\}$ execute Π_{jsh} on $[\alpha_b]_j^R$ to generate $[\![[\alpha_b]_j^R]\!]$.

– Servers execute $\Pi_{\text{mult}}(\mathcal{P}, [\alpha_b]_1^R, [\alpha_b]_2^R)$ to generate $[\![u]\!]$ where $u = [\alpha_b]_1^R [\alpha_b]_2^R$, followed by locally computing $[\![\alpha_b^R]\!] = [\![[\alpha_b]_1^R]\!] + [\![[\alpha_b]_2^R]\!] - 2[\![u]\!]$.

– Servers execute the preprocessing phase of $\Pi_{\text{mult}}(\mathcal{P}, \beta_b^R, \alpha_b^R)$ for $v = \beta_b^R \alpha_b^R$.

Online:

– P_1, P_2 execute $\Pi_{\text{jsh}}(P_1, P_2, \beta_b^R)$ to generate $[\![\beta_b^R]\!]$.

– Servers execute online phase of $\Pi_{\text{mult}}(\mathcal{P}, \beta_b^R, \alpha_b^R)$ to generate $[\![v]\!]$ where $v = \beta_b^R \alpha_b^R$, followed by locally computing $[\![b^R]\!] = [\![\beta_b^R]\!] + [\![\alpha_b^R]\!] - 2[\![v]\!]$.

Figure 13: 3PC: Bit2A Protocol

B.6.3 Bit Injection protocol

Given the binary sharing of a bit b, denoted as $[\![b]\!]^{\mathbf{B}}$, and the arithmetic sharing of $v \in \mathbb{Z}_{2^\ell}$, protocol Π_{BitInj} computes $[\![\cdot]\!]$-sharing of bv. Towards this, servers first execute Π_{bit2A} on $[\![b]\!]^{\mathbf{B}}$ to generate $[\![b]\!]$. This is followed by servers computing $[\![bv]\!]$ by executing Π_{mult} protocol on $[\![b]\!]$ and $[\![v]\!]$.

Lemma B.8 (Communication). *Protocol Π_{BitInj} requires an amortized communication cost of 12ℓ bits in the preprocessing phase and requires 2 rounds and an amortized communication of 7ℓ bits in the online phase.*

Proof. Protocol Π_{BitInj} is essentially an execution of Π_{bit2A} (Lemma B.7) followed by one invocation of Π_{mult} (Lemma B.4) and the costs follow. $\qquad\square$

B.7 Dot Product Protocol

The formal details for Π_{dotp} protocol appears in Fig. 14.

Protocol $\Pi_{\text{dotp}}(\mathcal{P}, \{[\![x_i]\!], [\![y_i]\!]\}_{i \in [n]})$

Preprocessing:

– P_0, P_j, for $j \in \{1,2\}$, together sample random $[\alpha_z]_j \in \mathbb{Z}_{2^\ell}$, while P_1, P_2 sample random $\gamma_z \in \mathbb{Z}_{2^\ell}$.

– Servers locally compute $\langle \cdot \rangle$-sharing of $\vec{d}, \vec{e}$ with $d_i = \gamma_{x_i} + \alpha_{x_i}$

and $e_i = \gamma_{y_i} + \alpha_{y_i}$ for $i \in [n]$ as follows:

$(\langle d_i \rangle_0 = ([\alpha_{x_i}]_2, [\alpha_{x_i}]_1), \langle d_i \rangle_1 = ([\alpha_{x_i}]_1, \gamma_{x_i}), \langle d_i \rangle_2 = (\gamma_{x_i}, [\alpha_{x_i}]_2))$

$(\langle e_i \rangle_0 = ([\alpha_{y_i}]_2, [\alpha_{y_i}]_1), \langle e_i \rangle_1 = ([\alpha_{y_i}]_1, \gamma_{y_i}), \langle e_i \rangle_2 = (\gamma_{y_i}, [\alpha_{y_i}]_2))$

– Servers execute $\Pi_{\text{dotpPre}}(\mathcal{P}, \langle \vec{d} \rangle, \langle \vec{e} \rangle)$ to generate $\langle f \rangle = \langle \vec{d} \odot \vec{e} \rangle$.

– P_0, P_1 locally set $[\chi]_1 = f_1$, while P_0, P_2 locally set $[\chi]_2 = f_0$. P_1, P_2 locally compute $\psi = f_2 - \sum_{i=1}^n \gamma_{x_i} \gamma_{y_i}$.

Online:

– P_0, P_j, for $j \in \{1,2\}$, compute $[\beta_z^\star]_j = -\sum_{i=1}^n ((\beta_{x_i} + \gamma_{x_i})[\alpha_{y_i}]_j + (\beta_{y_i} + \gamma_{y_i})[\alpha_{x_i}]_j) + [\alpha_z]_j + [\chi]_j$.

– P_0, P_1 jmp-send $[\beta_z^\star]_1$ to P_2 and P_0, P_2 jmp-send $[\beta_z^\star]_2$ to P_1.

– P_1, P_2 locally compute $\beta_z^\star = [\beta_z^\star]_1 + [\beta_z^\star]_2$ and set $\beta_z = \beta_z^\star + \sum_{i=1}^n (\beta_{x_i} \beta_{y_i}) + \psi$.

– P_1, P_2 jmp-send $\beta_z + \gamma_z$ to P_0.

Figure 14: 3PC: Dot Product Protocol ($z = \vec{x} \odot \vec{y}$)

The ideal world functionality for realizing Π_{dotpPre} is presented in Fig. 15.

Functionality $\mathcal{F}_{\text{DotPPre}}$

$\mathcal{F}_{\text{DotPPre}}$ interacts with the servers in $\mathcal{P}$ and the adversary $\mathcal{S}$. $\mathcal{F}_{\text{DotPPre}}$ receives $\langle \cdot \rangle$-shares of vectors $\vec{d} = (d_1, \ldots, d_n), \vec{e} = (e_1, \ldots, e_n)$ from the servers. Let $v_{j,s}$ for $j \in [n], s \in \{0,1,2\}$ denote the share of v_j such that $v_j = v_{j,0} + v_{j,1} + v_{j,2}$. Server P_s, for $s \in \{0,1,2\}$, holds $\langle d_j \rangle_s = (d_{j,s}, d_{j,(s+1)\%3})$ and $\langle e_j \rangle_s = (e_{j,s}, e_{j,(s+1)\%3})$ where $j \in [n]$. Let P_i denotes the server corrupted by $\mathcal{S}$. $\mathcal{F}_{\text{MulPre}}$ receives $\langle f \rangle_i = (f_i, f_{(i+1)\%3})$ from $\mathcal{S}$ where $f = \vec{d} \odot \vec{e}$. $\mathcal{F}_{\text{DotPPre}}$ proceeds as follows:

– Reconstructs d_j, e_j, for $j \in [n]$, using the shares received from honest servers and compute $f = \sum_{j=1}^n d_j e_j$.

– Compute $f_{(i+2)\%3} = f - f_i - f_{(i+1)\%3}$ and set the output shares as $\langle f \rangle_0 = (f_0, f_1), \langle f \rangle_1 = (f_1, f_2), \langle f \rangle_2 = (f_2, f_0)$.

– Send $(\mathsf{Output}, \langle f \rangle_s)$ to server $P_s \in \mathcal{P}$.

Figure 15: 3PC: Ideal functionality for Π_{dotpPre} protocol

B.7.1 Instantiating $\mathcal{F}_{\text{DotPPre}}$

To realize $\mathcal{F}_{\text{DotPPre}}$, the approach is to run a semi-honest dot product protocol followed by a verification phase to check the correctness of the output. For verification, the work of [7] provides techniques to verify the correctness of m multiplication triples (and degree-two relations) at a cost of $O(\sqrt{m})$ extended ring elements, albeit with abort security. While [8] improves their techniques to provide *robust* verification for multiplication, we show how to extend the techniques in [8] to robustly verify the correctness of m dot product tuples (dot product being a degree two relation), with vectors of dimension n, at a cost of $O(\sqrt{nm})$ extended ring elements. Thus, the cost to realize one instance of $\mathcal{F}_{\text{DotPPre}}$ can be brought down to only the cost of a semi-honest dot product computation (which is 3 ring elements and independent of the vector dimension), where the cost due to verification can be amortized away by setting n, m appropriately.

Given vectors $\vec{\mathbf{d}} = (\mathsf{d}_1, \dots, \mathsf{d}_n), \vec{\mathbf{e}} = (\mathsf{e}_1, \dots, \mathsf{e}_n)$, let server P_i, for $i \in \{0, 1, 2\}$, hold $\langle \mathsf{d}_j \rangle_i = (\mathsf{d}_{j,i}, \mathsf{d}_{j,(i+1)\%3})$ and $\langle \mathsf{e}_j \rangle_i = (\mathsf{e}_{j,i}, \mathsf{e}_{j,(i+1)\%3})$ where $j \in [\mathsf{n}]$ (henceforth, we omit the use of %3 in the subscript as it is understood from the context). The semi-honest dot product protocol proceeds as follows. The servers, using the shared key setup, non-interactively generate *3-out-of-3* additive shares of zero (as described in A.1), i.e P_i has ζ_i, such that $\zeta_0 + \zeta_1 + \zeta_2 = 0$. Then, each P_i locally computes *3-out-of-3* additive share of $\mathsf{f} = \vec{\mathbf{d}} \odot \vec{\mathbf{e}}$ as:

$$\mathsf{f}_i = \zeta_i + \sum_{j=1}^{n} (\mathsf{d}_{j,i} \cdot \mathsf{e}_{j,i} + \mathsf{d}_{j,i} \cdot \mathsf{e}_{j,i+1} + \mathsf{d}_{j,i+1} \cdot \mathsf{e}_{j,i}) \quad (5)$$

Now, to complete the $\langle \cdot \rangle$-sharing of f, P_i sends f_i to P_{i-1}. To check the correctness of the computation $\langle \mathsf{f} \rangle = \langle \vec{\mathbf{d}} \odot \vec{\mathbf{e}} \rangle$, each $P_i \in \mathcal{P}$ needs to prove that the f_i it sent in the semi-honest protocol satisfies 5, i.e.

$$\zeta_i + \sum_{j=1}^{n} (\mathsf{d}_{j,i} \cdot \mathsf{e}_{j,i} + \mathsf{d}_{j,i} \cdot \mathsf{e}_{j,i+1} + \mathsf{d}_{j,i+1} \cdot \mathsf{e}_{j,i}) - \mathsf{f}_i = 0 \quad (6)$$

This difference in the expected message that should be sent (computed using P_i's correct input shares) and actual message that is sent by P_i is captured by a circuit c, defined below.

$$c \left(\{ \mathsf{d}_{j,i}, \mathsf{d}_{j,i+1}, \mathsf{e}_{j,i}, \mathsf{e}_{j,i+1} \}_{j=1}^{n}, \zeta_i, \mathsf{f}_i \right)$$
$$= \zeta_i + \sum_{j=1}^{n} (\mathsf{d}_{j,i} \cdot \mathsf{e}_{j,i} + \mathsf{d}_{j,i} \cdot \mathsf{e}_{j,i+1} + \mathsf{d}_{j,i+1} \cdot \mathsf{e}_{j,i}) - \mathsf{f}_i \quad (7)$$

Here, c takes as input $u = 4\mathsf{n} + 2$ values: $\langle \cdot \rangle$-shares of $\vec{\mathbf{d}}, \vec{\mathbf{e}}$ held by P_i, i.e. $\{ \mathsf{d}_{j,i}, \mathsf{d}_{j,i+1}, \mathsf{e}_{j,i}, \mathsf{e}_{j,i+1} \}_{j=1}^{n}$, the additive share of zero, ζ_i, that P_i holds, and the additive share f_i sent by P_i. For correct computation with respect to P_i, we require the difference in the expected message and the actual message to be 0, i.e.,

$$c \left(\{ \mathsf{d}_{j,i}, \mathsf{d}_{j,i+1}, \mathsf{e}_{j,i}, \mathsf{e}_{j,i+1} \}_{j=1}^{n}, \zeta_i, \mathsf{f}_i \right) = 0 \quad (8)$$

We now explain how to verify the correctness for m dot product tuples assuming that the operations are carried out over a prime-order field. The verification can be extended to support operations over rings following the techniques of [7, 8]. To verify the correctness for m dot product tuples, $\{ \vec{\mathbf{d}}_k, \vec{\mathbf{e}}_k, \mathsf{f}_k \}_{k=1}^{m}$ where $\mathsf{f}_k = \vec{\mathbf{d}}_k \odot \vec{\mathbf{e}}_k$, the output of c (which is the difference in the expected and actual message sent) for each of the corresponding dot product tuple must be 0. To check correctness of all dot products *at once*, it suffices to check if a random linear combination of the output of each c (for each dot product) is 0. This is because the random linear combination of the differences will be 0 with high probability if $\mathsf{f}_k = \vec{\mathbf{d}}_k \odot \vec{\mathbf{e}}_k$ for each $k \in \{1, \dots, m\}$. We remark that the definition of $c(\cdot)$ in [8] enables the verification of only multiplication triples. With the re-definition of c as in 7, we can now verify the correctness of dot products while the

rest of the verification steps remain similar to that in [8]. We elaborate on the details, next.

A verification circuit, constructed as follows, enables P_i to prove the correctness of the additive share of f that it sent, for m instances of dot product at once. Note that the proof system is designed for the distributed-verifier setting where the proof generated by P_i will be shared among P_{i-1}, P_{i+1}, who can together verify its correctness. First, a sub-circuit g is defined as: group L small c circuits and take a random linear combination of the values on their output wires. Since each c circuit takes $u = 4\mathsf{n} + 2$ inputs as described earlier, g takes in uL inputs. Precisely, g is defined as follows:

$$g(x_1, \dots, x_{uL}) = \sum_{k=1}^{L} \theta_k \cdot c(x_{(k-1)u+1}, \dots, x_{(k-1)u+u})$$

Since there are total m dot products to be verified, there will be $M = m/L$ sub-circuits g. Looking ahead, this grouping technique enables obtaining a sub-linear communication cost for verification because the communication cost turns out to be $O(uL + M)$ and setting $uL = M$ gives the desired result. The sub-circuits g make up the circuit G which outputs a random linear combination of the values on the output wires of each g, i.e:

$$G(x_1, \dots, x_{um}) = \sum_{k=1}^{M} \eta_k \cdot g(x_{(k-1)uL+1}, \dots, x_{(k-1)uL+uL})$$

Here, θ_k and η_k are randomly sampled (non-interactively) by all parties. To prove correctness, P_i needs to prove that G outputs 0. For this, P_i defines $f_1 \dots, f_{uL}$ random polynomials of degree M, one for each input wire of g. For $\ell \in \{1, \dots, M\}$ and $j \in \{1, \dots, uL\}$, $f_j(0)$ is chosen randomly and $f_j(\ell) = x_{(\ell-1)u+j}$ (i.e the jth input of the ℓth g gate). P_i further defines a $2M$ degree polynomial $p(\cdot)$ on the output wires of g, i.e $p(\cdot) = g(f_1, \dots, f_{uL})$ where $p(\ell)$ for $\ell \in \{1, \dots, M\}$ is the output of the ℓth g gate. The additional $M + 1$ points required to interpolate the $2M$ degree polynomial p, are obtained by evaluating $f_1, \dots, f_{uL}$ on $M + 1$ additional points, followed by an application of g circuit. The proof generated by P_i consists of $f_1(0), \dots, f_{uL}(0)$ and the coefficients of p. Recall that since we are in the distributed-verifier setting, the prover P_i additively shares the proof with P_{i-1}, P_{i+1}. Note here, that shares of $f_1(0), \dots, f_{uL}(0)$ can be generated non-interactively.

To verify the proof, verifiers P_{i-1}, P_{i+1} need to check if the output of G is 0. This can be verified by computing the output of G as $b = \sum_{\ell=1}^{M} \eta_\ell \cdot p(\ell)$ and checking if $b = 0$, where η_ℓ's are non-interactively sampled by all after the proof is sent. If p is defined correctly, then this is indeed a random linear combination of the outputs of all the g-circuits. This necessitates the second check to verify the correctness of p as per its definition i.e $p(\cdot) = g(f_1(\cdot), \dots, f_{uL}(\cdot))$. This is performed by checking if $p(r) = g(f_1(r), \dots, f_{uL}(r))$ for a random $r \notin \{1, \dots, M\}$ (for privacy to hold) sampled non-interactively by all after the proof is sent. For the first check,

verifiers can locally compute additive shares of b (using the additive shares of coefficients of p obtained as part of the proof) and reconstruct b to check for equality with 0. For the second, verifiers locally compute additive shares of $p(r)$ using the shares of coefficients of p, and shares of $f_1(r), \ldots, f_{uL}(r)$ by interpolating $f_1, \ldots, f_{uL}$ using (P_i's) inputs to the c-circuits which are implicitly additively shared between them (owing to the replicated sharing property). Verifiers exchange these values among themselves, reconstruct it and check if $p(r) = g(f_1(r), \ldots, f_{uL}(r))$. Note that, the messages computed and exchanged by the verifiers, depend only on the proof sent by P_i and the random values (r, η) sampled by all. These messages can also be independently computed by P_i. Thus, in order to prevent a verifier from falsely rejecting a correct proof, we use jmp to exchange these messages. To optimize the communication cost further, it suffices if a single verifier computes the output of verification.

Setting the parameters: The proof sent by P_i consists of the constant terms $f_j(0)$ for $j \in \{1, \ldots, uL\}$ and $2M+1$ coefficients of p. The former can be can be generated non-interactively. Hence, P_i needs to communicate $2M+1$ elements to the verifiers (one of which can be performed non-interactively). The message sent by the verifier consists of the additive share of $\sum_{\ell=1}^{M} \eta_\ell \cdot p(\ell)$ (for the first check) and $f_1(r), \ldots, f_{uL}(r), p(r)$ (for the second check). Thus, the verifier communicates $uL+2$ elements. As the proof is executed three times, each time with one party acting as the prover and the other two acting as the verifiers, overall, each party communicates $uL+2M+3$ elements. Setting $uL=2M$ and $M = \frac{m}{L}$ results in the total communication required for verifying m dot products to be $O(\sqrt{\mathsf{n}m})$. Thus, verifying a single dot product has an amortized cost of $O\left(\sqrt{\frac{\mathsf{n}}{m}}\right)$ which can be made very small by appropriately setting the values of n, m. Thus, the (amortized) cost of a maliciously secure dot product protocol can be made equal to that of a semi-honest dot product protocol, which is 3 ring elements.

To support verification over rings [8], verification operations are carried out on the extended ring $\mathbb{Z}_{2^\ell}/f(x)$, which is the ring of all polynomials with coefficients in $\mathbb{Z}_{2^\ell}$ modulo a polynomial f, of degree d, irreducible over $\mathbb{Z}_2$. Each element in $\mathbb{Z}_{2^\ell}$ is lifted to a d-degree polynomial in $\mathbb{Z}_{2^\ell}[x]/f(x)$ (which results in blowing up the communication by a factor d). Thus, the per party communication amounts to $(uL+2M+3)d$ elements of $\mathbb{Z}_{2^\ell}$ for verifying m dot products of vector size n where $u = 4\mathsf{n}+2$. Further, the probability of a cheating prover is bounded by $\frac{2^{(\ell-1)d} \cdot 2M+1}{2^{\ell d}-M}$ (cf. Theorem 4.7 of [8]). Thus, if γ is such that $2^\gamma \geq 2M$, then the cheating probability is

$$\frac{2^{(\ell-1)d} \cdot 2M+1}{2^{\ell d}-M} \leq \frac{2^{(\ell-1)d} \cdot 2^\gamma+1}{2^{\ell d}-M} \approx 2^{-(d-\gamma)}$$

We note that both, [8] and our technique require a communication cost of $O(\sqrt{m\mathsf{n}})$ ring elements for verifying m dot products of vector size n. This is because multiplication is a

special case of dot product with $\mathsf{n}=1$. However, since our verification is for dot products, we can get away with performing only m semi-honest dot products whose cost is equivalent to computing m semi-honest multiplications, whereas [8] requires to execute $m\mathsf{n}$ multiplications (as their technique can only verify correctness of multiplications), resulting in a dot product cost dependent on the vector size. Concretely, to get 40 bits of statistical security and for a vector size of 2^{10} (CIFAR-10 [34] dataset), the aforementioned parameters can be set as given in Table 8.

m	M	γ	d	Cost (per dot product)
2^{20}	2^{16}	17	57	7.125
2^{30}	2^{21}	22	62	0.242
2^{40}	2^{26}	27	67	0.008
2^{50}	2^{31}	32	72	0.0002

Table 8: Cost of verification in terms of the number of ring elements communicated per dot product, and parameters for vector size $\mathsf{n} = 2^{10}$ and 40 bits of statistical security. Here, m - #dot products to be verified, M- #g sub-circuits, d-degree of extension.

It is possible to further bring down the communication cost required for verifying m dot product tuples to $O(\log(\mathsf{n}m))$ at the expense of requiring more rounds by further extending the technique of [7], which we leave as an exercise. We refer readers to [8] for formal details.

Lemma B.9 (Communication). *Protocol* Π_{dotp} *(Fig. 14) requires an amortized communication of* 3ℓ *bits in the preprocessing phase and requires* 1 *round and an amortized communication of* 3ℓ *bits in the online phase.*

Proof. During the preprocessing phase, servers execute Π_{dotpPre} as described in §B.7.1. This requires communicating 3ℓ bits for a single semi-honest dot product protocol and $O(\frac{\sqrt{\mathsf{n}}}{\sqrt{m}})$ extended ring elements for its verification. By appropriately setting the values of n, m, the cost of communicating $O(\frac{\sqrt{\mathsf{n}}}{\sqrt{m}})$ elements can be amortized away, thereby resulting in an amortized communication cost of 3ℓ bits in the preprocessing phase. The online phase follows similarly to that of Π_{mult}, the only difference being that servers combine their shares corresponding to all the n multiplications into one and then exchange. This requires one round and an amortized communication of 3ℓ bits. $\square$

B.8 Truncation

We now give details of how to generate $([\mathsf{r}], [\![\mathsf{r}^d]\!])$. For this, servers proceed as follows: P_0, P_j for $j \in \{1, 2\}$ sample random $\mathsf{r}_j \in \mathbb{Z}_{2^\ell}$. Recall that the bit at ith position in r is denoted as $\mathsf{r}[i]$. Define $\mathsf{r}[i] = \mathsf{r}_1[i] \oplus \mathsf{r}_2[i]$ for $i \in \{0, \ldots, \ell-1\}$. For r defined as above, we have $\mathsf{r}^d[i] = \mathsf{r}_1[i+d] \oplus \mathsf{r}_2[i+d]$ for $i \in \{0, \ldots, \ell-d-1\}$. Further,

$$r = \sum_{i=0}^{\ell-1} 2^i r[i] = \sum_{i=0}^{\ell-1} 2^i (r_1[i] \oplus r_2[i])$$

$$= \sum_{i=0}^{\ell-1} 2^i \left((r_1[i])^R + (r_2[i])^R - 2(r_1[i])^R \cdot (r_2[i])^R \right)$$

$$= \sum_{i=0}^{\ell-1} 2^i \left((r_1[i])^R + (r_2[i])^R \right) - \sum_{i=0}^{\ell-1} \left(2^{i+1}(r_1[i])^R \right) \cdot (r_2[i])^R \quad (9)$$

Similarly, for r^d we have the following,

$$r^d = \sum_{i=d}^{\ell-1} 2^{i-d} \left((r_1[i])^R + (r_2[i])^R \right) - \sum_{i=d}^{\ell-1} \left(2^{i-d+1}(r_1[i])^R \right) \cdot (r_2[i])^R \quad (10)$$

The servers non-interactively generate $\llbracket \cdot \rrbracket$-shares (arithmetic shares) for each bit of r_1 and r_2 as in Table 6. Given their $\llbracket \cdot \rrbracket$-shares, the servers execute Π_{dotp} twice to compute $\llbracket \cdot \rrbracket$-share of $A = \sum_{i=d}^{\ell-1}(2^{i-d+1}(r_1[i])^R) \cdot (r_2[i])^R$, and $B = \sum_{i=0}^{\ell-1}(2^{i+1}(r_1[i])^R) \cdot (r_2[i])^R$. Using these values, the servers can locally compute the $\llbracket \cdot \rrbracket$-shares for (r, r^d) pair following Equation 9 and 10, respectively. Note that servers need $[\cdot]$-shares of r and not $\llbracket \cdot \rrbracket$-shares. The $[\cdot]$-shares can be computed from the $\llbracket \cdot \rrbracket$-shares locally as follows. Let $(\alpha_r, \beta_r, \gamma_r)$ be the values corresponding to the $\llbracket \cdot \rrbracket$-shares of r. Since P_0 knows the entire value r in clear, and it knows α_r, it can locally compute β_r. Now, the servers set $[\cdot]$-shares as: $[r]_1 = -[\alpha_r]_1$ and $[r]_2 = \beta_r - [\alpha_r]_2$. The protocol appears in Fig. 16.

Protocol $\Pi_{\mathsf{trgen}}(\mathcal{P})$

– To generate each bit $r[i]$ of r for $i \in \{0, \ldots, \ell-1\}$, P_0, P_j for $j \in \{1, 2\}$ sample random $r_j[i] \in \mathbb{Z}_2$ and define $r[i] = r_1[i] \oplus r_2[i]$.
– Servers generate $\llbracket \cdot \rrbracket$-shares of $(r_j[i])^R$ for $i \in \{0, \ldots, \ell-1\}, j \in \{1, 2\}$ non-interactively following Table 6.
– Define $\vec{\mathbf{x}}$ and $\vec{\mathbf{y}}$ such that $x = 2^{i-d+1}(r_1[i])^R$ and $y_i = (r_2[i])^R$, respectively, for $i \in \{d, \ldots, \ell-1\}$. Define $\vec{\mathbf{p}}$ and $\vec{\mathbf{q}}$ such that $p_i = 2^{i+1}(r_1[i])^R$ and $q_i = (r_2[i])^R$, respectively, for $i \in \{0, \ldots, \ell-1\}$. Servers execute Π_{dotp} to compute $\llbracket \cdot \rrbracket$-shares of $A = \vec{\mathbf{x}} \odot \vec{\mathbf{y}}$ and $B = \vec{\mathbf{p}} \odot \vec{\mathbf{q}}$.
– Servers locally compute $\llbracket r^d \rrbracket = \sum_{i=d}^{\ell-1} 2^{i-d}(\llbracket (r_1[i])^R \rrbracket + \llbracket (r_2[i])^R \rrbracket) - \llbracket A \rrbracket$, and $\llbracket r \rrbracket = \sum_{i=0}^{\ell-1} 2^i (\llbracket (r_1[i])^R \rrbracket + \llbracket (r_2[i])^R \rrbracket) - \llbracket B \rrbracket$.
– P_0 locally computes $\beta_r = r + \alpha_r$. P_0, P_1 set $[r]_1 = -[\alpha_r]_1$ and P_0, P_2 set $[r]_2 = \beta_r - [\alpha_r]_2$.

Figure 16: 3PC: Generating Random Truncated Pair (r, r^d)

Lemma B.10 (Communication). *Protocol Π_{trgen} (Fig. 16) requires an amortized communication of 12ℓ bits.*

Proof. All the operations in Π_{trgen} are non-interactive except for the two dot product calls required to compute A, B. The cost thus follows from Lemma B.9. $\square$

B.9 Dot Product with Truncation

Lemma B.11 (Communication). *Protocol Π_{dotpt} (Fig. 17) requires an amortized communication of 15ℓ bits in the preprocessing phase and requires 1 round and an amortized communication of 3ℓ bits in the online phase.*

Proof. During the preprocessing phase, servers execute the preprocessing phase of Π_{dotp}, resulting in an amortized communication of 3ℓ bits (Lemma B.9). In parallel, servers execute one instance of Π_{trgen} protocol resulting in an additional communication of 12ℓ bits (Lemma B.10).

The online phase follows from that of Π_{dotp} protocol except that, now, P_1, P_2 compute additive shares of $z - r$, where $z = \vec{\mathbf{x}} \odot \vec{\mathbf{y}}$, which is achieved using two executions of Π_{jmp} in parallel. This requires one round and an amortized communication cost of 2ℓ bits. P_1, P_2 then jointly share the truncated value of $z - r$ with P_0, which requires one round and ℓ bits. However, this step can be deferred till the end for multiple dot product with truncation instances, which amortizes the cost. $\square$

The formal details for Π_{dotpt} protocol appear in Fig. 17.

Protocol $\Pi_{\mathsf{dotpt}}(\mathcal{P}, \{\llbracket x_i \rrbracket, \llbracket y_i \rrbracket\}_{i \in [n]})$

Preprocessing:
– Servers execute the preprocessing of $\Pi_{\mathsf{dotp}}(\mathcal{P}, \{\llbracket x_i \rrbracket, \llbracket y_i \rrbracket\}_{i \in [n]})$.
– In parallel, servers execute $\Pi_{\mathsf{trgen}}(\mathcal{P})$ to generate the truncation pair $([r], \llbracket r^d \rrbracket)$.

Online:
– P_0, P_j, for $j \in \{1, 2\}$, compute $[\Psi]_j = -\sum_{i=1}^n((\beta_{x_i} + \gamma_{x_i})[\alpha_{y_i}]_j + (\beta_{y_i} + \gamma_{y_i})[\alpha_{x_i}]_j) - [r]_j$ and set $[(z-r)^\star]_j = [\Psi]_j + [\chi]_j$.
– P_1, P_0 jmp-send $[(z-r)^\star]_1$ to P_2 and P_2, P_0 jmp-send $[(z-r)^\star]_2$ to P_1.
– P_1, P_2 locally compute $(z-r)^\star = [(z-r)^\star]_1 + [(z-r)^\star]_2$ and set $(z-r) = (z-r)^\star + \sum_{i=1}^n (\beta_{x_i}\beta_{y_i}) + \psi$.
– P_1, P_2 locally truncate $(z-r)$ to obtain $(z-r)^d$ and execute $\Pi_{\mathsf{jsh}}(P_1, P_2, (z-r)^d)$ to generate $\llbracket (z-r)^d \rrbracket$.
– Servers locally compute $\llbracket z \rrbracket = \llbracket (z-r)^d \rrbracket + \llbracket r^d \rrbracket$.

Figure 17: 3PC: Dot Product Protocol with Truncation

B.10 Activation Functions

Lemma B.12 (Communication). *Protocol relu requires an amortized communication of $21\ell - 6$ bits in the preprocessing phase and requires $\log \ell + 3$ rounds and an amortized communication of $16\ell - 6$ bits in the online phase.*

Proof. One instance of relu protocol comprises of execution of one instance of Π_{bitext}, followed by Π_{BitInj}. The cost, therefore, follows from Lemma B.6, and Lemma B.8. $\square$

The formal details of the MPC-friendly variant of the Sigmoid function [13, 39, 41] is given below:

$$\mathsf{sig}(v) = \begin{cases} 0 & v < -\frac{1}{2} \\ v + \frac{1}{2} & -\frac{1}{2} \le v \le \frac{1}{2} \\ 1 & v > \frac{1}{2} \end{cases}$$

Lemma B.13 (Communication). *Protocol* sig *requires an amortized communication of* $39\ell - 9$ *bits in the preprocessing phase and requires* $\log \ell + 4$ *rounds and an amortized communication of* $29\ell - 9$ *bits in the online phase.*

Proof. An instance of sig protocol involves the execution of the following protocols in order– i) two parallel instances of Π_{bitext} protocol, ii) once instance of Π_{mult} protocol over boolean value, and iii) one instance of Π_{BitInj} and Π_{bit2A} in parallel. The cost follows from Lemma B.6, Lemma B.7 and Lemma B.8. $\qquad\square$

B.11 Matrix Operations and Convolutions

Matrix Operations Linear matrix operations, such as addition of two matrices $\mathbf{A}, \mathbf{B}$ to generate matrix $\mathbf{C} = \mathbf{A} + \mathbf{B}$, can be computed by extending the scalar operations (addition, in this case) with respect to each element of the matrix. Matrix multiplication, on the other hand, can be expressed as a collection of dot products, where the element in the i^{th} row and j^{th} column of $\mathbf{C} = \mathbf{A} \times \mathbf{B}$, where $\mathbf{A}, \mathbf{B}$ are matrices of dimension $\mathsf{p} \times \mathsf{q}, \mathsf{q} \times \mathsf{r}$, respectively, can be computed as a dot product of the i^{th} row of $\mathbf{A}$ and the j^{th} column of $\mathbf{B}$. Thus, computing $\mathbf{C}$ of dimension $\mathsf{p} \times \mathsf{r}$ requires pr dot products whose communication cost (amortized) is equal to that of computing pr multiplications in our case. This improves the cost of matrix multiplication over the naive approach which requires pqr multiplications.

Convolutions These form an important building block in several neural network architectures and can be represented as matrix multiplications, as explained in the example below. Consider a 2-dimensional convolution (CV) of a 3×3 input matrix $\mathbf{X}$ with a kernel $\mathbf{K}$ of size 2×2. This can be represented as a matrix multiplication as follows.

$$\mathsf{CV}\left(\begin{bmatrix} \mathsf{x}_1 & \mathsf{x}_2 & \mathsf{x}_3 \\ \mathsf{x}_4 & \mathsf{x}_5 & \mathsf{x}_6 \\ \mathsf{x}_7 & \mathsf{x}_8 & \mathsf{x}_9 \end{bmatrix}, \begin{bmatrix} \mathsf{k}_1 & \mathsf{k}_2 \\ \mathsf{k}_3 & \mathsf{k}_4 \end{bmatrix}\right) = \begin{bmatrix} \mathsf{x}_1 & \mathsf{x}_2 & \mathsf{x}_4 & \mathsf{x}_5 \\ \mathsf{x}_2 & \mathsf{x}_3 & \mathsf{x}_5 & \mathsf{x}_6 \\ \mathsf{x}_4 & \mathsf{x}_5 & \mathsf{x}_7 & \mathsf{x}_8 \\ \mathsf{x}_5 & \mathsf{x}_6 & \mathsf{x}_8 & \mathsf{x}_9 \end{bmatrix} \begin{bmatrix} \mathsf{k}_1 \\ \mathsf{k}_2 \\ \mathsf{k}_3 \\ \mathsf{k}_4 \end{bmatrix}$$

Generally, convolving a $f \times f$ kernel over a $w \times h$ input with $p \times p$ padding using $s \times s$ stride having i input channels and o output channels, is equivalent to performing a matrix multiplication on matrices of dimension $(w' \cdot h') \times (i \cdot f \cdot f)$ and $(i \cdot f \cdot f) \times (o)$ where $w' = \dfrac{w - f + 2p}{s} + 1$ and $h' = \dfrac{h - f + 2p}{s} + 1$. We refer readers to [52] (cf. "Linear and Convolutional Layer") and [50] for more details.

C 4PC Protocols

In this section, we give the formal details for the 4PC protocols along with communication cost analysis.

C.1 4PC Joint Message Passing Primitive

Lemma C.1 (Communication). *Protocol* Π_{jmp4} *(Fig. 4) requires* 1 *round and an amortized communication of* ℓ *bits in the online phase.*

Proof. Server P_i sends the value v to P_k while P_j sends hash of the same to P_k. This accounts for one round of communication. Values sent by P_j for several instances can be concatenated and hashed to obtain a single value. Hence the cost of sending the hash gets amortized over multiple instances. Similarly, the two round exchange of inconsistency bits to establish a TTP can be combined for multiple instances, thereby amortizing this cost. Thus, the amortized cost of this protocol is ℓ bits. $\qquad\square$

The ideal functionality for jmp4 primitive appears in Fig. 18.

Functionality $\mathcal{F}_{\text{jmp4}}$

$\mathcal{F}_{\text{jmp4}}$ interacts with the servers in $\mathcal{P}$ and the adversary $\mathcal{S}$.

Step 1: $\mathcal{F}_{\text{jmp}}$ receives (Input, v_s) from senders P_s for $s \in \{i, j\}$, (Input, $\perp$) from receiver P_k and fourth server P_l, while it receives (Select, ttp) from $\mathcal{S}$. Here ttp is a boolean value, with a 1 indicating that $\mathsf{TTP} = P_l$ should be established.

Step 2: If $\mathsf{v}_i = \mathsf{v}_j$ and ttp $= 0$, or if $\mathcal{S}$ has corrupted P_l, set $\mathsf{msg}_i = \mathsf{msg}_j = \mathsf{msg}_l = \perp, \mathsf{msg}_k = \mathsf{v}_i$ and go to **Step 4**.

Step 3: Else : Set $\mathsf{msg}_i = \mathsf{msg}_j = \mathsf{msg}_k = \mathsf{msg}_l = P_l$.

Step 4: Send (Output, msg_s) to P_s for $s \in \{0, 1, 2, 3\}$.

Figure 18: 4PC: Ideal functionality for jmp4 primitive

C.2 Sharing Protocol

The protocol for sharing a value v by P_i is given in Fig. 19

Protocol $\Pi_{\text{sh4}}(P_i, \mathsf{v})$

Preprocessing:

– If $P_i = P_0 : P_0, P_3, P_j$, for $j \in \{1, 2\}$, together sample random $[\alpha_\mathsf{v}]_j \in \mathbb{Z}_{2^\ell}$, while $\mathcal{P}$ sample random $\gamma_\mathsf{v} \in \mathbb{Z}_{2^\ell}$.

– If $P_i = P_1 : P_0, P_3, P_1$ together sample random $[\alpha_\mathsf{v}]_1 \in \mathbb{Z}_{2^\ell}$, while $\mathcal{P}$ sample a random $[\alpha_\mathsf{v}]_2 \in \mathbb{Z}_{2^\ell}$. Also, P_1, P_2, P_3 sample random $\gamma_\mathsf{v} \in \mathbb{Z}_{2^\ell}$.

– If $P_i = P_2$: Analogous to the case when $P_i = P_1$.

– If $P_i = P_3 : P_0, P_3, P_j$, for $j \in \{1, 2\}$, sample random $[\alpha_\mathsf{v}]_j \in \mathbb{Z}_{2^\ell}$. P_1, P_2, P_3 together sample random $\gamma_\mathsf{v} \in \mathbb{Z}_{2^\ell}$.

Online:

– If $P_i = P_0 : P_0$ computes $\beta_\mathsf{v} = \mathsf{v} + \alpha_\mathsf{v}$ and sends β_v to P_1. P_0, P_1 jmp4-send β_v to P_2.

– If $P_i = P_j$, for $j \in \{1, 2\} : P_j$ computes $\beta_\mathsf{v} = \mathsf{v} + \alpha_\mathsf{v}$, sends β_v to P_{3-j}. P_1, P_2 jmp4-send $\beta_\mathsf{v} + \gamma_\mathsf{v}$ to P_0.

– If $P_i = P_3$: P_3 sends $\beta_\mathsf{v} + \gamma_\mathsf{v} = \mathsf{v} + \alpha_\mathsf{v} + \gamma_\mathsf{v}$ to P_0. P_3, P_0 jmp4-send $\beta_\mathsf{v} + \gamma_\mathsf{v}$ to both P_1 and P_2.

Figure 19: 4PC: Generating $[\![\mathsf{v}]\!]$-shares by server P_i

Lemma C.2 (Communication). *In the online phase, Π_{sh4} (Fig. 19) requires 2 rounds and an amortized communication of 2ℓ bits when P_0, P_1, P_2 share a value, whereas it requires an amortized communication of 3ℓ bits when P_3 shares a value.*

Proof. The proof for P_0, P_1, P_2 sharing a value follows from B.2. For the case when P_3 wants to share a value v, it first sends $\beta_\mathsf{v} + \gamma_\mathsf{v}$ to P_0 which requires one round and ℓ bits of communication. This is followed by 2 parallel calls to Π_{jmp4} which together require one round and an amortized communication of 2ℓ bits. $\qquad\square$

C.3 Joint Sharing Protocol

Protocol Π_{jsh4} enables a pair of (unordered) servers (P_i, P_j) to jointly generate a $[\![\cdot]\!]$-sharing of value $\mathsf{v} \in \mathbb{Z}_{2^\ell}$ known to both of them. In case of an inconsistency, the server outside the computation serves as a TTP. The protocol is described in Fig. 20.
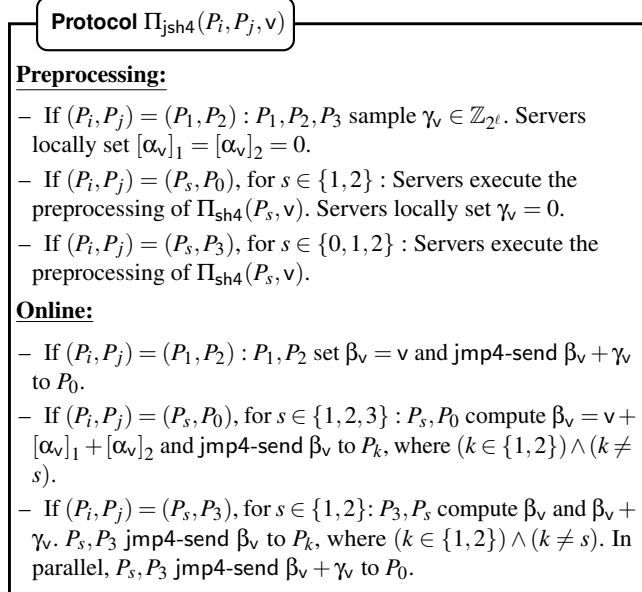

Protocol $\Pi_{\mathsf{jsh4}}(P_i, P_j, \mathsf{v})$

Preprocessing:

– If $(P_i, P_j) = (P_1, P_2)$: P_1, P_2, P_3 sample $\gamma_\mathsf{v} \in \mathbb{Z}_{2^\ell}$. Servers locally set $[\alpha_\mathsf{v}]_1 = [\alpha_\mathsf{v}]_2 = 0$.

– If $(P_i, P_j) = (P_s, P_0)$, for $s \in \{1, 2\}$: Servers execute the preprocessing of $\Pi_{\mathsf{sh4}}(P_s, \mathsf{v})$. Servers locally set $\gamma_\mathsf{v} = 0$.

– If $(P_i, P_j) = (P_s, P_3)$, for $s \in \{0, 1, 2\}$: Servers execute the preprocessing of $\Pi_{\mathsf{sh4}}(P_s, \mathsf{v})$.

Online:

– If $(P_i, P_j) = (P_1, P_2)$: P_1, P_2 set $\beta_\mathsf{v} = \mathsf{v}$ and jmp4-send $\beta_\mathsf{v} + \gamma_\mathsf{v}$ to P_0.

– If $(P_i, P_j) = (P_s, P_0)$, for $s \in \{1, 2, 3\}$: P_s, P_0 compute $\beta_\mathsf{v} = \mathsf{v} + [\alpha_\mathsf{v}]_1 + [\alpha_\mathsf{v}]_2$ and jmp4-send β_v to P_k, where $(k \in \{1, 2\}) \wedge (k \neq s)$.

– If $(P_i, P_j) = (P_s, P_3)$, for $s \in \{1, 2\}$: P_3, P_s compute β_v and $\beta_\mathsf{v} + \gamma_\mathsf{v}$. P_s, P_3 jmp4-send β_v to P_k, where $(k \in \{1, 2\}) \wedge (k \neq s)$. In parallel, P_s, P_3 jmp4-send $\beta_\mathsf{v} + \gamma_\mathsf{v}$ to P_0.

Figure 20: 4PC: $[\![\cdot]\!]$-sharing of a value $\mathsf{v} \in \mathbb{Z}_{2^\ell}$ jointly by P_i, P_j

When P_3, P_0 want to jointly share a value v which is available in the preprocessing phase, protocol Π_{jsh4} can be performed with a single element of communication (as opposed to 2 elements in Fig. 20). P_0, P_3 can jointly share v as follows. P_0, P_3, P_1 sample a random $\mathsf{r} \in \mathbb{Z}_{2^\ell}$ and set $[\alpha_\mathsf{v}]_1 = \mathsf{r}$. P_0, P_3 set $[\alpha_\mathsf{v}]_2 = -(\mathsf{r} + \mathsf{v})$ and jmp4-send $[\alpha_\mathsf{v}]_2$ to P_2. This is followed by servers locally setting $\gamma_\mathsf{v} = \beta_\mathsf{v} = 0$.

We further observe that servers can generate a $[\![\cdot]\!]$-sharing of v non-interactively when v is available with P_0, P_1, P_2. To do this, servers set $[\alpha_\mathsf{v}]_1 = [\alpha_\mathsf{v}]_2 = \gamma_\mathsf{v} = 0$ and $\beta_\mathsf{v} = \mathsf{v}$. We abuse notation and use $\Pi_{\mathsf{jsh4}}(P_0, P_1, P_2, \mathsf{v})$ to denote this sharing.

Lemma C.3 (Communication). *In the online phase, Π_{jsh4} (Fig. 20) requires 1 round and an amortized communication*

of 2ℓ bits when (P_3, P_s) for $s \in \{0, 1, 2\}$ share a value, and requires an amortized communication of ℓ bits, otherwise.

Proof. When (P_3, P_s) for $s \in \{0, 1, 2\}$ want to share a value v, there are two parallel calls to Π_{jmp4} which requires an amortized communication of 2ℓ bits and one round. In the other cases, Π_{jmp4} is invoked only once, resulting in an amortized communication of ℓ bits. $\qquad\square$

C.4 $\langle \cdot \rangle$-sharing Protocol

In some protocols, P_3 is required to generate $\langle \cdot \rangle$-sharing of a value v in the preprocessing phase, where $\langle \cdot \rangle$-sharing of v is same as that defined in 3PC (where $\mathsf{v} = \mathsf{v}_0 + \mathsf{v}_1 + \mathsf{v}_2$, and P_0 possesses $(\mathsf{v}_0, \mathsf{v}_1)$, P_1 possesses $(\mathsf{v}_1, \mathsf{v}_2)$, and P_2 possess $(\mathsf{v}_2, \mathsf{v}_0)$) with the addition that P_3 now possesses $(\mathsf{v}_0, \mathsf{v}_1, \mathsf{v}_2)$. We call the resultant protocol Π_{ash4} and it appears in Fig. 21.

Protocol $\Pi_{\mathsf{ash4}}(P_3, \mathsf{v})$

Preprocessing :

– Servers P_0, P_3, P_1 sample a random $\mathsf{v}_1 \in \mathbb{Z}_{2^\ell}$, while servers P_0, P_3, P_2 sample a random $\mathsf{v}_0 \in \mathbb{Z}_{2^\ell}$.

– P_3 computes $\mathsf{v}_2 = \mathsf{v} - \mathsf{v}_0 - \mathsf{v}_1$ and sends v_2 to P_2. P_3, P_2 jmp4-send v_2 to P_1.
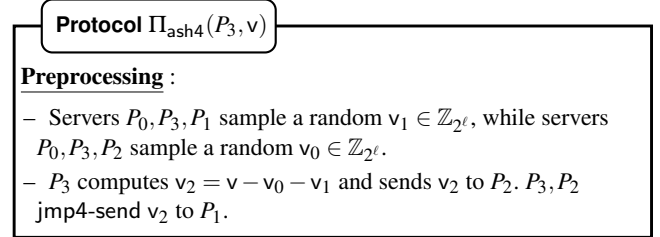
Figure 21: 4PC: $\langle \cdot \rangle$-sharing of value v by P_3

Note that servers can locally convert $\langle \mathsf{v} \rangle$ to $[\![\mathsf{v}]\!]$ by setting their shares as shown in Table 9.

	P_0	P_1	P_2	P_3
$[\![\mathsf{v}]\!]$	$(-\mathsf{v}_1, -\mathsf{v}_0, 0)$	$(-\mathsf{v}_1, \mathsf{v}_2, -\mathsf{v}_2)$	$(-\mathsf{v}_0, \mathsf{v}_2, -\mathsf{v}_2)$	$(-\mathsf{v}_0, -\mathsf{v}_1, -\mathsf{v}_2)$

Table 9: Local conversion of shares from $\langle \cdot \rangle$-sharing to $[\![\cdot]\!]$-sharing for a value v. Here, $[\alpha_\mathsf{v}]_1 = -\mathsf{v}_1, [\alpha_\mathsf{v}]_2 = -\mathsf{v}_0, \beta_\mathsf{v} = \mathsf{v}_2, \gamma_\mathsf{v} = -\mathsf{v}_2$.

Lemma C.4 (Communication). *Protocol Π_{ash4} (Fig. 21) requires 2 rounds and an amortized communication of 2ℓ bits.*

Proof. Communicating v_2 to P_2 requires ℓ bits and 1 round. This is followed by one invocation of Π_{jmp4} which requires ℓ bits and 1 round. Thus, the amortized communication cost is 2ℓ bits and two rounds. $\qquad\square$

C.5 Multiplication Protocol

Lemma C.5 (Communication). *Π_{mult4} (Fig. 5) requires an amortized communication of 3ℓ bits in the preprocessing phase, and 1 round with an amortized communication of 3ℓ bits in the online phase.*

Proof. In the preprocessing phase, the servers execute Π_{jmp4} to jmp4-send $[\Gamma_{\mathsf{xy}}]_2$ to P_2 resulting in amortized communication of ℓ bits. This is followed by 2 parallel invocations of Π_{jmp4} to jmp4-send $[\chi]_1, [\chi]_2$ to P_0 which require an amortized communication of 2ℓ bits. Thus, the amortized communication cost in preprocessing is 3ℓ bits. In the online

phase, there are 2 parallel invocations of Π_{jmp4} to jmp4-send $[\beta_z^\star]_1, [\beta_z^\star]_2$ to P_2, P_1, respectively, which requires amortized communication of 2ℓ bits and one round. This is followed by another call to Π_{jmp4} to jmp4-send $\beta_z + \gamma_z$ to P_0 which requires one more round and amortized communication of ℓ bits. However, jmp4-send of $\beta_z + \gamma_z$ can be delayed till the end of the protocol, and will require only one round for multiple multiplication gates and hence, can be amortized. Thus, the total number of rounds required for multiplication in the online phase is one with an amortized communication of 3ℓ bits. $\qquad\square$

C.6 Reconstruction Protocol

The formal protocol for reconstruction is given in Fig. 22.

Protocol $\Pi_{\mathsf{rec4}}(\mathcal{P}, [\![\mathsf{v}]\!])$

Online

- P_0 receives γ_v from P_1, P_2 and $\mathsf{H}(\gamma_\mathsf{v})$ from P_3.
- P_1 receives $[\alpha_\mathsf{v}]_2$ from P_2, P_3 and $\mathsf{H}([\alpha_\mathsf{v}]_2)$ from P_0.
- P_2 receives $[\alpha_\mathsf{v}]_1$ from P_0, P_3 and $\mathsf{H}([\alpha_\mathsf{v}]_1)$ from P_1.
- P_3 receives $\beta_\mathsf{v} + \gamma_\mathsf{v}$ from P_0, P_1 and $\mathsf{H}(\beta_\mathsf{v} + \gamma_\mathsf{v})$ from P_2.
- $P_i \in \mathcal{P}$ selects the missing share forming the majority among the values received and reconstructs the output.
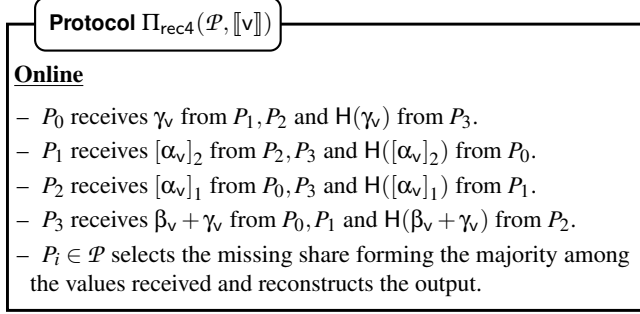---

Figure 22: 4PC: Reconstruction of v among the servers

Lemma C.6 (Communication). Π_{rec4} *(Fig. 22) requires an amortized communication of 8ℓ bits and 1 round in the online phase.*

Proof. Each P_s for $s \in \{0, 1, 2, 3\}$ receives the missing share in clear from two other servers, while the hash of it from the third. As before, the missing share sent by the third server can be concatenated over multiple instances and hashed to obtain a single value. Thus, the amortized communication cost is 2ℓ bits per server, resulting in a total cost of 8ℓ bits. $\qquad\square$

C.7 Special protocols

Here we provide details regarding the special protocols - i) Bit Extraction, ii) Bit2A, and iii) Bit Injection.

C.7.1 Bit Extraction Protocol

This protocol enables the servers to compute a boolean sharing of the most significant bit (MSB) of a value $\mathsf{v} \in \mathbb{Z}_{2^\ell}$ given the arithmetic sharing $[\![\mathsf{v}]\!]$. To compute the MSB, we use the optimized Parallel Prefix Adder (PPA) circuit from ABY3 [39], which takes as input two boolean values and outputs the MSB of the sum of the inputs. The circuit requires $2(\ell-1)$ AND gates and has a multiplicative depth of $\log \ell$. The protocol for bit extraction (Π_{bitext4}) involves computing the boolean PPA circuit using the protocols described in §4. The two inputs to this boolean circuit are generated as follows.

The value v whose MSB needs to be extracted can be represented as the sum of two values as $\mathsf{v} = \beta_\mathsf{v} + (-\alpha_\mathsf{v})$ where the first input to the circuit will be β_v and the second input will be $-\alpha_\mathsf{v}$. Since β_v is held by P_1, P_2, servers execute Π_{jsh4} to generate $[\![\beta_\mathsf{v}]\!]^{\mathbf{B}}$. Similarly, P_0, P_3 possess α_v, and servers execute Π_{jsh4} to generate $[\![-\alpha_\mathsf{v}]\!]^{\mathbf{B}}$. Servers in $\mathcal{P}$ use the $[\![\cdot]\!]^{\mathbf{B}}$-shares of these two inputs $(\beta_\mathsf{v}, -\alpha_\mathsf{v})$ to compute the optimized PPA circuit which outputs the $[\![\mathsf{msb}(\mathsf{v})]\!]^{\mathbf{B}}$.

Lemma C.7 (Communication). *The protocol Π_{bitext4} requires an amortized communication of $7\ell - 6$ bits in the preprocessing phase, and $\log \ell$ rounds with amortized communication of $7\ell - 6$ bits in the online phase.*

Proof. Generation of boolean sharing of α_v requires ℓ bits in the preprocessing phase (since Π_{jsh4} with P_0, P_3 can be achieved with ℓ bits of communication in the preprocessing phase), and generation of boolean sharing of β_v requires ℓ bits and one round (which can be deferred towards the end of the protocol thereby requiring one round for several instances) in the online phase. Further, the boolean PPA circuit to be computed requires $2(\ell-1)$ AND gates. Since each AND gate requires Π_{mult4} to be executed, it requires an amortized communication of $6\ell - 6$ bits in both the preprocessing phase and the online phase. Thus, the overall communication is $7\ell - 6$ bits, in both, the preprocessing and online phase. The circuit has a multiplicative depth of $\log \ell$ which results in $\log \ell$ rounds in the online phase. $\qquad\square$

C.7.2 Bit2A Protocol

This protocol enables servers to compute the arithmetic sharing of a bit b given its boolean sharing. Let b^{R} denote the value of b in the ring $\mathbb{Z}_{2^\ell}$. We observe that b^{R} can be written as follows. $\mathsf{b}^{\mathsf{R}} = (\alpha_\mathsf{b} \oplus \beta_\mathsf{b})^{\mathsf{R}} = \alpha_\mathsf{b}^{\mathsf{R}} + \beta_\mathsf{b}^{\mathsf{R}} - 2\alpha_\mathsf{b}^{\mathsf{R}}\beta_\mathsf{b}^{\mathsf{R}}$. Thus, to obtain an arithmetic sharing of b^{R}, the servers can compute an arithmetic sharing of $\beta_\mathsf{b}^{\mathsf{R}}$, $\alpha_\mathsf{b}^{\mathsf{R}}$ and $\beta_\mathsf{b}^{\mathsf{R}}\alpha_\mathsf{b}^{\mathsf{R}}$. This can be done as follows. P_0, P_3 execute Π_{jsh4} on $\alpha_\mathsf{b}^{\mathsf{R}}$ in the preprocessing phase to generate $[\![\alpha_\mathsf{b}^{\mathsf{R}}]\!]$. Similarly, P_1, P_2 execute Π_{jsh4} on $\beta_\mathsf{b}^{\mathsf{R}}$ in the online phase to generate $[\![\beta_\mathsf{b}^{\mathsf{R}}]\!]$. This is followed by Π_{mult4} on $[\![\beta_\mathsf{b}^{\mathsf{R}}]\!], [\![\alpha_\mathsf{b}^{\mathsf{R}}]\!]$, followed by local computation to obtain $[\![\mathsf{b}^{\mathsf{R}}]\!]$.

While the above approach serves the purpose, we now provide an improved version, which further helps in reducing the online cost. We observe that b^{R} can be written as follows. $\mathsf{b}^{\mathsf{R}} = (\alpha_\mathsf{b} \oplus \beta_\mathsf{b})^{\mathsf{R}} = ((\alpha_\mathsf{b} \oplus \gamma_\mathsf{b}) \oplus (\beta_\mathsf{b} \oplus \gamma_\mathsf{b}))^{\mathsf{R}} = (\mathsf{e} \oplus \mathsf{c})^{\mathsf{R}} = \mathsf{e}^{\mathsf{R}} + \mathsf{c}^{\mathsf{R}} - 2\mathsf{e}^{\mathsf{R}}\mathsf{c}^{\mathsf{R}}$ where $\mathsf{e} = \alpha_\mathsf{b} \oplus \gamma_\mathsf{b}$ and $\mathsf{c} = \beta_\mathsf{b} \oplus \gamma_\mathsf{b}$. Thus, to obtain an arithmetic sharing of b^{R}, P_3 generates $\langle\cdot\rangle$-sharing of e^{R}. To ensure the correctness of the shares, the servers P_0, P_1, P_2 check whether the following equation holds: $(\mathsf{e} \oplus \mathsf{r}_\mathsf{b})^{\mathsf{R}} = \mathsf{e}^{\mathsf{R}} + \mathsf{r}_\mathsf{b}^{\mathsf{R}} - 2\mathsf{e}^{\mathsf{R}}\mathsf{r}_\mathsf{b}^{\mathsf{R}}$. If the verification fails, a TTP is identified. Else, this is followed by servers locally converting $\langle\mathsf{e}^{\mathsf{R}}\rangle$-shares to $[\![\mathsf{e}^{\mathsf{R}}]\!]$ according to Table 9, followed by multiplying $[\![\mathsf{e}^{\mathsf{R}}]\!], [\![\mathsf{c}^{\mathsf{R}}]\!]$ and locally computing $[\![\mathsf{b}^{\mathsf{R}}]\!] = [\![\mathsf{e}^{\mathsf{R}}]\!] + [\![\mathsf{c}^{\mathsf{R}}]\!] - 2[\![\mathsf{e}^{\mathsf{R}}\mathsf{c}^{\mathsf{R}}]\!]$. Note that during $\Pi_{\mathsf{jsh4}}(P_0, P_1, P_2, \mathsf{c}^{\mathsf{R}})$ since $\alpha_{\mathsf{c}^{\mathsf{R}}}$ and $\gamma_{\mathsf{c}^{\mathsf{R}}}$ are set to 0, the preprocessing of multiplication can be performed locally.

Protocol $\Pi_{\text{bit2A4}}(\mathcal{P}, [\![b]\!]^{\mathbf{B}})$

Preprocessing :

– Servers execute $\Pi_{\text{ash4}}(P_3, e^R)$ (Fig. 21) where $e = \alpha_b \oplus \gamma_b$. Let the shares be $\langle e^R \rangle_0 = (e_0, e_1)$, $\langle e^R \rangle_1 = (e_1, e_2)$, $\langle e^R \rangle_2 = (e_2, e_0)$, $\langle e^R \rangle_3 = (e_0, e_1, e_2)$.

– Verification of $\langle e^R \rangle$-sharing is performed as follows:

  – P_1, P_2, P_3 sample a random $r \in \mathbb{Z}_{2^\ell}$ and a bit $r_b \in \mathbb{Z}_{2^1}$.
  – P_1, P_2 compute $x_1 = \gamma_b \oplus r_b$, and jmp4-send x_1 to P_0.
  – P_1, P_3 compute $y_1 = (e_1 + e_2)(1 - 2r_b^R) + r_b^R + r$, and jmp4-send y_1 to P_0.
  – P_2, P_3 compute $y_2 = e_0(1 - 2r_b^R) - r$, and jmp4-send $H(y_2)$ to P_0.
  – P_0 computes $x = e \oplus r_b = [\alpha_b]_1 \oplus [\alpha_b]_2 \oplus x_1$ and checks if $H(x^R - y_1) = H(y_2)$.
  – If verification fails, P_0 sets flag $= 1$, else it sets flag $= 0$. P_0 sends flag to P_1. Next, P_1, P_0 jmp4-send flag to P_2 and P_3. Servers set $\text{TTP} = P_1$ if flag $= 1$.

– If verification succeeds, servers locally convert $\langle e^R \rangle$ to $[\![e^R]\!]$ by setting their shares according to Table 9.

Online :

– Servers execute $\Pi_{\text{jsh4}}(P_0, P_1, P_2, c^R)$ where $c = \beta_b \oplus \gamma_b$.
– Servers execute $\Pi_{\text{mult4}}(\mathcal{P}, [\![e^R]\!], [\![c^R]\!])$ to generate $[\![e^R c^R]\!]$.
– Servers compute $[\![b^R]\!] = [\![e^R]\!] + [\![c^R]\!] - 2[\![e^R c^R]\!]$.

Figure 23: 4PC: Bit2A Protocol

Lemma C.8 (Communication). *Π_{bit2A4} (Fig. 23) requires an amortized communication of $3\ell + 4$ bits in the preprocessing phase, and 1 round with amortized communication of 3ℓ bits in the online phase.*

Proof. During preprocessing, one instance of Π_{ash4} requires 2ℓ bits of communication. Further, sending x_1 requires 1 bit, while sending y_1 requires ℓ bits. Sending of $H(y_2)$ can be amortized over several instances. Finally, communicating flag requires 3 bits. Thus, the overall amortized communication cost in preprocessing phase is $3\ell + 4$ bits. In the online phase, joint sharing of c^R can be performed non-interactively. The only cost is due to the online phase of multiplication which requires 3ℓ bits and one round. Thus, the amortized communication cost in the online phase is 3ℓ bits with one round of communication. $\qquad \square$

C.7.3 Bit Injection Protocol

Given the boolean sharing of a bit b, denoted as $[\![b]\!]^{\mathbf{B}}$, and the arithmetic sharing of $v \in \mathbb{Z}_{2^\ell}$, protocol Π_{bitinj4} computes $[\![\cdot]\!]$-sharing of bv. This can be naively computed by servers first executing Π_{bit2A4} on $[\![b]\!]^{\mathbf{B}}$ to generate $[\![b]\!]$, followed by servers computing $[\![bv]\!]$ by executing Π_{mult4} protocol on $[\![b]\!]$ and $[\![v]\!]$. Instead, we provide an optimized variant which helps in reducing the communication cost of the naive approach in, both, the preprocessing and online phase. We give the details next.

Let $z = b^R v$, where b^R denotes the value of b in $\mathbb{Z}_{2^\ell}$. Then, during the computation of $[\![z]\!]$, we observe the following:

$$
\begin{aligned}
z = b^R v &= (\alpha_b \oplus \beta_b)^R (\beta_v - \alpha_v) \\
&= ((\alpha_b \oplus \gamma_b) \oplus (\beta_b \oplus \gamma_b))^R ((\beta_v + \gamma_v) - (\alpha_v + \gamma_v)) \\
&= (c_b \oplus e_b)^R (c_v - e_v) = (c_b^R + e_b^R - 2c_b^R e_b^R)(c_v - e_v) \\
&= c_b^R c_v - c_b^R e_v + (c_v - 2c_b^R c_v)e_b^R + (2c_b^R - 1)e_b^R e_v
\end{aligned}
$$

where $c_b = \beta_b \oplus \gamma_b$, $e_b = \alpha_b \oplus \gamma_b$, $c_v = \beta_v + \gamma_v$ and $e_v = \alpha_v + \gamma_v$. The protocol proceeds with P_3 generating $\langle \cdot \rangle$-shares of e_b^R and $e_z = e_b^R e_v$, followed by verification of the same by P_0, P_1, P_2. If verification succeeds, then to enable P_2 to compute $\beta_z = z + \alpha_z$, P_1, P_0 jmp4-send the missing share of β_z to P_2. Similarly for P_1. Next, P_1, P_2 reconstruct β_z, and jmp4-send $\beta_z + \gamma_z$ to P_0 completing the protocol.

Protocol $\Pi_{\text{bitinj4}}(\mathcal{P}, [\![b]\!]^{\mathbf{B}}, [\![v]\!])$

Let $c_b = \beta_b \oplus \gamma_b$, $e_b = \alpha_b \oplus \gamma_b$, $c_v = \beta_v + \gamma_v$, $e_v = \alpha_v + \gamma_v$ and $e_z = e_b^R e_v$.

Preprocessing :

– P_0, P_3, P_j for $j \in \{1, 2\}$ sample $[\alpha_z]_1 \in \mathbb{Z}_{2^\ell}$ while P_1, P_2, P_3 sample $\gamma_z \in \mathbb{Z}_{2^\ell}$.

– Servers execute $\Pi_{\text{ash4}}(P_3, e_b^R)$ and $\Pi_{\text{ash4}}(P_3, e_z)$. Shares of $\langle e_v \rangle$ are set locally as $e_{v_0} = [\alpha_v]_2, e_{v_1} = [\alpha_v]_1, e_{v_3} = \gamma_v$.

– Servers verify correctness of $\langle e_b^R \rangle$ using steps similar to Π_{bit2A4} (Fig. 23). Correctness of $\langle e_z \rangle$ is verified as follows.

  – P_0, P_3, P_j for $j \in \{1, 2\}$ sample a random $r_j \in \mathbb{Z}_{2^\ell}$ while P_1, P_2, P_3 sample a random $r_0 \in \mathbb{Z}_{2^\ell}$. P_0, P_3 set $a_0 = r_1 - r_2$, P_1, P_3 set $a_1 = r_0 - r_1$ and P_2, P_3 set $a_2 = r_2 - r_0$.
  – P_1, P_3 compute $x_1 = e_{v_2} e_{b_2} + e_{v_1} e_{b_2} + e_{v_2} e_{b_1} + a_1$.
  – P_2, P_3 compute $x_2 = e_{v_0} e_{b_0} + e_{v_0} e_{b_2} + e_{v_2} e_{b_0} + a_2$.
  – P_0 computes $x_0 = e_{v_1} e_{b_1} + e_{v_1} e_{b_0} + e_{v_0} e_{b_1} + a_0$.
  – P_1, P_3 jmp4-send $y_1 = x_1 - e_{z_1}$ to P_0, while P_2, P_3 jmp4-send $H(-y_2)$ to P_0, where $y_2 = x_2 - e_{z_2}$.
  – P_0 computes $y_0 = x_0 - e_{z_0}$, and checks if $H(y_0 + y_1) = H(-y_2)$.
  – If verification fails, P_0 sets flag $= 1$, else it sets flag $= 0$. P_0 sends flag to P_1. Next, P_1, P_0 jmp4-send flag to P_2 and P_3. Servers set $\text{TTP} = P_1$ if flag $= 1$.

Online :

– P_0, P_1 compute $u_1 = -c_b^R e_{v_1} + (c_v - 2c_b^R c_v)e_{b_1}^R + (2c_b^R - 1)e_{z_1} + [\alpha_z]_1$, and jmp4-send u_1 to P_2.
– P_0, P_2 compute $u_2 = -c_b^R e_{v_0} + (c_v - 2c_b^R c_v)e_{b_0}^R + (2c_b^R - 1)e_{z_0} + [\alpha_z]_2$, and jmp4-send u_2 to P_1.
– P_1, P_2 compute $\beta_z = u_1 + u_2 - c_b^R e_{v_2} + (c_v - 2c_b^R c_v)e_{b_2}^R + (2c_b^R - 1)e_{z_2} + c_b^R c_v$.
– P_1, P_2 jmp4-send $\beta_z + \gamma_z$ to P_0.

Figure 24: 4PC: Bit Injection Protocol

Lemma C.9 (Communication). *Protocol Π_{bitinj4} requires an amortized communication cost of $6\ell + 4$ bits in the preprocessing phase, and requires 1 round with an amortized communication of 3ℓ bits in the online phase.*

Proof. The preprocessing phase requires two instances of Π_{ash4} which require 4ℓ bits of communication. Verifying correctness of $\langle e_b^R \rangle$ requires $\ell + 1$ bits, whereas for $\langle e_z \rangle$ we require ℓ bits. Finally, communicating the flag requires 3 bits. This results in the amortized communication of $6\ell + 4$ bits in the preprocessing phase. The online phase consists of three calls to Π_{jmp4} which requires 3ℓ bits of amortized communication. Note that the last call can be deferred towards the end of the computation, thereby requiring a single round for multiple instances. Thus, the number of rounds required in the online phase is one. □

C.8 Dot Product Protocol

The dot product protocol proceeds as follows. During the preprocessing phase, similar to the multiplication protocol P_0, P_1, P_3 sample a random $\left[\Gamma_{\vec{\mathbf{x}} \odot \vec{\mathbf{y}}}\right]_1$. P_0, P_3 compute $\Gamma_{\vec{\mathbf{x}} \odot \vec{\mathbf{y}}} = \sum_{i=1}^n \alpha_{\mathsf{x}_i} \alpha_{\mathsf{y}_i}$ and jmp4-send $\left[\Gamma_{\vec{\mathbf{x}} \odot \vec{\mathbf{y}}}\right]_2 = \Gamma_{\vec{\mathbf{x}} \odot \vec{\mathbf{y}}} - \left[\Gamma_{\vec{\mathbf{x}} \odot \vec{\mathbf{y}}}\right]_1$ to P_2. P_1, P_2, P_3 sample a random ψ, and generate its $[\cdot]$-shares locally. Servers P_3, P_j for $j \in \{1, 2\}$ then compute $[\chi]_j = \sum_{i=1}^n (\gamma_{\mathsf{x}_i} [\alpha_{\mathsf{y}_i}]_j + \gamma_{\mathsf{y}_i} [\alpha_{\mathsf{x}_i}]_j) + \left[\Gamma_{\vec{\mathbf{x}} \odot \vec{\mathbf{y}}}\right]_j - [\psi]_j$, and jmp4-send $[\chi]_j$ to P_0. The formal protocol is given in Fig. 25.

Protocol $\Pi_{\mathsf{dotp4}}(\mathcal{P}, \{[\![\mathsf{x}_i]\!], [\![\mathsf{y}_i]\!]\}_{i \in [n]})$

Preprocessing :

– P_0, P_3, P_j, for $j \in \{1, 2\}$, sample random $[\alpha_z]_j \in \mathbb{Z}_{2^\ell}$, while P_0, P_1, P_3 sample random $\left[\Gamma_{\vec{\mathbf{x}} \odot \vec{\mathbf{y}}}\right]_1 \in \mathbb{Z}_{2^\ell}$.

– P_1, P_2, P_3 together sample random $\gamma_z, \psi, r \in \mathbb{Z}_{2^\ell}$ and set $[\psi]_1 = r$, $[\psi]_2 = \psi - r$.

– P_0, P_3 compute $\left[\Gamma_{\vec{\mathbf{x}} \odot \vec{\mathbf{y}}}\right]_2 = \Gamma_{\vec{\mathbf{x}} \odot \vec{\mathbf{y}}} - \left[\Gamma_{\vec{\mathbf{x}} \odot \vec{\mathbf{y}}}\right]_1$, where $\Gamma_{\vec{\mathbf{x}} \odot \vec{\mathbf{y}}} = \sum_{i=1}^n \alpha_{\mathsf{x}_i} \alpha_{\mathsf{y}_i}$. P_0, P_3 jmp4-send $\left[\Gamma_{\vec{\mathbf{x}} \odot \vec{\mathbf{y}}}\right]_2$ to P_2.

– P_3, P_j, for $j \in \{1, 2\}$, set $[\chi]_j = \sum_{i=1}^n (\gamma_{\mathsf{x}_i} [\alpha_{\mathsf{y}_i}]_j + \gamma_{\mathsf{y}_i} [\alpha_{\mathsf{x}_i}]_j) + \left[\Gamma_{\vec{\mathbf{x}} \odot \vec{\mathbf{y}}}\right]_j - [\psi]_j$.

– P_1, P_3 jmp4-send $[\chi]_1$ to P_0, and P_2, P_3 jmp4-send $[\chi]_2$ to P_0.

Online :

– P_0, P_j, for $j \in \{1, 2\}$, compute $[\beta_z^\star]_j = -\sum_{i=1}^n ((\beta_{\mathsf{x}_i} + \gamma_{\mathsf{x}_i}) [\alpha_{\mathsf{y}_i}]_j + (\beta_{\mathsf{y}_i} + \gamma_{\mathsf{y}_i}) [\alpha_{\mathsf{x}_i}]_j) + [\alpha_z]_j + [\chi]_j$.

– P_1, P_0 jmp4-send $[\beta_z^\star]_1$ to P_2, while P_2, P_0 jmp4-send $[\beta_z^\star]_2$ to P_1.

– P_j for $j \in \{1, 2\}$ computes $\beta_z^\star = [\beta_z^\star]_1 + [\beta_z^\star]_2$ and sets $\beta_z = \beta_z^\star + \sum_{i=1}^n (\beta_{\mathsf{x}_i} \beta_{\mathsf{y}_i}) + \psi$.

– P_1, P_2 jmp4-send $\beta_z + \gamma_z$ to P_0.

Figure 25: 4PC: Dot Product Protocol ($z = \vec{\mathbf{x}} \odot \vec{\mathbf{y}}$)

Lemma C.10 (Communication). Π_{dotp4} *(Fig. 25) requires an amortized communication of 3ℓ bits in the preprocessing phase, and* 1 *round and an amortized communication of 3ℓ bits in the online phase.*

Proof. The preprocessing phase requires three calls to Π_{jmp4}, one to jmp4-send $\left[\Gamma_{\vec{\mathbf{x}} \odot \vec{\mathbf{y}}}\right]_2$ to P_2, and two to jmp4-send $[\chi]_1, [\chi]_2$ to P_0. Each invocation of Π_{jmp4} requires ℓ bits resulting in the amortized communication cost of preprocessing

phase to be 3ℓ bits. In the online phase, there are 2 parallel invocations of Π_{jmp4} to jmp4-send $[\beta_z^\star]_1, [\beta_z^\star]_2$ to P_2, P_1, respectively, which require amortized communication of 2ℓ bits and one round. This is followed by another call to Π_{jmp4} to jmp4-send $\beta_z + \gamma_z$ to P_0 which requires one more round and amortized communication of ℓ bits. As in the multiplication protocol, this step can be delayed till the end of the protocol and clubbed for multiple instances. Thus, the online phase requires one round and an amortized communication of 3ℓ bits. □

C.9 Truncation

Given the $[\![\cdot]\!]$-sharing of a value v, this protocol enables the servers to compute the $[\![\cdot]\!]$-sharing of the truncated value v^d (right shifted value by, say, d positions). Given $[\![\mathsf{v}]\!]$ and a random truncation pair $([\mathsf{r}], [\![\mathsf{r}^d]\!])$, the value $(\mathsf{v} - \mathsf{r})$ is opened, truncated and added to $[\![\mathsf{r}^d]\!]$ to obtain $[\![\mathsf{v}^d]\!]$. The protocol for generating the truncation pair $([\mathsf{r}], [\![\mathsf{r}^d]\!])$ is described in Fig. 26.

Protocol $\Pi_{\mathsf{trgen4}}(\mathcal{P})$

– P_0, P_3, P_j, for $j \in \{1, 2\}$ sample random $R_j \in \mathbb{Z}_{2^\ell}$. P_0, P_3 sets $\mathsf{r} = R_1 + R_2$ while P_j sets $[\mathsf{r}]_j = R_j$.

– P_0, P_3 locally truncate r to obtain r^d and execute $\Pi_{\mathsf{jsh4}}(P_0, P_3, \mathsf{r}^d)$ to generate $[\![\mathsf{r}^d]\!]$.
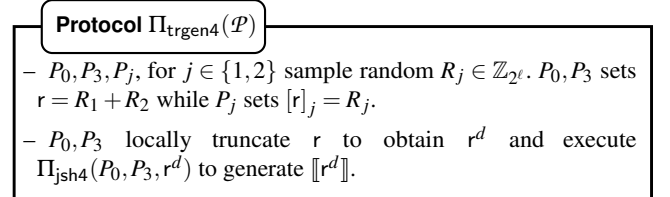
Figure 26: 4PC: Generating Random Truncated Pair $(\mathsf{r}, \mathsf{r}^d)$

Lemma C.11 (Communication). Π_{trgen4} *(Fig. 26) requires* 1 *round and an amortized communication of ℓ bits in the online phase.*

Proof. The cost follows directly from that of Π_{jmp4} (Lemma C.1 and §C.3). □

C.10 Dot Product with Truncation

Protocol Π_{dotpt4} (Fig. 27) enables servers to generate $[\![\cdot]\!]$-sharing of the truncated value of $z = \vec{\mathbf{x}} \odot \vec{\mathbf{y}}$, denoted as z^d, given the $[\![\cdot]\!]$-sharing of n-sized vectors $\vec{\mathbf{x}}$ and $\vec{\mathbf{y}}$.

Lemma C.12 (Communication). Π_{dotpt4} *(Fig. 27) requires an amortized communication of 4ℓ bits in the preprocessing phase, and* 1 *round with amortized communication of 3ℓ bits in the online phase.*

Proof. The preprocessing phase comprises of the preprocessing phase of Π_{dotp4} and Π_{trgen4} which results in an amortized communication of $3\ell + \ell = 4\ell$ bits. The online phase follows from that of Π_{dotp4} protocol except that, now, P_1, P_2 compute $[\cdot]$-shares of $z - \mathsf{r}$. This requires one round and an amortized communication cost of 2ℓ bits. P_1, P_2 then jointly share the truncated value of $z - \mathsf{r}$ with P_0, which requires 1 round and ℓ bits. However, this step can be deferred till the end for multiple instances, which results in amortizing this round. The total amortized communication is thus 3ℓ bits in online phase. □

Protocol $\Pi_{\text{dotpt4}}(\mathcal{P}, \{[\![x_i]\!], [\![y_i]\!]\}_{i \in [n]})$

Preprocessing :

– Servers execute the preprocessing phase of $\Pi_{\text{dotp4}}(\mathcal{P},$ $\{[\![x_i]\!], [\![y_i]\!]\}_{i \in [n]})$.
– Servers execute $\Pi_{\text{trgen4}}(\mathcal{P})$ to generate the truncation pair $([\![r]\!], [\![r^d]\!])$. P_0 obtains the value r in clear.

Online :

– P_0, P_j, for $j \in \{1,2\}$, compute $[\![\Psi]\!]_j = -\sum_{i=1}^{n}((\beta_{x_i} + \gamma_{x_i})[\alpha_{y_i}]_j + (\beta_{y_i} + \gamma_{y_i})[\alpha_{x_i}]_j) - [r]_j$ and sets $[(z-r)^{\star}]_j = [\![\Psi]\!]_j + [\chi]_j$.
– P_1, P_0 jmp-send $[(z-r)^{\star}]_1$ to P_2 and P_2, P_0 jmp4-send $[(z-r)^{\star}]_2$ to P_1.
– P_1, P_2 locally compute $(z-r)^{\star} = [(z-r)^{\star}]_1 + [(z-r)^{\star}]_2$ and set $(z-r) = (z-r)^{\star} + \sum_{i=1}^{n}(\beta_{x_i}\beta_{y_i}) + \psi$.
– P_1, P_2 locally truncate $(z-r)$ to obtain $(z-r)^d$ and execute $\Pi_{\text{jsh4}}(P_1, P_2, (z-r)^d)$ to generate $[\![(z-r)^d]\!]$.
– Servers locally compute $[\![z^d]\!] = [\![(z-r)^d]\!] + [\![r^d]\!]$.

Figure 27: 4PC: Dot Product Protocol with Truncation

C.11 Activation Functions

Lemma C.13 (Communication). *Protocol for* relu *requires an amortized communication of* $13\ell - 2$ *bits in the preprocessing phase and requires* $\log \ell + 1$ *rounds and an amortized communication of* $10\ell - 6$ *bits in the online phase.*

Proof. One instance of relu protocol comprises of execution of one instance of Π_{bitext4}, followed by Π_{bitinj4}. The cost, therefore, follows from Lemma C.7, and Lemma C.9. $\square$

Lemma C.14 (Communication). *Protocol for* sig *requires an amortized communication of* $23\ell - 1$ *bits in the preprocessing phase and requires* $\log \ell + 2$ *rounds and an amortized communication of* $20\ell - 9$ *bits in the online phase.*

Proof. An instance of sig protocol involves the execution of the following protocols in order– i) two parallel instances of Π_{bitext4} protocol, ii) one instance of Π_{mult4} protocol over boolean value, and iii) one instance of Π_{bitinj4} and Π_{bit2A4} in parallel. The cost follows from Lemma C.7, Lemma C.8 and Lemma C.9. $\square$

D Security Analysis of Our Protocols

In this section, we provide detailed security proofs for our constructions in both the 3PC and 4PC domains. We prove security using the real-world/ ideal-word simulation based technique. We provide proofs in the $\mathcal{F}_{\text{setup}}$-hybrid model for the case of 3PC, where $\mathcal{F}_{\text{setup}}$ (Fig. 6) denotes the ideal functionality for the three server shared-key setup. Similarly, the 4PC proofs are provided in the $\mathcal{F}_{\text{setup4}}$-hybrid model (Fig. 7).

Let $\mathcal{A}$ denote the real-world adversary corrupting at most one server in $\mathcal{P}$, and $\mathcal{S}$ denote the corresponding ideal world

adversary. The strategy for simulating the computation of function f (represented by a circuit ckt) is as follows: The simulation begins with the simulator emulating the shared-key setup ($\mathcal{F}_{\text{setup}}/\mathcal{F}_{\text{setup4}}$) functionality and giving the respective keys to the adversary. This is followed by the input sharing phase in which $\mathcal{S}$ extracts the input of $\mathcal{A}$, using the known keys, and sets the inputs of the honest parties to be 0. $\mathcal{S}$ now knows all the inputs and can compute all the intermediate values for each of the building blocks in clear. Also, $\mathcal{S}$ can obtain the output of the ckt in clear. $\mathcal{S}$ now proceeds simulating each of the building block in topological order using the aforementioned values (inputs of $\mathcal{A}$, intermediate values and circuit output).

In some of our sub protocols, adversary is able to decide on which among the honest parties should be chosen as the Trusted Third Party (TTP) in that execution of the protocol. To capture this, we consider *corruption-aware* functionalities [4] for the sub-protocols, where the functionality is provided the identity of the corrupt server as an auxiliary information.

For modularity, we provide the simulation steps for each of the sub-protocols separately. These steps, when carried out in the respective order, result in the simulation steps for the entire 3/4PC protocol. If a TTP is identified during the simulation of any of the sub-protocols, simulator will stop the simulation at that step. In the next round, the simulator receives the input of the corrupt party in clear on behalf of the TTP for the 3PC case, whereas it receives the input shares from adversary for 4PC.

D.1 Security Proofs for 3PC protocols

The ideal functionality $\mathcal{F}_{\text{3PC}}$ for evaluating ckt in the 3PC setting appears in Fig. 28.

Functionality $\mathcal{F}_{\text{3PC}}$

$\mathcal{F}_{\text{3PC}}$ interacts with the servers in $\mathcal{P}$ and the adversary $\mathcal{S}$. Let f denote the functionality to be computed. Let x_s be the input corresponding to the server P_s, and y_s be the corresponding output, i.e $(y_0, y_1, y_2) = f(x_0, x_1, x_2)$.

Step 1: $\mathcal{F}_{\text{3PC}}$ receives (Input, x_s) from $P_s \in \mathcal{P}$, and computes $(y_0, y_1, y_2) = f(x_0, x_1, x_2)$.

Step 2: $\mathcal{F}_{\text{3PC}}$ sends (Output, y_s) to $P_s \in \mathcal{P}$.

Figure 28: 3PC: Ideal functionality for evaluating a function f

D.1.1 Joint Message Passing (jmp) Protocol

This section provides the security proof for the jmp primitive, which forms the crux for achieving GOD in our constructions. Let $\mathcal{F}_{\text{jmp}}$ (Fig. 8) denote the ideal functionality and let $\mathcal{S}_{\text{jmp}}^{P_s}$ denote the corresponding simulator for the case of corrupt $P_s \in \mathcal{P}$. We begin with the case for a corrupt sender, P_i. The case for a corrupt P_j is similar and hence we omit details for the same.

Simulator $\mathcal{S}_{\mathsf{jmp}}^{P_i}$

- $\mathcal{S}_{\mathsf{jmp}}^{P_i}$ initializes $\mathsf{ttp} = \perp$ and receives v_i from $\mathcal{A}$ on behalf of P_k.

- In case, $\mathcal{A}$ fails to send a value $\mathcal{S}_{\mathsf{jmp}}^{P_i}$ broadcasts "$(\mathtt{accuse}, P_i)$", sets $\mathsf{ttp} = P_j$, $\mathsf{v}_i = \perp$, and skip to the last step.

- Else, it checks if $\mathsf{v}_i = \mathsf{v}$, where v is the value computed by $\mathcal{S}_{\mathsf{jmp}}^{P_i}$ based on the interaction with $\mathcal{A}$, and using the knowledge of the shared keys. If the values are equal, $\mathcal{S}_{\mathsf{jmp}}^{P_i}$ sets $b_k = 0$, else, sets $b_k = 1$, and sends the same to $\mathcal{A}$ on the behalf of P_k.

- If $\mathcal{A}$ broadcasts "$(\mathtt{accuse}, P_k)$", $\mathcal{S}_{\mathsf{jmp}}^{P_i}$ sets $\mathsf{v}_i = \perp$, $\mathsf{ttp} = P_j$, and skips to the last step.

- $\mathcal{S}_{\mathsf{jmp}}^{P_i}$ computes and sends b_j to $\mathcal{A}$ on behalf of P_j and receives $b_{\mathcal{A}}$ from $\mathcal{A}$ on behalf of honest P_j.

- If $\mathcal{S}_{\mathsf{jmp}}^{P_i}$ does not receive a $b_{\mathcal{A}}$ on behalf of P_j, it broadcasts "$(\mathtt{accuse}, P_i)$", sets $\mathsf{v}_i = \perp$, $\mathsf{ttp} = P_k$. If $\mathcal{A}$ broadcasts "$(\mathtt{accuse}, P_j)$", $\mathcal{S}_{\mathsf{jmp}}^{P_i}$ sets $\mathsf{v}_i = \perp$, $\mathsf{ttp} = P_k$. If ttp is set, skip to the last step.

- If $(\mathsf{v}_i = \mathsf{v})$ and $b_{\mathcal{A}} = 1$, $\mathcal{S}_{\mathsf{jmp}}^{P_i}$ broadcasts $\mathsf{H}_j = \mathsf{H}(\mathsf{v})$ on behalf of P_j.

- Else if $\mathsf{v}_i \neq \mathsf{v}_j : \mathcal{S}_{\mathsf{jmp}}^{P_i}$ broadcasts $\mathsf{H}_j = \mathsf{H}(\mathsf{v})$ and $\mathsf{H}_k = \mathsf{H}(\mathsf{v}_i)$ on behalf of P_j and P_k, respectively. If $\mathcal{A}$ does not broadcast, $\mathcal{S}_{\mathsf{jmp}}^{P_i}$ sets $\mathsf{ttp} = P_k$. Else if, $\mathcal{A}$ broadcasts a value $\mathsf{H}_{\mathcal{A}}$:

  - If $\mathsf{H}_{\mathcal{A}} \neq \mathsf{H}_j : \mathcal{S}_{\mathsf{jmp}}^{P_i}$ sets $\mathsf{ttp} = P_k$.
  - Else if $\mathsf{H}_{\mathcal{A}} \neq \mathsf{H}_k : \mathcal{S}_{\mathsf{jmp}}^{P_i}$ sets $\mathsf{ttp} = P_j$.

- $\mathcal{S}_{\mathsf{jmp}}^{P_i}$ invokes $\mathcal{F}_{\mathsf{jmp}}$ on $(\mathsf{Input}, \mathsf{v}_i)$ and $(\mathsf{Select}, \mathsf{ttp})$ on behalf of $\mathcal{A}$.

Figure 29: Simulator $\mathcal{S}_{\mathsf{jmp}}^{P_i}$ for corrupt sender P_i

The case for a corrupt receiver, P_k is provided in Fig. 30.

Simulator $\mathcal{S}_{\mathsf{jmp}}^{P_k}$

- $\mathcal{S}_{\mathsf{jmp}}^{P_k}$ initializes $\mathsf{ttp} = \perp$, computes v honestly and sends v and $\mathsf{H}(\mathsf{v})$ to $\mathcal{A}$ on behalf of P_i and P_j, respectively.

- If $\mathcal{A}$ broadcasts "$(\mathtt{accuse}, P_i)$", set $\mathsf{ttp} = P_j$, else if $\mathcal{A}$ broadcasts "$(\mathtt{accuse}, P_j)$", set $\mathsf{ttp} = P_i$. If both messages are broadcast, set $\mathsf{ttp} = P_i$. If ttp is set skip to the last step.

- On behalf of P_i, P_j, $\mathcal{S}_{\mathsf{jmp}}^{P_k}$ receives $b_{\mathcal{A}}$ from $\mathcal{A}$. Let b_i (resp. b_j) denote the bit received by P_i (resp. P_j) from $\mathcal{A}$.

- If $\mathcal{A}$ failed to send bit $b_{\mathcal{A}}$ to P_i, $\mathcal{S}_{\mathsf{jmp}}^{P_k}$ broadcasts "$(\mathtt{accuse}, P_k)$", set $\mathsf{ttp} = P_j$. Similarly, for P_j. If both P_i, P_j broadcast "$(\mathtt{accuse}, P_k)$", set $\mathsf{ttp} = P_i$. If ttp is set, skip to the last step.

- If $b_i \vee b_j = 1 : \mathcal{S}_{\mathsf{jmp}}^{P_k}$ broadcasts $\mathsf{H}_i, \mathsf{H}_j$ where $\mathsf{H}_i = \mathsf{H}_j = \mathsf{H}(\mathsf{v})$ on behalf of P_i, P_j, respectively.

- If $\mathcal{A}$ does not broadcast $\mathcal{S}_{\mathsf{jmp}}^{P_k}$ sets $\mathsf{ttp} = \perp$. If $\mathcal{A}$ broadcasts a value $\mathsf{H}_{\mathcal{A}}$:

  - If $\mathsf{H}_{\mathcal{A}} \neq \mathsf{H}_i : \mathcal{S}_{\mathsf{jmp}}^{P_k}$ sets $\mathsf{ttp} = P_j$.
  - Else if $\mathsf{H}_{\mathcal{A}} = \mathsf{H}_i = \mathsf{H}_j : \mathcal{S}_{\mathsf{jmp}}^{P_k}$ sets $\mathsf{ttp} = P_i$.

- $\mathcal{S}_{\mathsf{jmp}}^{P_k}$ invokes $\mathcal{F}_{\mathsf{jmp}}$ on $(\mathsf{Input}, \perp)$ and $(\mathsf{Select}, \mathsf{ttp})$ on behalf of $\mathcal{A}$.

Figure 30: Simulator $\mathcal{S}_{\mathsf{jmp}}^{P_k}$ for corrupt receiver P_k

D.1.2 Sharing Protocol

The case for a corrupt P_0 is provided in Fig. 36.

Simulator $\mathcal{S}_{\mathsf{sh}}^{P_0}$

Preprocessing: $\mathcal{S}_{\mathsf{sh}}^{P_0}$ emulates $\mathcal{F}_{\mathsf{setup}}$ and gives the keys $(k_{01}, k_{02}, k_{\mathcal{P}})$ to $\mathcal{A}$. The values that are commonly held along with $\mathcal{A}$ are sampled using appropriate shared key. Otherwise, values are sampled randomly.

Online:

- If the dealer $P_s = P_0$:
  - $\mathcal{S}_{\mathsf{sh}}^{P_0}$ receives β_v on behalf of P_1 and sets $\mathsf{msg} = \mathsf{v}$ accordingly.
  - Steps for Π_{jmp} protocol are simulated according to $\mathcal{S}_{\mathsf{jmp}}^{P_i}$ (Fig. 29), where P_0 plays the role of one of the senders.
- If the dealer $P_s = P_1$:
  - $\mathcal{S}_{\mathsf{sh}}^{P_0}$ sets $\mathsf{v} = 0$ by assigning $\beta_\mathsf{v} = \alpha_\mathsf{v}$.
  - Steps for Π_{jmp} protocol are simulated similar to $\mathcal{S}_{\mathsf{jmp}}^{P_k}$ (Fig. 30), with P_0 acting as the receiver.
- If the dealer if P_2 : Similar to the case when $P_s = P_1$.

Figure 31: Simulator $\mathcal{S}_{\mathsf{sh}}^{P_0}$ for corrupt P_0

The case for a corrupt P_1 is provided in Fig. 32. The case for a corrupt P_2 is similar.

Simulator $\mathcal{S}_{\mathsf{sh}}^{P_1}$

Preprocessing: $\mathcal{S}_{\mathsf{jsh}}^{P_1}$ emulates $\mathcal{F}_{\mathsf{setup}}$ and gives the keys $(k_{01}, k_{12}, k_{\mathcal{P}})$ to $\mathcal{A}$. The values that are commonly held along with $\mathcal{A}$ are sampled using appropriate shared key. Otherwise, values are sampled randomly.

Online:

- If dealer $P_s = P_1 : \mathcal{S}_{\mathsf{sh}}^{P_1}$ receives β_v from $\mathcal{A}$ on behalf of P_2.
- If $P_s = P_0 : \mathcal{S}_{\mathsf{sh}}^{P_1}$ sets $\mathsf{v} = 0$ by assigning $\beta_\mathsf{v} = \alpha_\mathsf{v}$ and sends β_v to $\mathcal{A}$ on behalf of P_s.
- If $P_s = P_2$: Similar to the case where $P_s = P_0$.
- Steps of Π_{jmp}, in all the steps above, are simulated similar to $\mathcal{S}_{\mathsf{jmp}}^{P_i}$ (Fig. 29), ie. the case of corrupt sender.

Figure 32: Simulator $\mathcal{S}_{\mathsf{sh}}^{P_1}$ for corrupt P_1

D.1.3 Multiplication Protocol

The case for a corrupt P_0 is provided in Fig. 33.

Simulator $\mathcal{S}_{\mathsf{mult}}^{P_0}$

Preprocessing:

- $\mathcal{S}_{\mathsf{mult}}^{P_0}$ samples $[\alpha_\mathsf{z}]_1, [\alpha_\mathsf{z}]_2$ and γ_z on behalf of P_1, P_2 and generates the $\langle \cdot \rangle$-shares of d, e honestly.

- $\mathcal{S}_{\mathsf{mult}}^{P_0}$ emulates $\mathcal{F}_{\mathsf{MulPre}}$, and extracts $\psi, [\chi]_1, [\chi]_2$ on behalf of P_1, P_2.

Online:

- $\mathcal{S}_{\mathsf{mult}}^{P_0}$ computes $[\beta_\mathsf{z}^\star]_1, [\beta_\mathsf{z}^\star]_2$ and steps of Π_{jmp} are simulated

according to $\mathcal{S}_{\mathsf{jmp}}^{P_i}$ with $\mathcal{A}$ as one of the sender for both $[\beta_{\mathsf{z}}^{\star}]_1$, and $[\beta_{\mathsf{z}}^{\star}]_2$.

– $\mathcal{S}_{\mathsf{mult}}^{P_0}$ computes $\beta_{\mathsf{z}} + \gamma_{\mathsf{z}}$ on behalf of P_1, P_2 and steps of Π_{jmp} are simulated according to $\mathcal{S}_{\mathsf{jmp}}^{P_k}$ with $\mathcal{A}$ as the receiver for $\beta_{\mathsf{z}} + \gamma_{\mathsf{z}}$.

Figure 33: Simulator $\mathcal{S}_{\mathsf{mult}}^{P_0}$ for corrupt P_0

The case for a corrupt P_1 is provided in Fig. 34. The case for a corrupt P_2 is similar.

Simulator $\mathcal{S}_{\mathsf{mult}}^{P_1}$

Preprocessing:

– $\mathcal{S}_{\mathsf{mult}}^{P_1}$ samples $[\alpha_{\mathsf{z}}]_1, \gamma_{\mathsf{z}}$ and $[\alpha_{\mathsf{z}}]_2$ on behalf of P_0, P_2. $\mathcal{S}_{\mathsf{mult}}^{P_1}$ generates the $\langle \cdot \rangle$-shares of d, e honestly.

– $\mathcal{S}_{\mathsf{mult}}^{P_1}$ emulates $\mathcal{F}_{\mathsf{MulPre}}$, and extracts $\psi, [\chi]_1, [\chi]_2$ on behalf of P_0, P_2.

Online:

– $\mathcal{S}_{\mathsf{mult}}^{P_1}$ computes $[\beta_{\mathsf{z}}^{\star}]_1, [\beta_{\mathsf{z}}^{\star}]_2$ on behalf of P_0, P_2, and steps of Π_{jmp} are simulated according to $\mathcal{S}_{\mathsf{jmp}}^{P_i}$ with $\mathcal{A}$ as one of the sender for $[\beta_{\mathsf{z}}^{\star}]_1$, and as the receiver for $[\beta_{\mathsf{z}}^{\star}]_2$.

– $\mathcal{S}_{\mathsf{mult}}^{P_1}$ computes $\beta_{\mathsf{z}} + \gamma_{\mathsf{z}}$ on behalf of P_2 and steps of Π_{jmp} are simulated according to $\mathcal{S}_{\mathsf{jmp}}^{P_i}$ with $\mathcal{A}$ one of the senders for $\beta_{\mathsf{z}} + \gamma_{\mathsf{z}}$.

Figure 34: Simulator $\mathcal{S}_{\mathsf{mult}}^{P_1}$ for corrupt P_1

D.1.4 Reconstruction Protocol

The case for a corrupt P_0 is provided in Fig. 51. The case for a corrupt P_1, P_2 is similar.

Simulator $\mathcal{S}_{\mathsf{rec}}$

Preprocessing:

– $\mathcal{S}_{\mathsf{rec}}$ computes commitments on $[\alpha_{\mathsf{v}}]_1, [\alpha_{\mathsf{v}}]_2$ and γ_{v} on behalf of P_1, P_2, using the respective shared keys.

– The steps of Π_{jmp} are simulated similar to $\mathcal{S}_{\mathsf{jmp}}^{P_k}$ with $\mathcal{A}$ acting as the receiver for $\mathsf{Com}(\gamma_{\mathsf{v}})$, and $\mathcal{S}_{\mathsf{jmp}}^{P_i}$ with $\mathcal{A}$ acting as one of the senders for $\mathsf{Com}([\alpha_{\mathsf{v}}]_1)$ and $\mathsf{Com}([\alpha_{\mathsf{v}}]_2)$.

Online:

– $\mathcal{S}_{\mathsf{rec}}$ receives openings for $\mathsf{Com}([\alpha_{\mathsf{v}}]_1), \mathsf{Com}([\alpha_{\mathsf{v}}]_2)$ on behalf of P_2 and P_1, respectively.

– $\mathcal{S}_{\mathsf{rec}}$ opens $\mathsf{Com}(\gamma_{\mathsf{v}})$ to $\mathcal{A}$ on behalf of P_1, P_2.

Figure 35: Simulator $\mathcal{S}_{\mathsf{rec}}$ for corrupt P_0

D.1.5 Joint Sharing Protocol

The case for a corrupt P_0 is provided in Fig. 36. The case for a corrupt P_1, P_2 is similar.

Simulator $\mathcal{S}_{\mathsf{jsh}}$

Preprocessing: $\mathcal{S}_{\mathsf{sh}}^{P_0}$ emulates $\mathcal{F}_{\mathsf{setup}}$ and gives the keys $(k_{01}, k_{02}, k_{\mathcal{P}})$ to $\mathcal{A}$. The values that are commonly held along with $\mathcal{A}$ are sampled using appropriate shared key. Otherwise, values are sampled randomly.

Online:

– If $(P_i, P_j) = (P_1, P_0)$: $\mathcal{S}_{\mathsf{jsh}}$ computes $\beta_{\mathsf{v}} = \mathsf{v} + \alpha_{\mathsf{v}}$ on behalf of P_1. The steps of Π_{jmp} are simulated similar to $\mathcal{S}_{\mathsf{jmp}}^{P_i}$, where the $\mathcal{A}$ acts as one of the senders.

– If $(P_i, P_j) = (P_2, P_0)$: Similar to the case when $(P_i, P_j) = (P_1, P_0)$.

– If $(P_i, P_j) = (P_1, P_2)$: $\mathcal{S}_{\mathsf{jsh}}$ sets $\mathsf{v} = 0$ by setting $\beta_{\mathsf{v}} = \alpha_{\mathsf{v}}$. The steps of Π_{jmp} are simulated similar to $\mathcal{S}_{\mathsf{jmp}}^{P_k}$, where the $\mathcal{A}$ acts as the receiver.

Figure 36: Simulator $\mathcal{S}_{\mathsf{jsh}}$ for corrupt P_0

D.1.6 Dot Product Protocol

The case for a corrupt P_0 is provided in Fig. 37.

Simulator $\mathcal{S}_{\mathsf{dotp}}^{P_0}$

Preprocessing: $\mathcal{S}_{\mathsf{dotp}}$ emulates $\mathcal{F}_{\mathsf{DotPPre}}$ and derives ψ and respective $[\cdot]$-shares of χ honestly on behalf of P_1, P_2.

Online:

– $\mathcal{S}_{\mathsf{dotp}}^{P_0}$ computes $[\cdot]$-shares of $\beta_{\mathsf{z}}^{\star}$ on behalf of P_1, P_2. The steps of Π_{jmp}, required to provide P_1 with $[\beta_{\mathsf{z}}^{\star}]_2$, and P_2 with $[\beta_{\mathsf{z}}^{\star}]_1$, are simulated similar to $\mathcal{S}_{\mathsf{jmp}}^{P_i}$, where P_0 acts as one of the sender in both cases.

– $\mathcal{S}_{\mathsf{dotp}}^{P_0}$ computes $\beta_{\mathsf{z}}^{\star}$ and β_{z} on behalf of P_1, P_2. The steps of the Π_{jmp}, required to provide P_0 with $\beta_{\mathsf{z}} + \gamma_{\mathsf{z}}$, are simulated similar to $\mathcal{S}_{\mathsf{jmp}}^{P_k}$, where P_0 acts as the receiver.

Figure 37: Simulator $\mathcal{S}_{\mathsf{dotp}}$ for corrupt P_0

The case for a corrupt P_1 is provided in Fig. 38. The case for a corrupt P_2 is similar.

Simulator $\mathcal{S}_{\mathsf{dotp}}^{P_1}$

Preprocessing: $\mathcal{S}_{\mathsf{dotp}}^{P_1}$ emulates $\mathcal{F}_{\mathsf{DotPPre}}$ and derives $[\cdot]$-shares of ψ, χ honestly on behalf of P_0, P_2.

Online:

– $\mathcal{S}_{\mathsf{dotp}}^{P_1}$ computes $[\beta_{\mathsf{z}}^{\star}]_1$ on behalf of P_0, and $[\beta_{\mathsf{z}}^{\star}]_2$ on behalf of P_0 and P_2. The steps of Π_{jmp}, required to provide P_1 with $[\beta_{\mathsf{z}}^{\star}]_2$, and P_2 with $[\beta_{\mathsf{z}}^{\star}]_1$, are simulated similar to $\mathcal{S}_{\mathsf{jmp}}^{P_i}$, where $\mathcal{A}$ acts as one of the sender in the former case, and as a receiver in the latter case.

– $\mathcal{S}_{\mathsf{dotp}}^{P_1}$ computes $\beta_{\mathsf{z}}^{\star}$ and β_{z} on behalf of P_2. The steps of Π_{jmp}, required to provide P_0 with $\beta_{\mathsf{z}} + \gamma_{\mathsf{z}}$, are simulated similar to $\mathcal{S}_{\mathsf{jmp}}^{P_i}$, where $\mathcal{A}$ acts as one of the sender.

Figure 38: Simulator $\mathcal{S}_{\mathsf{dotp}}$ for corrupt P_1

D.1.7 Truncation Protocol

The case for a corrupt P_0 is provided in Fig. 39.

Simulator $\mathcal{S}_{\mathsf{trgen}}^{P_0}$

- For $i \in \{0, \ldots, \ell-1\}$, for $j \in \{1,2\}$, $\mathcal{S}_{\mathsf{trgen}}^{P_0}$ samples $\mathsf{r}_j[i]$ on behalf of P_j along with $\mathcal{A}$ using respective keys.
- $\mathcal{S}_{\mathsf{trgen}}^{P_0}$ acting on behalf of P_1, P_2 generates $\llbracket \cdot \rrbracket$-shares of $(\mathsf{r}_j[i])^{\mathsf{R}}$ for $i \in \{0, \ldots, \ell-1\}, j \in \{1,2\}$ non-interactively.
- $\mathcal{S}_{\mathsf{trgen}}^{P_0}$ defines $\vec{\mathbf{x}}, \vec{\mathbf{y}}, \vec{\mathbf{p}}$ and $\vec{\mathbf{q}}$ as per Fig. 16. The steps for Π_{dotp} are simulated similar to $\mathcal{S}_{\Pi_{\mathsf{dotp}}}^{P_0}$ for generating A, B.
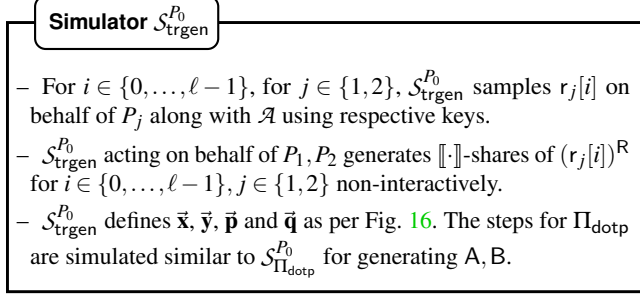

Figure 39: Simulator $\mathcal{S}_{\mathsf{trgen}}^{P_0}$ for corrupt P_0

The case for a corrupt P_1 is provided in Fig. 40. The case for a corrupt P_2 is similar.

Simulator $\mathcal{S}_{\mathsf{trgen}}^{P_1}$

- For $i \in \{0, \ldots, \ell-1\}$, $\mathcal{S}_{\mathsf{trgen}}^{P_1}$ samples $\mathsf{r}_1[i]$ on behalf of P_0 along with $\mathcal{A}$, using respective keys, and it samples $\mathsf{r}_2[i]$ randomly on behalf of P_0, P_2.
- $\mathcal{S}_{\mathsf{trgen}}^{P_1}$ acting on behalf of P_0, P_2 generates $\llbracket \cdot \rrbracket$-shares of $(\mathsf{r}_j[i])^{\mathsf{R}}$ for $i \in \{0, \ldots, \ell-1\}, j \in \{1,2\}$ non-interactively.
- $\mathcal{S}_{\mathsf{trgen}}^{P_1}$ defines $\vec{\mathbf{x}}, \vec{\mathbf{y}}, \vec{\mathbf{p}}$ and $\vec{\mathbf{q}}$ as per Fig. 16. The steps for Π_{dotp} are simulated similar to $\mathcal{S}_{\Pi_{\mathsf{dotp}}}^{P_1}$ for generating A, B.
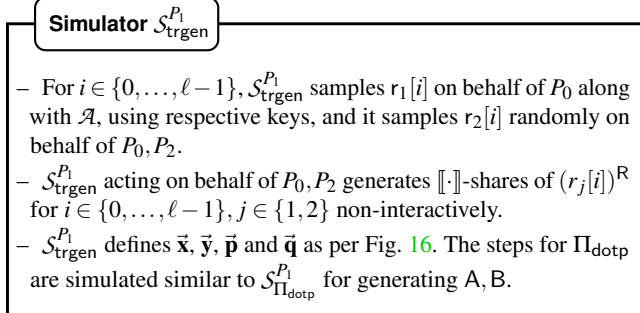

Figure 40: Simulator $\mathcal{S}_{\mathsf{trgen}}^{P_1}$ for corrupt P_1

D.2 Security Proofs for 4PC protocols

The ideal functionality $\mathcal{F}_{\mathsf{4PC}}$ for evaluating a function f to be computed by ckt in 4PC appears in Fig. 41.

Functionality $\mathcal{F}_{\mathsf{4PC}}$

$\mathcal{F}_{\mathsf{4PC}}$ interacts with the servers in $\mathcal{P}$ and the adversary $\mathcal{S}$. Let f denote the function to be computed. Let x_s be the input corresponding to the server P_s, and y_s be the corresponding output, i.e $(\mathsf{y}_0, \mathsf{y}_1, \mathsf{y}_2, \mathsf{y}_3) = f(\mathsf{x}_0, \mathsf{x}_1, \mathsf{x}_2, \mathsf{x}_3)$.

Step 1: $\mathcal{F}_{\mathsf{4PC}}$ receives $(\mathsf{Input}, \mathsf{x}_s)$ from $P_s \in \mathcal{P}$, and computes $(\mathsf{y}_0, \mathsf{y}_1, \mathsf{y}_3, \mathsf{y}_3) = f(\mathsf{x}_0, \mathsf{x}_1, \mathsf{x}_2, \mathsf{x}_3)$.

Step 2: $\mathcal{F}_{\mathsf{4PC}}$ sends $(\mathsf{Output}, \mathsf{y}_s)$ to $P_s \in \mathcal{P}$.
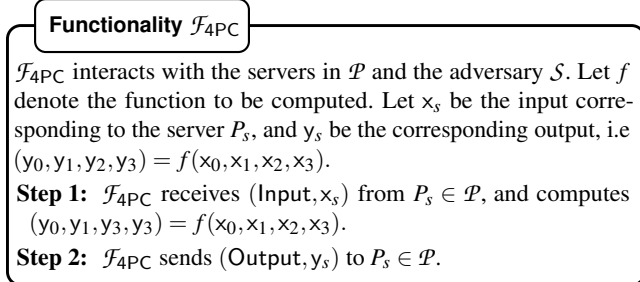

Figure 41: 4PC: Ideal functionality for computing f in 4PC setting

D.2.1 Joint Message Passing

Let $\mathcal{F}_{\mathsf{jmp4}}$ Fig. 18 denote the ideal functionality and let $\mathcal{S}_{\mathsf{jmp4}}^{P_s}$ denote the corresponding simulator for the case of corrupt $P_s \in \mathcal{P}$.

We begin with the case for a corrupt sender, P_i, which is provided in Fig. 42. The case for a corrupt P_j is similar and hence we omit details for the same.
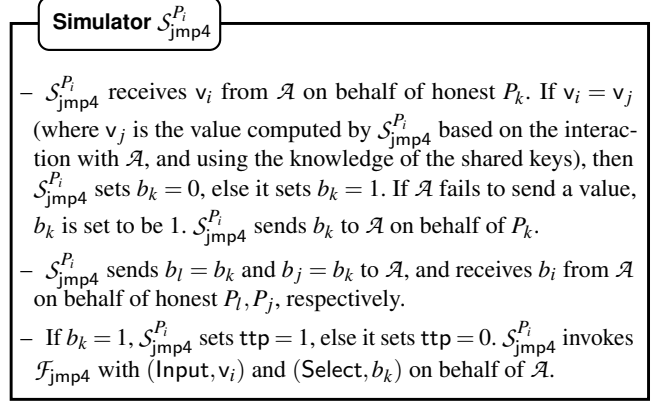

Simulator $\mathcal{S}_{\mathsf{jmp4}}^{P_i}$

- $\mathcal{S}_{\mathsf{jmp4}}^{P_i}$ receives v_i from $\mathcal{A}$ on behalf of honest P_k. If $\mathsf{v}_i = \mathsf{v}_j$ (where v_j is the value computed by $\mathcal{S}_{\mathsf{jmp4}}^{P_i}$ based on the interaction with $\mathcal{A}$, and using the knowledge of the shared keys), then $\mathcal{S}_{\mathsf{jmp4}}^{P_i}$ sets $b_k = 0$, else it sets $b_k = 1$. If $\mathcal{A}$ fails to send a value, b_k is set to be 1. $\mathcal{S}_{\mathsf{jmp4}}^{P_i}$ sends b_k to $\mathcal{A}$ on behalf of P_k.
- $\mathcal{S}_{\mathsf{jmp4}}^{P_i}$ sends $b_l = b_k$ and $b_j = b_k$ to $\mathcal{A}$, and receives b_i from $\mathcal{A}$ on behalf of honest P_l, P_j, respectively.
- If $b_k = 1$, $\mathcal{S}_{\mathsf{jmp4}}^{P_i}$ sets $\mathsf{ttp} = 1$, else it sets $\mathsf{ttp} = 0$. $\mathcal{S}_{\mathsf{jmp4}}^{P_i}$ invokes $\mathcal{F}_{\mathsf{jmp4}}$ with $(\mathsf{Input}, \mathsf{v}_i)$ and (Select, b_k) on behalf of $\mathcal{A}$.

Figure 42: Simulator $\mathcal{S}_{\mathsf{jmp4}}^{P_i}$ for corrupt sender P_i

The case for a corrupt receiver, P_k is provided in Fig. 43.

Simulator $\mathcal{S}_{\mathsf{jmp4}}^{P_k}$

- $\mathcal{S}_{\mathsf{jmp4}}^{P_k}$ sends $\mathsf{v}, \mathsf{H}(\mathsf{v})$ (where v is the value computed by $\mathcal{S}_{\mathsf{jmp4}}^{P_k}$ based on the interaction with $\mathcal{A}$, and using the knowledge of the shared keys) to $\mathcal{A}$ on behalf of honest P_i, P_j, respectively.
- $\mathcal{S}_{\mathsf{jmp4}}^{P_k}$ receives b_{ki}, b_{kj}, b_{kl} from $\mathcal{A}$ on behalf of P_i, P_j, P_l, respectively. If $\mathcal{A}$ fails to send a value, it is assumed to be 1.
- $\mathcal{S}_{\mathsf{jmp4}}^{P_k}$ sets b_k to be majority value in b_{ki}, b_{kj}, b_{kl}. If $b_k = 1$, $\mathcal{S}_{\mathsf{jmp4}}^{P_k}$ sets $\mathsf{ttp} = 1$, else it sets $\mathsf{ttp} = 0$. $\mathcal{S}_{\mathsf{jmp4}}^{P_k}$ invokes $\mathcal{F}_{\mathsf{jmp4}}$ with $(\mathsf{Input}, \perp)$ and (Select, b_k) on behalf of $\mathcal{A}$.
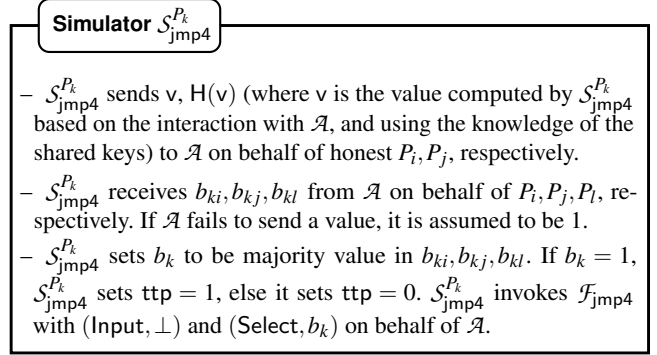

Figure 43: Simulator $\mathcal{S}_{\mathsf{jmp4}}^{P_k}$ for corrupt receiver P_k

The case for a corrupt receiver, P_l, which is the server outside the computation involved in Π_{jmp4}, is provided in Fig. 44.

Simulator $\mathcal{S}_{\mathsf{jmp4}}^{P_l}$

- $\mathcal{S}_{\mathsf{jmp4}}^{P_l}$ sends $b_k = 0$ followed by $b_i = 0, b_j = 0$ to $\mathcal{A}$ on behalf of P_k and P_i, P_j, respectively.
- $\mathcal{S}_{\mathsf{jmp4}}^{P_l}$ invokes $\mathcal{F}_{\mathsf{jmp4}}$ with $(\mathsf{Input}, \perp)$ and (Select, b_k) on behalf of $\mathcal{A}$.
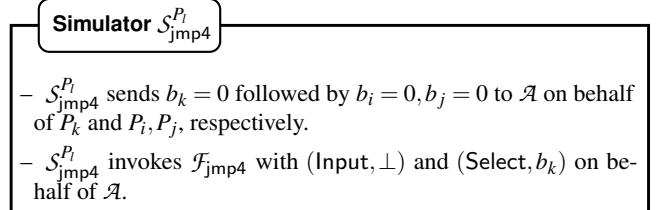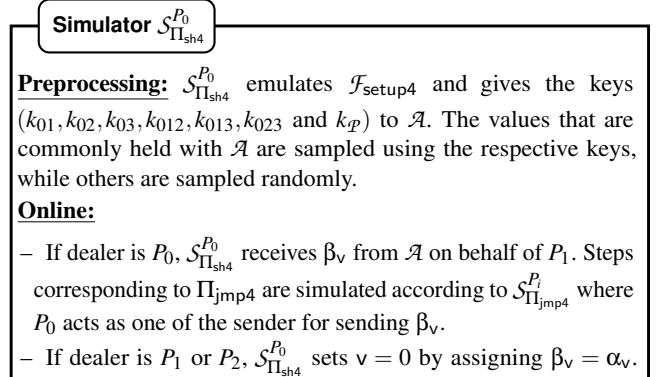

Figure 44: Simulator $\mathcal{S}_{\mathsf{jmp4}}^{P_l}$ for corrupt fourth server P_l

D.2.2 Sharing Protocol

The case for corrupt P_0 is given in Fig. 45.

Simulator $\mathcal{S}_{\Pi_{\mathsf{sh4}}}^{P_0}$

Preprocessing: $\mathcal{S}_{\Pi_{\mathsf{sh4}}}^{P_0}$ emulates $\mathcal{F}_{\mathsf{setup4}}$ and gives the keys $(k_{01}, k_{02}, k_{03}, k_{012}, k_{013}, k_{023}$ and $k_{\mathcal{P}})$ to $\mathcal{A}$. The values that are commonly held with $\mathcal{A}$ are sampled using the respective keys, while others are sampled randomly.

Online:

- If dealer is P_0, $\mathcal{S}_{\Pi_{\mathsf{sh4}}}^{P_0}$ receives β_v from $\mathcal{A}$ on behalf of P_1. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_i}$ where P_0 acts as one of the sender for sending β_v.
- If dealer is P_1 or P_2, $\mathcal{S}_{\Pi_{\mathsf{sh4}}}^{P_0}$ sets $\mathsf{v} = 0$ by assigning $\beta_\mathsf{v} = \alpha_\mathsf{v}$.

Steps corresponding to Π_{jmp4} for sending $\beta_\mathsf{v} + \gamma_\mathsf{v}$ to $\mathcal{A}$ are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_k}$ where P_0 acts as the receiver.

– If dealer is P_3, $\mathcal{S}_{\Pi_{\mathsf{sh4}}}^{P_0}$ sets $\mathsf{v} = 0$ by assigning $\beta_\mathsf{v} = \alpha_\mathsf{v}$. $\mathcal{S}_{\Pi_{\mathsf{sh4}}}^{P_0}$ sends $\beta_\mathsf{v} + \gamma_\mathsf{v}$ to $\mathcal{A}$ on behalf of P_3. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_j}$ where P_0 acts as one of the sender with P_1, P_2 as the receivers, separately.

Figure 45: Simulator $\mathcal{S}_{\Pi_{\mathsf{sh4}}}^{P_0}$ for corrupt P_0

The case for corrupt P_1 is given in Fig. 46. The case for a corrupt P_2 is similar.

Simulator $\mathcal{S}_{\Pi_{\mathsf{sh4}}}^{P_1}$

Preprocessing: $\mathcal{S}_{\Pi_{\mathsf{sh4}}}^{P_1}$ emulates $\mathcal{F}_{\mathsf{setup4}}$ and gives the keys $(k_{01}, k_{12}, k_{13}, k_{012}, k_{013}, k_{123}$ and $k_\mathcal{P})$ to $\mathcal{A}$. The values that are commonly held with $\mathcal{A}$ are sampled using the respective keys, while others are sampled randomly.

Online:

– If dealer is P_1, $\mathcal{S}_{\Pi_{\mathsf{sh4}}}^{P_1}$ receives β_v from $\mathcal{A}$ on behalf of P_2. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_i}$ where P_1 acts as one of the sender for sending $\beta_\mathsf{v} + \gamma_\mathsf{v}$ to P_0.

– If dealer is P_0 or P_2, $\mathcal{S}_{\Pi_{\mathsf{sh4}}}^{P_1}$ sets $\mathsf{v} = 0$ by assigning $\beta_\mathsf{v} = \alpha_\mathsf{v}$.

  • If dealer is P_0, $\mathcal{S}_{\Pi_{\mathsf{sh4}}}^{P_1}$ sends β_v to $\mathcal{A}$ on behalf of P_0. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_j}$ where P_1 acts as one of the sender to send β_v.

  • If dealer is P_2, $\mathcal{S}_{\Pi_{\mathsf{sh4}}}^{P_1}$ sends β_v to $\mathcal{A}$ on behalf of P_2. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_i}$ where P_1 acts as one of the sender to send $\beta_\mathsf{v} + \gamma_\mathsf{v}$.

– If dealer is P_3, $\mathcal{S}_{\Pi_{\mathsf{sh4}}}^{P_1}$ sets $\mathsf{v} = 0$ by assigning $\beta_\mathsf{v} = \alpha_\mathsf{v}$. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_k}$ where P_1 acts as the receiver for receiving $\beta_\mathsf{v} + \gamma_\mathsf{v}$.

Figure 46: Simulator $\mathcal{S}_{\Pi_{\mathsf{sh4}}}^{P_1}$ for corrupt P_1

The case for corrupt P_3 is given in Fig. 47.

Simulator $\mathcal{S}_{\Pi_{\mathsf{sh4}}}^{P_3}$

Preprocessing: $\mathcal{S}_{\Pi_{\mathsf{sh4}}}^{P_3}$ emulates $\mathcal{F}_{\mathsf{setup4}}$ and gives the keys $(k_{03}, k_{13}, k_{23}, k_{013}, k_{023}, k_{123}$ and $k_\mathcal{P})$ to $\mathcal{A}$. The values that are commonly held with $\mathcal{A}$ are sampled using the respective keys, while others are sampled randomly.

Online:

– If dealer is P_3, $\mathcal{S}_{\Pi_{\mathsf{sh4}}}^{P_3}$ receives $\beta_\mathsf{v} + \gamma_\mathsf{v}$ from $\mathcal{A}$ on behalf of P_0. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_i}$ where P_3 acts as one of the sender with P_1, P_2 as the receivers, separately.

– If dealer is P_0 or P_1 or P_2, steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_l}$ where P_3 acts as the server outside the computation.

Figure 47: Simulator $\mathcal{S}_{\Pi_{\mathsf{sh4}}}^{P_3}$ for corrupt P_3

D.2.3 Multiplication Protocol

The case for corrupt P_0 is given in Fig. 48.

Simulator $\mathcal{S}_{\Pi_{\mathsf{mult4}}}^{P_0}$

Preprocessing:

– $\mathcal{S}_{\Pi_{\mathsf{mult4}}}^{P_0}$ samples $[\alpha_z]_1, [\alpha_z]_2, [\Gamma_{xy}]_1$ using the respective keys with $\mathcal{A}$. $\mathcal{S}_{\Pi_{\mathsf{mult4}}}^{P_0}$ samples γ_z, ψ, r randomly on behalf of the respective honest parties, and computes $[\Gamma_{xy}]_2$ honestly.

– Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_i}$ where P_0 acts as one of the sender for sending $[\Gamma_{xy}]_2$.

– $\mathcal{S}_{\Pi_{\mathsf{mult4}}}^{P_0}$ computes $[\chi]_1, [\chi]_2$ honestly. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_k}$ where P_0 acts as the receiver for $[\chi]_1, [\chi]_2$.

Online:

– $\mathcal{S}_{\Pi_{\mathsf{mult4}}}^{P_0}$ computes $[\beta_z^\star]_1, [\beta_z^\star]_2$ honestly. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_j}$ where P_0 acts as one of the sender for sending $[\beta_z^\star]_1, [\beta_z^\star]_2$.

– $\mathcal{S}_{\Pi_{\mathsf{mult4}}}^{P_0}$ computes $\beta_z + \gamma_z$. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_k}$ where P_0 acts as the receiver for receiving $\beta_z + \gamma_z$.

Figure 48: Simulator $\mathcal{S}_{\Pi_{\mathsf{mult4}}}^{P_0}$ for corrupt P_0

The case for corrupt P_1 is given in Fig. 49. The case for a corrupt P_2 is similar.

Simulator $\mathcal{S}_{\Pi_{\mathsf{mult4}}}^{P_1}$

Preprocessing:

– $\mathcal{S}_{\Pi_{\mathsf{mult4}}}^{P_1}$ samples $[\alpha_z]_1, \gamma_z, \psi, r, [\Gamma_{xy}]_1$ using the respective keys with $\mathcal{A}$. $\mathcal{S}_{\Pi_{\mathsf{mult4}}}^{P_1}$ samples $[\alpha_z]_2$ randomly on behalf of the respective honest parties.

– Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_l}$ where P_1 acts as the server outside the computation while communicating $[\Gamma_{xy}]_2$.

– $\mathcal{S}_{\Pi_{\mathsf{mult4}}}^{P_1}$ computes $[\chi]_1$. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_i}$ where P_1 acts as one of the sender for $[\chi]_1$.

– Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_l}$ where P_1 acts as the server outside the computation while communicating $[\chi]_2$.

Online:

– $\mathcal{S}_{\Pi_{\mathsf{mult4}}}^{P_1}$ computes $[\beta_z^\star]_1, [\beta_z^\star]_2$. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_i}$ and $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_k}$, where P_1 acts as one of the sender for sending $[\beta_z^\star]_1$, and P_1 acts as the receiver for receiving $[\beta_z^\star]_2$, respectively.

– $\mathcal{S}_{\Pi_{\mathsf{mult4}}}^{P_1}$ computes $\beta_z + \gamma_z$. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_i}$ where P_1 acts as one of the sender for sending $\beta_z + \gamma_z$.

Figure 49: Simulator $\mathcal{S}_{\Pi_{\mathsf{mult4}}}^{P_1}$ for corrupt P_1

The case for corrupt P_3 is given in Fig. 50.

Simulator $\mathcal{S}_{\Pi_{\mathsf{mult4}}}^{P_3}$

Preprocessing:

- $\mathcal{S}_{\Pi_{\mathsf{mult4}}}^{P_3}$ samples $[\alpha_z]_1, [\alpha_z]_2, \gamma_z, \psi, r, [\Gamma_{xy}]_1$ using the respective keys with $\mathcal{A}$. $\mathcal{S}_{\Pi_{\mathsf{mult4}}}^{P_3}$ computes $[\Gamma_{xy}]_2$ honestly.
- Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_j}$ where P_3 acts as one of the sender for sending $[\Gamma_{xy}]_2$.
- $\mathcal{S}_{\Pi_{\mathsf{mult4}}}^{P_3}$ computes $[\chi]_1, [\chi]_2$. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_j}$ where P_3 acts as one of the sender for sending $[\chi]_1$ and $[\chi]_2$.

Online:

- Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_l}$ where P_3 acts as the server outside the computation involving $[\beta_z^\star]_1, [\beta_z^\star]_2$ and $\beta_z + \gamma_z$.

Figure 50: Simulator $\mathcal{S}_{\Pi_{\mathsf{mult4}}}^{P_3}$ for corrupt P_3

D.2.4 Reconstruction Protocol

The case for corrupt P_0 is given in Fig. 51. The cases for corrupt P_1, P_2, P_3 are similar.

Simulator $\mathcal{S}_{\Pi_{\mathsf{rec4}}}^{P_0}$

- $\mathcal{S}_{\Pi_{\mathsf{rec4}}}^{P_0}$ sends γ_v to $\mathcal{A}$ on behalf of P_1, P_2, and $\mathsf{H}(\gamma_v)$ on behalf of P_3, respectively.
- $\mathcal{S}_{\Pi_{\mathsf{rec4}}}^{P_0}$ receives $\mathsf{H}([\alpha_v]_1), \mathsf{H}([\alpha_v]_2), \beta_v + \gamma_v$ from $\mathcal{A}$ on behalf of P_2, P_1, P_3, respectively.

Figure 51: Simulator $\mathcal{S}_{\Pi_{\mathsf{rec4}}}^{P_0}$ for corrupt P_0

D.2.5 Joint Sharing Protocol

The case for corrupt P_0 is given in Fig. 52.

Simulator $\mathcal{S}_{\Pi_{\mathsf{jsh4}}}^{P_0}$

Preprocessing:

- $\mathcal{S}_{\Pi_{\mathsf{jsh4}}}^{P_0}$ has knowledge of α_v and γ_v, which it obtains while emulating $\mathcal{F}_{\mathsf{setup4}}$. The common values shared with the $\mathcal{A}$ are sampled using the appropriate shared keys, while other values are sampled at random.

Online:

- If dealers are (P_0, P_1): $\mathcal{S}_{\Pi_{\mathsf{jsh4}}}^{P_0}$ computes β_v using v. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_j}$ where P_0 acts as one of the sender for β_v.
- If dealers are (P_0, P_2) or (P_0, P_3): Analogous to the above case.
- If dealers are (P_1, P_2): $\mathcal{S}_{\Pi_{\mathsf{jsh4}}}^{P_0}$ sets $v = 0$ and $\beta_v = [\alpha_v]_1 + [\alpha_v]_2$. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_k}$ where P_0 acts as the receiver for $\beta_v + \gamma_v$.
- If dealers are (P_3, P_1): $\mathcal{S}_{\Pi_{\mathsf{jsh4}}}^{P_0}$ sets $v = 0$ and $\beta_v = [\alpha_v]_1 + [\alpha_v]_2$.

Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_l}$ where P_0 acts as the server outside the computation for β_v, and according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_k}$ where P_0 acts as the receiver for $\beta_v + \gamma_v$.

- If dealers are (P_3, P_2): Analogous to the above case.

Figure 52: Simulator $\mathcal{S}_{\Pi_{\mathsf{jsh4}}}^{P_0}$ for corrupt P_0

The case for corrupt P_1 is given in Fig. 53. The case for corrupt P_2 is similar.

Simulator $\mathcal{S}_{\Pi_{\mathsf{jsh4}}}^{P_1}$

Preprocessing:

- $\mathcal{S}_{\Pi_{\mathsf{jsh4}}}^{P_1}$ has knowledge of α-values and γ corresponding to v which it obtains while emulating $\mathcal{F}_{\mathsf{setup4}}$. The common values shared with the $\mathcal{A}$ are sampled using the appropriate shared keys, while other values are sampled at random.

Online:

- If dealers are (P_0, P_1): $\mathcal{S}_{\Pi_{\mathsf{jsh4}}}^{P_1}$ computes β_v using v. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_i}$ where P_1 acts as one of the sender for β_v.
- If dealers are (P_1, P_2): Analogous to the previous case, except that now $\beta_v + \gamma_v$ is sent instead of β_v.
- If dealers are (P_3, P_1): $\mathcal{S}_{\Pi_{\mathsf{jsh4}}}^{P_1}$ computes β_v and $\beta_v + \gamma_v$. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_i}$ where P_1 acts as one of the sender for $\beta_v, \beta_v + \gamma_v$.
- If dealers are (P_0, P_2) or (P_0, P_3) or (P_3, P_2): $\mathcal{S}_{\Pi_{\mathsf{jsh4}}}^{P_1}$ sets $v = 0$ and $\beta_v = [\alpha_v]_1 + [\alpha_v]_2$. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_k}$ where P_1 acts as the receiver for β_v.

Figure 53: Simulator $\mathcal{S}_{\Pi_{\mathsf{jsh4}}}^{P_1}$ for corrupt P_1

The case for corrupt P_3 is given in Fig. 54.

Simulator $\mathcal{S}_{\Pi_{\mathsf{jsh4}}}^{P_3}$

Preprocessing:

- $\mathcal{S}_{\Pi_{\mathsf{jsh4}}}^{P_3}$ has knowledge of α-values and γ corresponding to v which it obtains while emulating $\mathcal{F}_{\mathsf{setup4}}$. The common values shared with the $\mathcal{A}$ are sampled using the appropriate shared keys, while other values are sampled at random.

Online:

- If dealers are (P_1, P_2): $\mathcal{S}_{\Pi_{\mathsf{jsh4}}}^{P_3}$ sets $v = 0$. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_l}$ where P_3 acts as the server outside the computation for $\beta_v + \gamma_v$.
- If dealers are (P_0, P_1) or (P_0, P_2): Analogous to the above case.
- If dealers are (P_0, P_3): $\mathcal{S}_{\Pi_{\mathsf{jsh4}}}^{P_3}$ computes β_v using v. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_i}$ where P_3 acts as one of the sender for sending β_v.
- If dealers are (P_3, P_1): $\mathcal{S}_{\Pi_{\mathsf{jsh4}}}^{P_3}$ computes β_v and $\beta_v + \gamma_v$. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}_{\Pi_{\mathsf{jmp4}}}^{P_j}$ where P_3 acts as one of the sender for sending $\beta_v, \beta_v + \gamma_v$.

– If dealers are (P_3, P_2): Analogous to the above case.

Figure 54: Simulator $\mathcal{S}^{P_3}_{\Pi_{jsh4}}$ for corrupt P_3

D.2.6 Dot Product Protocol

The case for corrupt P_0 is given in Fig. 55.

Simulator $\mathcal{S}^{P_0}_{\Pi_{dotp4}}$

Preprocessing:

– $\mathcal{S}^{P_0}_{\Pi_{dotp4}}$ samples $[\alpha_z]_1, [\alpha_z]_2, [\Gamma_{\vec{x}\odot\vec{y}}]_1$ using the respective keys with $\mathcal{A}$. $\mathcal{S}^{P_0}_{\Pi_{dotp4}}$ samples γ_z, ψ, r randomly on behalf of the respective honest parties, and computes $[\Gamma_{\vec{x}\odot\vec{y}}]_2$ honestly.

– Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}^{P_i}_{\Pi_{jmp4}}$ where P_0 acts as one of the sender for $[\Gamma_{\vec{x}\odot\vec{y}}]_2$.

– $\mathcal{S}^{P_0}_{\Pi_{dotp4}}$ computes χ_1, χ_2 honestly. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}^{P_k}_{\Pi_{jmp4}}$ where P_0 acts as the receiver for χ_1 and χ_2.

Online:

– $\mathcal{S}^{P_0}_{\Pi_{dotp4}}$ computes $[\beta_z^\star]_1, [\beta_z^\star]_2$ honestly. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}^{P_j}_{\Pi_{jmp4}}$ where P_0 acts as one of the sender for $[\beta_z^\star]_1, [\beta_z^\star]_2$.

– $\mathcal{S}^{P_0}_{\Pi_{dotp4}}$ computes $\beta_z + \gamma_z$. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}^{P_k}_{\Pi_{jmp4}}$ where P_0 acts as the receiver for $\beta_z + \gamma_z$.

Figure 55: Simulator $\mathcal{S}^{P_0}_{\Pi_{dotp4}}$ for corrupt P_0

The case for corrupt P_1 is given in Fig. 56. The case for corrupt P_2 is similar.

Simulator $\mathcal{S}^{P_1}_{\Pi_{dotp4}}$

Preprocessing:

– $\mathcal{S}^{P_1}_{\Pi_{dotp4}}$ samples $[\alpha_z]_1, \gamma_z, \psi, r, [\Gamma_{\vec{x}\odot\vec{y}}]_1$ using the respective keys with $\mathcal{A}$. $\mathcal{S}^{P_1}_{\Pi_{dotp4}}$ samples $[\alpha_z]_2$ randomly on behalf of the respective honest parties.

– Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}^{P_l}_{\Pi_{jmp4}}$ where P_1 acts as the server outside the computation for $[\Gamma_{\vec{x}\odot\vec{y}}]_2$.

– $\mathcal{S}^{P_1}_{\Pi_{dotp4}}$ computes χ_1. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}^{P_i}_{\Pi_{jmp4}}$ where P_1 acts as one of the sender for χ_1.

– Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}^{P_l}_{\Pi_{jmp4}}$ where P_1 acts as the server outside the computation for χ_2.

Online:

– $\mathcal{S}^{P_1}_{\Pi_{dotp4}}$ computes $[\beta_z^\star]_1, [\beta_z^\star]_2$. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}^{P_i}_{\Pi_{jmp4}}$ and $\mathcal{S}^{P_k}_{\Pi_{jmp4}}$, where P_1 acts as one of the sender for $[\beta_z^\star]_1$, and P_1 acts as the receiver for $[\beta_z^\star]_2$.

– $\mathcal{S}^{P_1}_{\Pi_{dotp4}}$ computes $\beta_z + \gamma_z$. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}^{P_i}_{\Pi_{jmp4}}$ where P_1 acts as one of the sender for $\beta_z + \gamma_z$.

Figure 56: Simulator $\mathcal{S}^{P_1}_{\Pi_{dotp4}}$ for corrupt P_1

The case for corrupt P_3 is given in Fig. 57.

Simulator $\mathcal{S}^{P_3}_{\Pi_{dotp4}}$

Preprocessing:

– $\mathcal{S}^{P_3}_{\Pi_{dotp4}}$ samples $[\alpha_z]_1, [\alpha_z]_2, \gamma_z, \psi, r, [\Gamma_{\vec{x}\odot\vec{y}}]_1$ using the respective keys with $\mathcal{A}$. $\mathcal{S}^{P_3}_{\Pi_{dotp4}}$ computes $[\Gamma_{\vec{x}\odot\vec{y}}]$ honestly.

– Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}^{P_j}_{\Pi_{jmp4}}$ where P_3 acts as one of the sender for $[\Gamma_{\vec{x}\odot\vec{y}}]_2$.

– $\mathcal{S}^{P_3}_{\Pi_{dotp4}}$ computes χ_1, χ_2. Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}^{P_j}_{\Pi_{jmp4}}$ where P_3 acts as one of the sender for χ_1, χ_2.

Online:

– Steps corresponding to Π_{jmp4} are simulated according to $\mathcal{S}^{P_l}_{\Pi_{jmp4}}$ where P_3 acts as the server outside the computation for $[\beta_z^\star]_1, [\beta_z^\star]_2, \beta_z + \gamma_z$.

Figure 57: Simulator $\mathcal{S}^{P_3}_{\Pi_{dotp4}}$ for corrupt P_3

D.2.7 Truncation Pair Generation

Here we give the simulation steps for Π_{trgen4}. The case for corrupt P_0 is given in Fig. 58. The case for corrupt P_3 is similar.

Simulator $\mathcal{S}^{P_0}_{\Pi_{trgen4}}$

– $\mathcal{S}^{P_0}_{\Pi_{trgen4}}$ samples R_1, R_2 using the respective keys with $\mathcal{A}$.

– Steps corresponding to Π_{jsh4} are simulated according to $\mathcal{S}^{P_0}_{\Pi_{jsh4}}$ (Fig. 52).

Figure 58: Simulator $\mathcal{S}^{P_0}_{\Pi_{trgen4}}$ for corrupt P_0

The case for corrupt P_1 is given in Fig. 59. The case for corrupt P_2 is similar.

Simulator $\mathcal{S}^{P_1}_{\Pi_{trgen4}}$

– $\mathcal{S}^{P_1}_{\Pi_{trgen4}}$ samples R_1 using the respective keys with $\mathcal{A}$, and samples R_2 randomly.

– Steps corresponding to Π_{jsh4} are simulated according to $\mathcal{S}^{P_1}_{\Pi_{jsh4}}$ (Fig. 53).

Figure 59: Simulator $\mathcal{S}^{P_1}_{\Pi_{trgen4}}$ for corrupt P_1