

Asterisk: Super-fast MPC with a Friend

Banashri Karmakar*, Nishat Koti*, Arpita Patra*, Sikhar Patranabis[†], Protik Paul* and Divya Ravi[‡]

* IISc Bangalore

Email: {banashrik,kotis,arpita,protikpaul}@iisc.ac.in

[†] IBM IRL

Email: sikhar.patranabis@ibm.com

[‡] Aarhus University

Email: divya@cs.au.dk

Abstract—Secure multiparty computation (MPC) enables privacy-preserving collaborative computation over sensitive data held by multiple mutually distrusting parties. Unfortunately, in the most natural setting where a majority of the parties are maliciously corrupt (also called the *dishonest majority* setting), traditional MPC protocols incur high overheads and offer weaker security guarantees than are desirable for practical applications. In this paper, we explore the possibility of circumventing these drawbacks and achieving practically efficient dishonest majority MPC protocols with strong security guarantees by assuming an additional semi-honest, non-colluding helper party HP¹. We believe that this is a more realistic alternative to assuming an honest majority, since many real-world applications of MPC involving potentially large numbers of parties (such as dark pools) are typically enabled by a central governing entity that can be modeled as the HP.

In the above model, we are the first to design, implement and benchmark a practically-efficient and general multi-party framework, **Asterisk**. Our framework requires invoking HP only a constant number of times, achieves the strong security guarantee of *fairness* (either all parties learn the output or none do), scales to hundreds of parties, outperforms all existing dishonest majority MPC protocols, and is, in fact, competitive with state-of-the-art honest majority MPC protocols. Our experiments show that **Asterisk** achieves $288 - 228 \times$ speedup in preprocessing as compared to the best dishonest majority MPC protocol. With respect to online time, **Asterisk** supports 100-party evaluation of a circuit with 10^6 multiplication gates in approximately 20 seconds. We also implement and benchmark practically efficient and highly scalable instance of dark pools using **Asterisk**. The run times showcase the effectiveness of **Asterisk** in enabling efficient realizations of real-world privacy-preserving applications with strong security guarantees.

1. Introduction

In today’s digital landscape, vast amounts of user data are being collected and analyzed. This motivates the need

1. We refer to this entity as a friend in the title.

for privacy-preserving computation that enable computing on sensitive data without compromising data privacy. Secure multiparty computation (MPC) [42] is one such technology that facilitates computation over sensitive data without disclosing any underlying information. Informally, an MPC protocol allows a set of n distrusting parties with private inputs to jointly evaluate a function on these inputs, while ensuring that a (centralized) adversary corrupting a subset of at most $t < n$ parties learns nothing beyond the function output.

We motivate the practical applicability of MPC with the example use-case of privacy-preserving dark pools. Dark pools are private exchanges that allow institutional investors to trade large volumes of financial instruments (such as stocks, bonds, shares, etc.) away from the prying eyes of the public. Unlike public exchanges, trade requests submitted to a dark pool remain hidden until they are matched. At a high-level, dark pools operate as follows. Investors submit trade requests to a central dark pool operator who is responsible for matching buy orders to sell orders, and publishing the list of matched (or satisfied) orders. Submitted requests that are unmatched remain hidden with the dark pool operator until they get satisfied. This solution necessitates that the centralized dark pool operator (a) respects the privacy of the client’s trade interests, and (b) executes the identification of matches among the buy and sell orders honestly. MPC, on the other hand, obviates the need for such trust assumptions. To guarantee privacy and ensure that the trade interest is not disclosed in the clear to the dark pool operator, one can move to the distributed setting. Here, the clients run an MPC protocol among themselves (with their trade requests as the private input) to determine the matching orders². The protocol would identify the matching orders while guaranteeing that no corrupt subset of $t < n$ clients can reconstruct any of the honest client’s inputs (and, in fact, does not learn anything beyond the matching orders).

Honest vs dishonest majority. While MPC appears to be an ideal solution, one needs to consider more fine-grained

2. Depending on the dark pool algorithm, an alternative approach is to employ an MPC protocol to enable multiple parties to collectively effectuate the role of a dark pool operator [10], [33], [17].

aspects, namely the *corruption threshold* that the MPC protocol can withstand, and the *notion of security* that the protocol achieves. In practical applications such as privacy-preserving dark pools, one would expect a majority of the clients to be corrupt. In the context of MPC, this translates to the adversary corrupting up to $t = (n - 1)$ out of n parties, and is termed as the *dishonest majority setting*. The counterpart to this setting is the *honest majority setting*, where the adversary corrupts only a minority of the parties, i.e., $t < n/2$. Our focus in this paper is on MPC in the dishonest majority setting.

Traditional dishonest majority MPC. Unfortunately, traditional MPC protocols in dishonest majority *necessarily* rely on cryptographic hardness assumptions, which renders them practically inefficient in comparison to their honest majority counterparts, and potentially degrades the performance of *real-time, throughput-sensitive* applications (e.g., dark pools, privacy-preserving machine learning, etc.). Additionally, known impossibility results make it impossible for traditional dishonest majority MPC protocols to provide stronger security notions of fairness or full security/guaranteed output delivery (GOD) in the presence of a malicious³ adversary [15]. Fairness guarantees that corrupt parties do not gain any unfair advantage in learning the output (i.e., either all parties learn the output of the computation or none do), while GOD provides the stronger guarantee that, irrespective of the corrupt parties’ misbehaviour, all parties obtain the output. In the absence of fairness (or GOD), a dishonest majority protocol empowers an adversary to abort the computation at will as well as learn the output. This can be disastrous in applications such as dark pools since this allows the adversary to repeatedly abort the computation if the outcome of matching is not favourable to the adversary (unlike in a fair protocol where the adversary does not gain any unfair advantage in learning the output). These drawbacks tend to render traditional dishonest majority MPC protocols insufficient for many practical applications.

Honest majority MPC. MPC protocols in honest majority setting [6], [7], [24], [34] avoid the above drawbacks and yield practically efficient solutions with strong security guarantees. However, the assumption of an honest majority is, in general, stronger than its counterpart dishonest majority and may not hold in many application scenarios.

MPC with a friend. The above discussion motivates the need for *alternative MPC models* that could bypass the inherent drawbacks of traditional dishonest majority protocols, while also avoiding the strong corruption threshold assumptions that characterize honest-majority solutions. In this paper, we consider such a model that (we believe) realistically captures corruption behaviors in certain real-world applications. We dub this model as *MPC with a friend*, where the parties are aided by the presence of an additional helper party HP (which is not trusted but is semi-honest, i.e., does not deviate arbitrarily from the protocol). As we discuss below, in certain scenarios, as compared to assuming an

honest majority within a (potentially large) set of parties, it seems more practical to identify a single (semi-honest) party, while allowing a majority of the remaining parties to be (maliciously) corrupt. For example, in the case of dark pools, the Securities and Exchange Commission (SEC) [10] is a trusted entity that governs the integrity of the dark pool operations (including that the right algorithm is used to perform matching of buy and sell orders). The SEC can instead be modelled as semi-honest (thereby reducing the points of trust in the overall system) and yet non-colluding with the clients. Thus, it can enact the role of HP, whereas the clients can be maliciously corrupt under a dishonest majority assumption. Analogous to dark pools, another motivating application is inventory matching [39], which is a standard mechanism for trading financial stocks by which buyers and sellers can be paired so that the related stocks can be traded without adversely impacting the market price for either client. In the financial world, banks often undertake the task of finding such matches between their clients, thus a bank can enact the role of the HP for its clients. In general, the trusted, central entity, that is responsible for governing the execution of an application, can be modelled as the HP. *Related work.* We note that several works have considered using a “trusted” HP to achieve MPC protocols with stronger security guarantees [21], [22], [27]. However, these works typically assumed that the HP was a small and fully trusted entity (e.g., tamper-proof hardware tokens [9], [29], [14], [18], [13], [12], [25], [4]), sometimes doing minimal work. On the other hand, modeling the HP as a semi-honest entity was introduced in a more recent work [28]. As shown in [28], this model circumvents the impossibility of [15], even when the HP is semi-honest, as long as it does not collude with the remaining parties (with a maliciously corrupt majority). It is additionally shown in [28] that this non-collusion assumption is somewhat necessary for strong security notions such as fairness/GOD that we want for practical applications (we expand more on this later). Finally, the authors of [28] propose certain concrete MPC protocols in this setting that achieve GOD; however, these are primarily feasibility results based on strong cryptographic machinery and are practically inefficient.

A related MPC model called Assisted MPC (based on a probabilistic semi-honest, non-colluding helper party) was studied in [37]; however, this model differs from the HP-aided MPC model of [28] in two ways: (a) the authors of [37] assume that the semi-honest behavior of the helper party is *probabilistic* (i.e., it could also be malicious with some finite non-zero probability), and (b) it is not studied in [37] if this model circumvents known impossibilities and enables strong security guarantees such as fairness/GOD (the protocols in [37], while practically efficient, still achieve the weaker notion of abort security – same as traditional dishonest majority MPC protocols). This leads to the question: can we bridge this gap between inefficient protocols with strong security guarantees and efficient protocols with weak security guarantees using a non-colluding, semi-honest HP?

Relation to FaF security. We briefly comment on the non-

3. A malicious adversary may deviate arbitrarily from protocol specification.

colluding HP assumption above, and its relation to another notion of security for MPC called friends-and-foes (FaF) security [1], which requires security to hold not only against the adversarial parties (foes), but also against quorums of honest parties (friends). [1] proposed modelling this using a *decentralized* adversary consisting of two different *non-colluding* adversaries—(i) a malicious adversary that corrupts any subset of at most t out of n parties, and (ii) a semi-honest adversary that corrupts any subset of at most h^* out of the remaining $(n - t)$ parties. Further, the FaF model requires security to hold even when the malicious adversary sends its view to the semi-honest adversary. A natural adaptation of this notion to the HP-aided setting would translate to the malicious adversary corrupting up to $(n - 1)$ parties and the semi-honest adversary corrupting the HP (this is clearly weaker than traditional FaF security because the HP is a designated semi-honest party, but stronger than the fully non-colluding HP assumption of [28], since the HP can be given access to the views of the corrupt parties). The impossibility result of [28] does not rule out the possibility of getting stronger security guarantees in this model (their impossibility assumes that the *same centralized* adversary corrupts both the HP and a majority of the remaining parties). Unfortunately, we rule out the possibility of achieving fair MPC protocols in this model (see Section 3), thus establishing that the non-colluding model of HP-aided MPC from [28] is the best we can hope for if we wish to achieve fairness/GOD in dishonest majority. In the rest of the paper, we refer to this setting as the *dishonest majority HP-aided setting*. Concretely, we ask the following: *Can we achieve highly efficient MPC protocols with strong security in the dishonest majority HP-aided setting?*

1.1. Our contributions

We answer the above question in the affirmative. Our contributions are as described below.

Asterisk. We put forth *Asterisk*—a highly efficient MPC protocol in the preprocessing paradigm that allows computing arithmetic as well as Boolean circuits over secret-shared inputs in the dishonest majority HP-aided setting (i.e., in the presence of a semi-honest, non-colluding HP) while achieving fairness. Unlike standard dishonest-majority MPC protocols [19], [31], [5] that only achieve abort security due to known impossibilities for achieving fairness [15], *Asterisk* achieves stronger security (fairness) by leveraging the HP to bypass this impossibility. In particular, *Asterisk* uses the (semi-honest, non-colluding) HP in both the preprocessing and online phases, and, unlike existing dishonest majority protocols, achieves an extremely fast and lightweight (function-dependent) preprocessing phase, as well as a highly efficient online phase, while maintaining privacy against the (semi-honest) HP. *Asterisk* makes overall only *four* calls to the HP, which is reasonably small. We expand more on this below.

Efficient preprocessing. We note that traditional dishonest-majority MPC protocols [19], [31], [32], [5] crucially rely on

cryptographic machinery (such as homomorphic encryption or oblivious transfer) in the preprocessing phase, which leads to greater computational and communication overheads as compared to honest-majority protocols. *Asterisk*, on the other hand, achieves a highly efficient (function-dependent) pre-processing phase that uses no additional cryptographic machinery, and only requires communicating 3 elements per multiplication gate (this is a significant improvement over the $\mathcal{O}(n^2)$ complexity of state-of-the-art dishonest majority protocols, such as [31], [32], [5]).

In fact, *Asterisk* also outperforms Assisted MPC [37] in terms of preprocessing efficiency. Recall that, like *Asterisk*, Assisted MPC also relies on an external non-colluding semi-honest party (however, the helper party is only probabilistically semi-honest in the case of Assisted MPC, and the resulting protocol is only abort secure). However, Assisted MPC incurs a communication overhead of $\mathcal{O}(n)$ per multiplication gate in its preprocessing phase, which is significantly higher than the requirement of communicating 3 elements per multiplication gate in the preprocessing phase of *Asterisk*.

Technically, these efficiency gains stem from careful use of the HP in the preprocessing phase of *Asterisk*, which constitutes a major technical novelty of our work. The (function-dependent) preprocessing phase of *Asterisk* uses the HP to: (a) produce authenticated additive shares of the masks for all wires in the circuit representation of the function to be computed, and (b) generate shares of multiplication triples. In particular, the preprocessing phase *exclusively* involves the HP generating and sending messages to the parties, while the parties themselves only perform local computations without communicating with each other. This minimizes the communication overhead, while also ensuring that the preprocessing phase is *inherently error-free* (thereby avoiding any additional overheads for verification checks). See Section 4 for details.

Efficient online phase. *Asterisk* achieves better online efficiency in the dishonest majority (all but one) protocols [31], [32]. Also, it does not rely on additional primitives such as a commitment scheme, which is required in the prior works. Furthermore, we utilize the HP in the online phase to achieve a stronger security guarantee of fairness.

Table 1 provides a comparison of *Asterisk* with other dishonest majority protocols [31], [32], [37] as well as state-of-the-art honest majority protocols [24], [34] in terms of efficiency and security. It is easy to see that *Asterisk* achieves better efficiency (in both the preprocessing and online phases) as well as stronger security guarantees (fairness as opposed to abort security) as compared to dishonest majority protocols [31], [32], [37]. In fact, *Asterisk* is actually closer to the honest majority protocols [24], [34] in terms of efficiency and security guarantees.

Impossibility of fairness in FaF with HP. We justify our choice of assuming the dishonest majority HP-aided setting (with a semi-honest, *fully* non-colluding HP) for the design of *Asterisk* by proving that this assumption is, in fact, necessary to achieve dishonest majority MPC protocols

Protocol	Model	Security	Preprocessing		Online	
			Comm.	Rounds	Comm.	
Mascot [31]	DM	Abort	$O(n^2)$	2	$4n$	
Overdrive † [32]		Abort	$O(n^2)$	2	$4n$	
Assisted MPC † [37]		Abort	$2n$	2	$4n$	
ATLAS [24] *	HM	Abort	NA	2	$2n$	
MPClan [34]		Fairness	$1.5n$	2	$1.5n$	
Asterisk	DM with HP	Fairness	3	2	$2n$	

‡ For preprocessing, [31] relies on OT and [32] relies on SHE.

† [37] uses an additional helper party in the preprocessing phase.

* The costs reported are for the computational variant. We assume parties have common PRF keys that facilitate the non-interactive generation of common random values between them, similar to Asterisk.

TABLE 1: Comparison of Asterisk with the state-of-the-art dishonest majority (DM) and honest majority (HM) protocols.

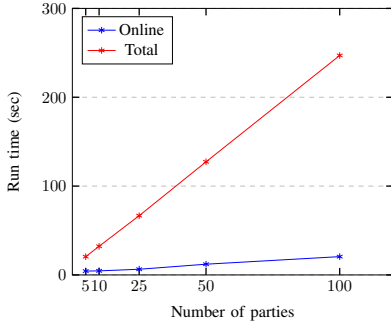


Figure 1: Online and total run time over a LAN network of Asterisk for a circuit of size 10^6 and the depth of the circuit is 10, for varying number of parties.

satisfying fairness. Concretely, we prove that it is impossible to design a fair protocol when considering the model of FaF security with HP. This result, coupled with the impossibility of attaining fairness using a colluding HP [28], effectively rules out the possibility of designing fair protocols unless the HP is *fully* non-colluding (which is precisely the model we use for designing Asterisk). See Section 3 for more details.

Applications and building blocks. We build upon Asterisk to design secure and efficient protocols in the HP-aided setting for privacy-preserving dark pools. Along the way, we design sub-protocols for various granular functionalities, such as equality, comparison, dot product, and shuffle (among others). We believe that these building blocks are of independent interest (with potentially other applications such as privacy-preserving machine learning, anonymous broadcast, etc.) and underline the versatility of Asterisk as a framework for enabling a wide range of privacy-preserving applications in a secure yet efficient manner.

Implementation and benchmarks. We implement and benchmark Asterisk to showcase its efficiency as well as scalability. When evaluating a circuit comprising 10^6 multiplication gates, Asterisk has a response time (online) of around 20 seconds for as many as 100 parties. Fig. 1 summarizes Asterisk’s performance for varying number of parties. Fig. 2 presents a comparison between Asterisk and Assisted MPC [37] (both of which use some kind of external helper party). We highlight that Asterisk’s preprocessing

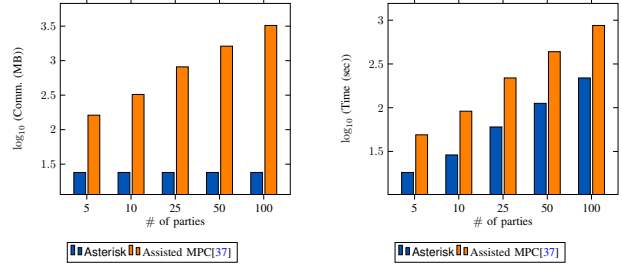


Figure 2: Preprocessing comparison between Assisted MPC [37] and Asterisk for a circuit of size 10^6 (the values are scaled in the logarithmic scale with base 10).

phase has improvements of up to $4\times$ and $134\times$, in terms of run time and communication overheads, respectively, as compared to that of Assisted MPC. We present additional experimental results in Section 5 showcasing that, compared to state-of-the-art dishonest majority MPC protocols [31], [32], Asterisk achieves gains of up to $6\times$ and $4\times$ in total and online communication, respectively, as well as gains of up to $288\times$ in throughput of multiplicative triple-generation in the preprocessing phase.

We also showcase the practicality of Asterisk by benchmarking the application of privacy-preserving dark pools. We consider two dark pool algorithms—volume matching and continuous double auctions (CDA). We observe that in the presence of up to 100 clients, volume matching has a response time of 22 seconds. In the case of CDA, we observe that processing a new buy (or sell) request takes around 10 seconds given a sell (buy) list of size 500.

Comparison with FaF-secure protocols. As noted earlier, the dishonest majority HP-aided model bears some resemblance to the FaF-security model [1], with the difference that in the former model, the HP is a designated semi-honest party (there is no such designated party in FaF-security), and in FaF-security, the semi-honest adversary may be given access to the view of the malicious adversary (this is not allowed in the dishonest majority HP-aided model). While [1] presented some feasibility results for FaF-secure MPC without honest majority, more recent works have proposed concretely efficient FaF-secure MPC protocols [26], [33] that achieve stronger-than-abort security, albeit for specific numbers of parties (specifically, 4 parties in [26] and 5 parties in [33]) and specific corruption thresholds (at most one malicious corruption). It is not immediately obvious how one might extend these schemes for any general number of parties n while retaining comparable efficiency and security guarantees. On the other hand, the efficiency and security guarantees of Asterisk hold for any general number of parties (this is well suited for applications such as dark pools, where the number of parties could be large in practice). We also note that [26] has a significantly more expensive preprocessing phase as compared to Asterisk, and both [26], [33] cannot tolerate a dishonest majority. Finally, we believe that for the application that we consider in this paper, the dishonest majority HP-aided model more aptly captures the expected behavior of corrupt parties, which do not seem to have any particular incentive to share their views with the

semi-honest HP.

Organization. Section 2 introduces notations and formally defines the dishonest majority HP-aided model. Section 3 justifies this model by presenting our impossibility result for fairness in the FaF model with a semi-honest HP. Section 4 describes our main contribution—the Asterisk framework. Section 5 presents an outline of how Asterisk is used to realize the target application (and the building blocks thereof), followed by implementation details and benchmarking results. Finally, Section 6 concludes and discusses some open questions. Certain missing details are deferred to the appendix.

2. Preliminaries

Threat model. We begin with a description of the dishonest majority HP-aided model of MPC that we use throughout the paper. We consider n parties $\mathcal{P} = \{P_1, \dots, P_n\}$, where each party P_i is initialized with input $x_i \in \{0, 1\}^*$ and random coins $r_i \in \{0, 1\}^*$. These parties interact in synchronous rounds. In every round parties can communicate either over a fully connected point-to-point (P2P) network, where we additionally assume all communication to be private and ideally authenticated. Further, we assume that there exists a special party HP, outside \mathcal{P} , called a “helper party” (abbreviated henceforth as HP) such that each party P_i can interact with HP via private and authenticated point-to-point channels. The HP does not typically hold any inputs and also does not obtain any output at the end of the protocol. Depending on whether we allow the HP to keep any state in between its invocations (where an invocation corresponds to all parties interacting with the HP in a single round) or not, we refer to the HP as being ‘stateful’ or ‘stateless’ respectively.

We consider a non-colluding adversary who either corrupts up to $n - 1$ among the n parties maliciously or the HP in a semi-honest manner. We prove the security of our protocols based on the standard real world / ideal world paradigm [8]. Essentially, the security of a protocol is analyzed by comparing what an adversary can do in the real execution of the protocol to what it can do in an ideal execution, that is considered secure by definition (in the presence of an incorruptible trusted party). For a detailed description of the security model of non-colluding, colluding and a related notion of friends-and-foes (FaF), refer to Appendix A.

Notations. We also introduce some notations that we use throughout the paper.

- *Data representation:* We consider secure computation over finite fields \mathbb{F} . \mathbb{F} can be either a binary field \mathbb{F}_2 or a prime field \mathbb{F}_p where p is a k -bit prime. For computation, we consider signed values, that is, where the boolean representation of a value $x \in \mathbb{F}_p$ is in the 2’s complement form. A positive value $x \in [0, \frac{p}{2}]$ has the most significant bit (MSB) as 0 and a negative value $x \in [-\frac{p}{2}, 0)$ has 1 in the MSB.
- $[1, n]$ denotes the set $\{1, \dots, n\}$.

- $(b)^A$ is used to denote arithmetic equivalent of bit b .
- $b = 1(x \stackrel{?}{=} 0)$ denotes $b = 1$ if $x = 0$, otherwise $b = 0$.
- $b = 1(x \stackrel{?}{\leq} 0)$ denotes $b = 1$ if $x \leq 0$, otherwise $b = 0$.
- $\mathcal{F}_{\text{Rand}}$ is a random coin-tossing functionality. It allows parties and HP to sample common random values such that a corrupt party guesses this value with negligible probability.

3. Impossibility of fairness in FaF with HP

In this section, we show that it is impossible to achieve fairness with FaF security in our HP-aided setting. The impossibility result, at a high level, follows similarly to the impossibility of attaining fairness for $n = 3, t = 1, h^* = 1$ with FaF security, as described in [1]. Note that in the standard definition of (t, h^*) -FaF security, any subset of h^* parties can be semi-honestly corrupt. On the other hand, when considering FaF security with HP, it is only a designated entity (namely the HP) which is semi-honest (see Appendix A.1 for the formal definition). In this way, the standard notion of FaF security is stronger than FaF security with HP. Hence, the impossibility of FaF security with HP does not follow directly from the impossibility of [1]. We show how the proof of [1] can be modified to extend to the notion of FaF security with HP.

The main difference in our proof and that of [1], is that in our case, the semi-honest party (HP) does not have any input, while the semi-honest party in the case of [1] does. Further, the proof of [1] relied on this input being unchanged while the simulator invokes the ideal functionality. To make this proof work for our setting, we first modify the functionality to account for the fact that HP does not have an input. Further, instead of relying on the input of the semi-honest party remaining unchanged, we rely on the fact that the honest party’s input to the ideal functionality remains unchanged. We next provide a high-level overview of the proof.

First, we work out the proof for two parties in which one will be maliciously corrupt and in addition there is a semi-honestly corrupt HP. The argument for n parties and one HP follows from player partitioning technique [36]. Consider 2 parties A and B , and a HP. We show that there exists a function F (which takes input from A and B), that cannot have a fair protocol in the considered model of FaF security with HP. To prove the above claim, we proceed via contradiction. That is, we show that if there exists a fair protocol, then there exists a polynomial time algorithm that can invert a one-way permutation. For this, we define the function F to be Swap_κ using a one-way permutation f_κ (κ is the security parameter). To be specific, we use the following functionality

$$\text{Swap}_\kappa((a, y_B), (b, y_A)) = \begin{cases} (a, b) & \text{if } f_\kappa(a) = y_A \text{ and } f_\kappa(b) = y_B \\ 0^\kappa & \text{otherwise} \end{cases}$$

Each party inputs a pair of values $(a, y_B), (b, y_A)$, where a, b is in the domain of f_κ , and y_A, y_B is in the range of f_κ . Swap_κ then outputs (a, b) if $f_\kappa(b) = y_A$ and $f_\kappa(a) = y_B$. Let there exist a fair protocol to compute Swap_κ that comprises r rounds, where $r \leq p(\kappa)$, for some polynomial $p(\cdot)$. We show that there exists a round $i \leq r$ such that either (A, HP) gain an advantage in learning the output than B or it is the case that (B, HP) gain an advantage in learning the output than A . Without loss of generality, consider the case that A is malicious. We then show that this advantage of the pair (A, HP) over (B, HP) cannot be simulated, thereby breaking the fairness property. For this, we show that if this can be simulated, then there exists a polynomial time algorithm which can break the one-wayness property of the one-way permutation.

To showcase that the advantage of (A, HP) over (B, HP) cannot be simulated, consider the strategy of a malicious A . A acts honestly (using the original input it held) until this round i , after which it aborts. A then sends its entire view to HP. Because of the aforementioned claim on the existence of round i , receiving A 's view allows the HP to recover the correct output (corresponding to an all-honest execution) with a significantly higher probability than B . If there exists a simulator for such an (A, HP) pair, then we can construct a polynomial time algorithm M , which can be used to break the one-wayness of the underlying one-way permutation. Elaborately, M is the same algorithm performed by HP to obtain the output at the end of round i when A aborts. Observe that M takes as input the view of HP (which includes the view of corrupt A as well as its input y_B), and generates the output of Swap_κ , which comprises the inverse of a one-way permutation (b) . Hence, using M , we can construct a polynomial-time algorithm to invert a one-way permutation.

The formal theorem appears below and detailed proof in Appendix B.

Theorem 3.1. *There exists an n -party functionality such that no protocol computes it with fairness in the FaF security model with HP.*

4. Efficient n -PC with HP

Here, we describe the MPC protocol of *Asterisk*, which is designed in the preprocessing paradigm. Computation begins by executing a one-time shared key setup phase, where common PRF keys are established between the parties to facilitate the non-interactive generation of common random values among the parties during the protocol execution. Next, the preprocessing phase kicks in, where the necessary preprocessing (input-independent) data is generated. This is followed by the online phase once the input is available.

4.1. Shared key setup

Let $F : \{0, 1\}^\kappa \times \{0, 1\}^\kappa \rightarrow \mathbb{F}$ be a secure pseudo-random function (PRF). Parties rely on a one-time setup [3], [2], [34], [35] to establish the following PRF keys

among different subsets of the parties (including the HP). Specifically, the set of keys established between the parties and HP for our protocol is as follows:

- Every party $P_i \in \mathcal{P}$ holds a common key with HP, denoted by $k_i, i \in [1, n]$.
- \mathcal{P} hold a common key $k_{\mathcal{P}}$, unknown to HP.
- $\mathcal{P} \cup \{\text{HP}\}$ hold a common key k_{all} .

The ideal functionality for the same appears in Fig. 3. Discussion on how to instantiate $\mathcal{F}_{\text{Setup}}$ is deferred to Appendix C.1.

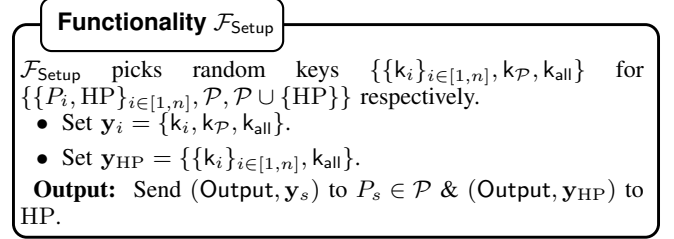


Figure 3: Ideal functionality for shared-key setup

4.2. Sharing semantics

To achieve (malicious) security in the dishonest majority setting, *Asterisk* uses *authenticated* additive secret-sharing of the inputs. Let $v \in \mathbb{F}$ be a secret. We say that v is additively shared among parties in \mathcal{P} if there exist $v_i \in \mathbb{F}$ such that $v = \sum_{i=1}^n v_i$ and $P_i \in \mathcal{P}$ holds v_i . For authentication, we use an information-theoretic (IT) MAC (message authentication code) under a global MAC key $\Delta \in \mathbb{F}$. The authentication tag t for a value v is defined as $t = \Delta \cdot v$, where the MAC key Δ and the tag t are additively shared among the parties in \mathcal{P} . The following are the various sharing semantics used by *Asterisk*.

1. $[\cdot]$ -sharing: A value $v \in \mathbb{F}$ is said to be $[\cdot]$ -shared (additively shared) among parties in \mathcal{P} if the i th party P_i holds $[v]_i \in \mathbb{F}$ such that $v = \sum_{i=1}^n [v]_i$.
2. $\langle \cdot \rangle$ -sharing: A value $v \in \mathbb{F}$ is said to be $\langle \cdot \rangle$ -shared among parties in \mathcal{P} if, the i th party P_i holds $\langle v \rangle_i = \{[v]_i, [t_v]_i\}$, where $t_v = \Delta \cdot v$ is the tag of v , as discussed above. Along with the share of value v and the tag t_v , parties also hold an additive sharing of the MAC key Δ , i.e., P_i holds $[\Delta]_i$.
3. $\llbracket \cdot \rrbracket$ -sharing: A value $v \in \mathbb{F}$ is said to be $\llbracket \cdot \rrbracket$ -shared if
 - there exists a mask $\delta_v \in \mathbb{F}$ that is $\langle \cdot \rangle$ -shared among the parties in \mathcal{P} , and
 - there exists a masked value $m_v = v + \delta_v$ that is held by each $P_i \in \mathcal{P}$.

We denote P_i 's $\llbracket \cdot \rrbracket$ -share of v as $\llbracket v \rrbracket_i = \{m_v, \langle \delta_v \rangle_i\}$.

Note that all these sharing schemes are linear, i.e., given shares of values a_1, \dots, a_m and public constants c_1, \dots, c_m , shares of $\sum_{i=1}^m c_i a_i$ can be computed non-interactively for an integer m . Further, Boolean sharing of a secret bit is

analogous to arithmetic sharing as described above, with the difference that addition/subtraction operations are replaced by XOR (\oplus). We use the superscript **B** to denote Boolean sharing of a bit. The Boolean world is analogous to the arithmetic world where addition/subtraction operations are replaced with XORs, and multiplication is replaced with AND (\wedge).

4.3. Design of Asterisk

We elaborate on the preprocessing and online phase of Asterisk in the preprocessing model.

Preprocessing phase. In this phase, parties obtain $\langle \cdot \rangle$ -shares of the masks associated with all the wires in the circuit representing the function to be computed, while HP gets them all in clear. Specifically, the HP begins by sampling the MAC key Δ and generates its $[\cdot]$ -shares. For an input wire with value v where P_d is the input provider, parties generate $\langle \cdot \rangle$ -shares of a random mask δ_v such that δ_v is known to P_d and HP. For wires that are the output of an addition gate, HP and all parties locally add the $\langle \cdot \rangle$ -shares of masks of the input wires of the addition gate to obtain the $\langle \cdot \rangle$ -shares of the mask for the output wire. With respect to a multiplication gate, let x, y be the inputs to this gate and z be the output. Parties generate $\langle \cdot \rangle$ -shares of a random mask δ_z such that it is known to the HP. Moreover, HP computes $\delta_{xy} = \delta_x \cdot \delta_y$, where δ_x, δ_y are the masks corresponding to x and y , respectively. Then it generates $\langle \cdot \rangle$ -share of δ_{xy} . Looking ahead, δ_{xy} will be used to generate the masked value for the output of the multiplication gate. Finally, for an output gate where v is the output, HP generates and stores the mask δ_v . During the online phase, it will send δ_v to all the parties to facilitate output reconstruction. The preprocessing phase is described in Fig. 7 which uses several sub-protocols for generating the $\langle \cdot \rangle$ -shares of different values. We start with the latter.

The preprocessing phase involves generation of $\langle \cdot \rangle$ -sharing of the following kind of secrets– (i) a random mask that is known by a designated party and the HP, (ii) a random mask that is known only to the HP, (iii) a secret value (not necessarily uniformly random) known to the HP. Note that to generate $\langle \cdot \rangle$ -shares of a value v , the HP can sample $n-1$ shares $[v]_i$ for $i \in \{1, \dots, n-1\}$, followed by computing $[v]_n = v - \sum_{i=1}^{n-1} [v]_i$. Similarly, it can compute the tag t_v , then sample $n-1$ shares for t_v followed by computing the last share for t_v . The HP can then send the i th share of v and t_v to P_i . This would require $2n$ elements of communication. We improve this cost by using the common keys established during the setup phase.

$\Pi_{\langle \cdot \rangle\text{-Sh}}(P_d, \text{Rand})$ (Fig. 4) indicates generating a $\langle \cdot \rangle$ -sharing of a random value v such that the dealer, P_d , and HP know the value in clear. This requires communication of 2 elements.

Protocol $\Pi_{\langle \cdot \rangle\text{-Sh}}(P_d, \text{Rand})$

Input: $\forall i \in [1, n]$, P_i 's input k_i , and HP's input $\{\{k_i\}_{i \in [1, n]}, \Delta\}$.

Output: P_d 's output v and $\langle v \rangle_d$ and $P_i (i \neq d)$ gets $\langle v \rangle_i$.

Protocol:

- HP and P_d sample a random value v using the k_d .
- For all $i \in \{1, \dots, n-1\}$, HP and P_i sample a random share $[v]_i$ and a random share of the tag $[t_v]_i$ using the k_i .
- HP computes $[v]_n = v - \sum_{i \in [n-1]} [v]_i$ and $[t_v]_n = (\Delta \cdot v) - \sum_{i \in [n-1]} [t_v]_i$.
- HP sends $([v]_n, [t_v]_n)$ to P_n .

Figure 4: Generating $\langle \cdot \rangle$ of a random value known to P_d

$\Pi_{\langle \cdot \rangle\text{-Sh}}(\text{HP}, v)$ (Fig. 5) indicates generating a $\langle \cdot \rangle$ -sharing of a random value v sampled by HP. This requires communication of 1 element.

Protocol $\Pi_{\langle \cdot \rangle\text{-Sh}}(\text{HP}, \text{Rand})$

Input: $\forall i \in [1, n]$, P_i 's input k_i , and HP's input $\{\{k_i\}_{i \in [1, n]}, \Delta\}$.

Output: P_i gets $\langle v \rangle_i$ for all $i \in [1, n]$.

Protocol:

- For all $i \in \{1, \dots, n-1\}$, HP and P_i sample a random share $[v]_i$ and a random share of the tag $[t_v]_i$ using k_i .
- HP and P_n sample a random value $[v]_n$ using k_n .
- HP computes $v = \sum_{i=1}^n [v]_i$ and $[t_v]_n = (\Delta \cdot v) - \sum_{i \in [n-1]} [t_v]_i$.
- HP sends $[t_v]_n$ to P_n .

Figure 5: Generating $\langle \cdot \rangle$ of a random value known only to HP

$\Pi_{\langle \cdot \rangle\text{-Sh}}(\text{HP}, v)$ (Fig. 6) indicates generating a $\langle \cdot \rangle$ -sharing of v held by HP. This requires communication of 2 elements.

Protocol $\Pi_{\langle \cdot \rangle\text{-Sh}}(\text{HP}, v)$

Input: $\forall i \in [1, n]$, P_i 's input k_i , and HP's input $\{\{k_i\}_{i \in [1, n]}, \Delta, v\}$.

Output: P_i gets $\langle v \rangle_i$ for all $i \in [1, n]$.

Protocol:

- For all $i \in \{1, \dots, n-1\}$, HP and P_i sample a random share $[v]_i$ and a random share of the tag $[t_v]_i$ using k_i .
- HP computes $[v]_n = v - \sum_{i \in [n-1]} [v]_i$ and $[t_v]_n = (\Delta \cdot v) - \sum_{i \in [n-1]} [t_v]_i$.
- HP sends $([v]_n, [t_v]_n)$ to P_n .

Figure 6: Generating $\langle \cdot \rangle$ of a value v known to HP

We emphasize that the complete preprocessing for a circuit C can be performed at once. Only one round of communication is required from HP to P_n such that all the parties obtain the preprocessing data for the circuit C . For the i th party, preproc_i be the preprocessing data. preproc_i consists of $\langle \cdot \rangle$ -sharing of all wires of C , additionally, for multiplication gates, it contains $\langle \cdot \rangle$ -sharing of the product of the masks of the input wires for the corresponding gate. HP's $\text{preproc}_{\text{HP}}$ only contains the mask of the output wires,

which is used for the output reconstruction. We describe the preprocessing phase of our protocol in Fig. 7.

Protocol Π_{Prep}

Input: For all $i \in [1, n]$, P_i has input the PRF key k_i and the circuit C . HP has $\{k_i\}_{i \in [1, n]}$ and C as input.

Output: P_i 's output preproc_i and HP's output $\text{preproc}_{\text{HP}}$.

Protocol:

- *KeyGen:* For all $i \in \{1, \dots, n\}$, P_i and HP sample $[\Delta]_i$ using their common key k_i . HP computes $\Delta = \sum_{i=1}^n [\Delta]_i$.
- *Input gates:* For a gate with input v and input provider P_d , all parties and HP execute $\Pi_{(\cdot)\text{-Sh}}(P_d, \text{Rand})$ to enable generation of mask δ_v for v in clear among HP and P_d , and generation of its $\langle \cdot \rangle$ -shares among the parties in \mathcal{P} .
- *Addition gates:* For a gate with input x and y , HP and P_i locally compute $\langle \delta_{x+y} \rangle_i = \langle \delta_x \rangle_i + \langle \delta_y \rangle_i$, where $\langle \delta_x \rangle_i + \langle \delta_y \rangle_i = ([\delta_x]_i + [\delta_y]_i, [t_x]_i + [t_y]_i)$.
- *Multiplication gates:* For a gate with input x and y , all parties and HP execute $\Pi_{(\cdot)\text{-Sh}}(\text{HP}, \text{Rand})$ to generate $\langle \cdot \rangle$ -shares of the mask δ_z for the output wire $z = x \cdot y$, where δ_z is known in clear to HP. HP computes $\delta_{xy} = \delta_x \cdot \delta_y$. All parties and HP execute $\Pi_{(\cdot)\text{-Sh}}(\text{HP}, \delta_{xy})$ to generate $\langle \delta_{xy} \rangle$.

Figure 7: Preprocessing phase of Asterisk for a circuit C

Note that the preprocessing phase exclusively involves the (semi-honest) HP generating and sending certain values to the parties. All other parties' tasks are limited to local computation, and they are not involved in the communication. Hence, *the preprocessing phase is inherently error-free*, and we do not need any additional verification checks.

Online phase. Given $\{\{\text{preproc}_i\}_{i \in [1, n]}\}$ and $\text{preproc}_{\text{HP}}$ (generated in the preprocessing phase), once the input is available, the online phase involves generating the masked values for all the wires in the circuit, as per the $[\cdot]$ -sharing semantics. Recall that the masked value m_v for a $[\cdot]$ -shared value v is defined as $m_v = v + \delta_v$ where $\langle \delta_v \rangle$ is generated in the preprocessing phase. At a high level, to generate the masked value m_v , each party is required to perform local computations to generate its $[\cdot]$ -share of m_v . The parties are then expected to send their $[\cdot]$ -share of m_v consistently to every other party. Each party uses the received $[\cdot]$ -shares of m_v to reconstruct it. Note that a malicious party may send inconsistent $[\cdot]$ -shares to the other parties. To ensure consistency of the communicated messages, we resort to performing MAC verification before the output reconstruction. To reconstruct the masked value, m_v , we use P_{King} based approach [20], where parties in \mathcal{P} send their $[\cdot]$ -shares of a value to the P_{King} (which is a designated party among $\{P_1, \dots, P_n\}$), followed by P_{King} sending messages to all the parties. In practice, one would balance the workload by choosing the players as P_{King} suitably. While the formal details of the online phase for evaluating the circuit representing the function to be computed appear in Fig. 11, we next give an overview of the same.

Input gates. For an input gate g_{in} , let P_d be the designated party, referred to as the dealer, providing the input v . Observe that, from the preprocessing phase, P_d holds a random

value δ_v as a mask for v , while the i -th party P_i holds $\langle \delta_v \rangle_i$. All parties, excluding the HP, sample a common random value r using their common PRF key $k_{\mathcal{P}}$. P_d computes $q_v = v + \delta_v + r$ and sends q_v to HP. HP sends q_v to all the parties. Each party locally computes $m_v = q_v - r$, thereby generating its $[\cdot]$ -share of v as $[\![v]\!]_i = \{m_v, \langle \delta_v \rangle_i\}$. The formal details are provided in Fig. 8.

Protocol $\Pi_{[\cdot]\text{-Sh}}$

Input:

- Every input provider/dealer P_d , has input v and mask δ_v .
- For all $i \in [1, n]$, P_i has the input gate g_{in} , $\langle \delta_v \rangle_i$ generated in Fig. 7 for g_{in} and common key $k_{\mathcal{P}}$ generated in the setup phase.

Output: For all $i \in [1, n]$, P_i 's output $[\![v]\!]_i = (m_v, \langle \delta_v \rangle_i)$ where $m_v = v + \delta_v$.

Protocol:

- All parties, excluding HP, sample a random value r using their common key $k_{\mathcal{P}}$.
- P_d computes $m_v = v + \delta_v + r$ and sends it to HP.
- HP sends q_v to all the parties.
- For all $i \in [1, n]$, P_i gets $m_v = q_v - r$ and outputs $[\![v]\!]_i = \{m_v, \langle \delta_v \rangle_i\}$.

Figure 8: Online phase of input sharing of Asterisk

Addition gates. For an addition gate g_{add} , the parties locally add the $[\cdot]$ -shares of the input wires to obtain the $[\cdot]$ -sharing for the output wire. Parties complete this step by locally adding the masked values of the input wires of g_{add} and the addition of the masks is already done in preprocessing.

Multiplication gates. For a multiplication gate g_{mul} , let x and y be the inputs, and let z be the output. To generate $[\cdot]$ -shares of z , parties need to generate m_z , and $\langle \cdot \rangle$ -shares of the mask δ_z . Recall that $\langle \cdot \rangle$ -shares of δ_z are generated in the preprocessing phase. With respect to m_z , it can be computed as follows.

$$\begin{aligned} m_z &= z + \delta_z = xy + \delta_z = (m_x - \delta_x)(m_y - \delta_y) + \delta_z \\ &= m_x m_y - m_x \delta_y - m_y \delta_x + \delta_{xy} + \delta_z \end{aligned}$$

Given the $\langle \cdot \rangle$ -shares of $\delta_x, \delta_y, \delta_z$ and δ_{xy} , which are generated in the preprocessing, and relying on the linearity of $\langle \cdot \rangle$ -sharing, parties can locally generate $\langle \cdot \rangle$ -shares of m_z using the above equation. Elaborately, since $\langle \cdot \rangle$ -share of m_z comprises of $[\cdot]$ -shares of m_z and its tag t_{m_z} , parties compute $[m_z] = m_x m_y - m_x [\delta_y] - m_y [\delta_x] + [\delta_{xy}] + [\delta_z]$, and $[t_{m_z}]_i = [\Delta]_i m_x \cdot m_y - m_x \cdot [t_y]_i - m_y \cdot [t_x]_i + [t_{xy}]_i + [t_z]_i$. To reconstruct m_z , all parties to send their share of m_z to the P_{King} . It reconstructs and sends m_z to all the parties in \mathcal{P} . Thus all parties obtain $[\![z]\!]_i$. Note that, a corrupt party may send incorrect shares to P_{King} and a corrupt P_{King} may send inconsistent values to the parties. However, such malicious behavior will get detected in the MAC verification step. The formal details appear in Fig. 9.

Protocol Π_{mult}

Input:

- For all $i \in [1, n]$, P_i has the multiplication gate g_{mul} . It has the $\langle \cdot \rangle$ of the following values generated in Fig. 7 for g_{mul} $\delta_x, \delta_y, \delta_{xy}, \delta_z$.

Output: For all $i \in [1, n]$, P_i 's output $m_z = z + \delta_z$, where $z = x \cdot y$.

Protocol:

- For all $i \in [1, n]$, P_i holds $\llbracket x \rrbracket_i, \llbracket y \rrbracket_i$ and $\langle \delta_{xy} \rangle_i, \langle \delta_z \rangle_i$.
- Each party P_i computes $\llbracket t_{m_z} \rrbracket_i = -m_x \cdot \llbracket t_y \rrbracket_i - m_y \cdot \llbracket t_x \rrbracket_i + \llbracket t_{xy} \rrbracket_i + \llbracket t_z \rrbracket_i + \llbracket \Delta \rrbracket_i \cdot (m_x \cdot m_y)$.
- P_1 computes $\llbracket m_z \rrbracket_1 = m_x \cdot m_y - m_x \cdot \llbracket \delta_y \rrbracket_1 - m_y \cdot \llbracket \delta_x \rrbracket_1 + \llbracket \delta_{xy} \rrbracket_1 + \llbracket \delta_z \rrbracket_1$ and sends it to P_{King} .
- For all $i \in [1, n] \setminus \{1\}$, P_i computes $\llbracket m_z \rrbracket_i = -m_x \cdot \llbracket \delta_y \rrbracket_i - m_y \cdot \llbracket \delta_x \rrbracket_i + \llbracket \delta_{xy} \rrbracket_i + \llbracket \delta_z \rrbracket_i$ and sends it to P_{King} .
- P_{King} reconstructs m_z and sends it to all the parties (excluding HP).
- For all $i \in [1, n]$, P_i outputs $\llbracket z \rrbracket = (m_z, \langle \delta_z \rangle_i)$.

Figure 9: Online phase of multiplication of Asterisk

Verification. To verify the correctness of the computed m_z , it suffices to check if $\llbracket t_{m_z} \rrbracket - m_z \cdot \llbracket \Delta \rrbracket$ is a sharing of 0. To attain an efficient realization of this check, we combine the verification with respect to multiple multiplication gates into a single check, and this check is performed at the end of the computation, before output reconstruction. This allows us to amortize the cost due to this verification check across many multiplication gates. To combine the checks corresponding to m multiplication gates into one check, parties proceed as follows.

- All parties (including the HP) sample a random vector $\rho = (\rho_0, \rho_1, \dots, \rho_m) \leftarrow \mathbb{F}^{m+1}$ using the common coin functionality $\mathcal{F}_{\text{Rand}}$. Recall that the $\mathcal{F}_{\text{Rand}}$ functionality allows parties and HP to sample common random values such that a corrupt party can guess the sampled value with negligible probability. To realize $\mathcal{F}_{\text{Rand}}$, HP samples a key k_{ver} and sends it to all the parties in \mathcal{P} . Parties can evaluate a PRF with key k_{ver} on a common input ctr. Note that, we can not use k_{all} for $\mathcal{F}_{\text{Rand}}$ since an adversarial party knows the key beforehand and can compute ρ . Thus the soundness of the verification check does not hold. Since the key k_{ver} is not known to any party, a malicious party can guess ρ with negligible probability. Therefore, it can cheat and pass the verification with negligible probability.
- All parties (excluding the HP) sample a random sharing of 0, i.e., P_i gets α_i such that $\sum_{i=1}^n \alpha_i = 0$, using their common key $k_{\mathcal{P}}$.

Let $\{m_{z_1}, \dots, m_{z_m}\}$ be the reconstructed values sent by HP to all the parties, and let $\{\llbracket t_{m_{z_1}} \rrbracket_i, \dots, \llbracket t_{m_{z_m}} \rrbracket_i\}$ be the shares of the tags held by party P_i . For $j \in \{1, \dots, m\}$, P_i computes $\llbracket \omega_{z_j} \rrbracket_i = \llbracket t_{m_{z_j}} \rrbracket_i - m_{z_j} \cdot \llbracket \Delta \rrbracket_i$, and sends the following to the HP: $\llbracket \omega_{z_j} \rrbracket = \rho_0 \cdot \alpha_i + \sum_{j \in [m]} \rho_j \cdot \llbracket \omega_{z_j} \rrbracket_i$. Finally, the HP checks if $\sum_{i=1}^n \llbracket \omega_{z_j} \rrbracket_i = 0$, and sends the outcome to all the parties. The details of the batched MAC verification are given in Fig. 10.

Protocol Π_{verify}

Input:

- For every $i \in [1, n]$, P_i has the reconstructed values $\{m_{z_1}, m_{z_2}, \dots, m_{z_m}\}$.

- For every $i \in [1, n]$, P_i has $\llbracket t_{m_{z_j}} \rrbracket_i$ as the MAC share of m_{z_j} and $\llbracket \Delta \rrbracket_i$ as the MAC key share for the i th party.

Output: All parties and HP output 1 if the verification passes, else 0.

Protocol:

- HP samples a new key k_{ver} and send it to all the parties in \mathcal{P} .
- For all $i \in [1, n]$, P_i computes $\llbracket \omega_{z_j} \rrbracket_i = \llbracket t_{m_{z_j}} \rrbracket_i - m_{z_j} \cdot \llbracket \Delta \rrbracket_i$.
- For all $i \in [1, n]$, P_i samples $\rho = (\rho_0, \rho_1, \dots, \rho_m) \leftarrow \mathbb{F}^{m+1}$ using the key k_{ver} .
- All parties (excluding the HP) sample a random sharing of 0, i.e., P_i gets α_i such that $\sum_{i=1}^n \alpha_i = 0$, using the common key $k_{\mathcal{P}}$.
- For all $i \in [1, n]$, P_i computes $\llbracket \omega_{z_j} \rrbracket = \rho_0 \cdot \alpha_i + \sum_{j=1}^m \rho_j \cdot \llbracket \omega_{z_j} \rrbracket_i$ and sends to HP.
- HP checks if $\sum_{i=1}^n \llbracket \omega_{z_j} \rrbracket_i = 0$. If the check fails, HP sends 0 to all, else 1.

Figure 10: Amortized MAC-Verification of Asterisk

Output reconstruction. If the parties pass the above verification, then HP sends the masks for the output wires. Let v be an output, and δ_v be the mask generated in the preprocessing phase. HP sends δ_v to all the parties. Each party, holding m_v computes $v = m_v - \delta_v$ and gets the output.

Remark. Note that throughout our protocol, the HP is invoked a total of 4 times (once in the offline phase and thrice in the online phase). Note that by a “single” invocation of the HP, we refer to a single instance of the HP receiving messages from all the parties, performing some local computations based on these inputs, and sending messages to the parties. In the online phase, the HP receives messages from the input providers and sends the respective values to all the parties. For verification, HP generates k_{ver} , performs the verification, and sends either the output masks or “abort” to all the parties. We note that a small (constant) number of invocations reduces the overall reliance of our protocol on the HP. This is particularly helpful in practical scenarios where the HP could be a busy resource that is engaged across multiple protocol executions simultaneously.

Remark. Note that HP sends the masks of the output wires to all the parties only if the verification is passed. Therefore if a malicious party deviates from the protocol, resulting in abort, the malicious party fails to learn the output of the protocol. Hence, our protocol achieves the fairness security guarantee. Since HP stores the masks of the output wires, this requires the HP to be a stateful machine that stores data in memory. However, we can also adapt our protocol to use a stateless HP instead using known techniques from [23], which we elaborate in the full version.

Protocol Π_{Online}

Input:

- For every input provider/dealer P_d , it has input v .
- For all $i \in [1, n]$, P_i has the circuit C , the preprocessing data preproc_i generated in Fig. 7 and common keys $\{k_i, k_{\mathcal{P}}, k_{\text{all}}\}$ generated in the setup phase.
- HP's input $\text{preproc}_{\text{HP}}$.

Output: For all $i \in [1, n]$, P_i 's output $C(x_1, x_2, \dots, x_n)$.

Protocol:

- **Input gates:** Parties execute $\Pi_{[\cdot]\text{-Sh}}$ (Fig. 8) for every input gates of C .
- **Addition gates:** Let g_{add} be an addition gate, where x and y are the inputs.
 - For all $i \in [1, n]$, P_i holds $[\![x]\!]_i, [\![y]\!]_i$.
 - For all $i \in [1, n]$, P_i locally computes $[\![x + y]\!]_i = [\![x]\!]_i + [\![y]\!]_i$, where $[\![x]\!]_i + [\![y]\!]_i = ((m_x + m_y), (\delta_x)_i + (\delta_y)_i)$.
- **Multiplication gates:** Parties execute Π_{mult} (Fig. 9) for every multiplication gates of C .
- **Amortized verification:** At the end of all the multiplication gates, all parties and HP execute Π_{verify} (Fig. 10). If parties output \perp , then continue, else output \perp .
- **Output reconstruction:** Let g be an output wire and v be an output. HP holds the mask δ_v of v . HP sends δ_v to all the parties. Parties get the output $v = m_v - \delta_v$.

Figure 11: Online phase of Asterisk for circuit C

4.4. Proof of security

Functionality \mathcal{F}_{MPC}

- **Initialize:** On input $(\text{Init}, \mathbb{F})$ from parties and HP, store \mathbb{F} .
- **Input:** On input $(\text{Input}, P_i, \text{id}, x)$ from P_i and $(\text{Input}, P_i, \text{id})$ from all other parties, with a fresh identifier id and $x \in \mathbb{F}$, store (id, x) .
- **Add:** On command $(\text{Add}, \text{id}_1, \text{id}_2, \text{id}_3)$ from all parties (where id_1, id_2 are present in memory), retrieve (id_1, x) , (id_2, y) and store $(\text{id}_3, x + y)$.
- **Multiply:** On command $(\text{Mult}, \text{id}_1, \text{id}_2, \text{id}_3)$ from all parties (where id_1, id_2 are present in memory), retrieve (id_1, x) , (id_2, y) and store $(\text{id}_3, x \cdot y)$.
- **Output:** On input $(\text{Output}, \text{id})$ from all the parties (where id is present in the memory), output id to the adversary. Wait for an input from the adversary; if this is deliver then retrieves (id, y) and sends it to all the parties, otherwise if the adversary sends abort then it outputs \perp to all the parties including the adversary.

Figure 12: Ideal functionality for evaluating f with fairness

Lemma 4.1. *Assuming existence of a PRF, protocol Asterisk designed in the preprocessing (Fig. 7)-online (Fig. 11) realizes \mathcal{F}_{MPC} (Fig. 12) with computational non-colluding security in the presence of a semi-honest HP.*

Below we briefly discuss the proof of Lemma 4.1 and the formal proof is given in Appendix C.2.

Case I: Consider a malicious adversary corrupting all but one party from \mathcal{P} . In the preprocessing, the adversary

performs only local computation thus the preprocessing is error-free. The online phase is similar to the online phase of [5]. Before the output reconstruction, the verification check ensures that the adversary cannot cheat (except with negligible probability) by sending incorrect values during the reconstruction phase. Finally, to obtain the output, HP sends the masks of the output wires only if the verification check is passed. Thus, either all parties get the output or none which ensures fairness.

Case II: Consider the HP to be semi-honest. In the preprocessing phase, it generates the preprocessing data to evaluate the circuit C . In the online phase, all the values sent to HP are padded with a randomly sampled value. Therefore, to a semi-honest HP, the messages it received in the online phase are indistinguishable from randomly sampled messages. Hence, our protocol is secure against the semi-honest HP.

5. Applications and benchmarks

In this section, we discuss the building blocks required to realize practically efficient privacy-preserving dark pools via Asterisk. We then describe a prototype implementation of Asterisk and compare its performance with other MPC protocols [37], [31], [32]. Finally, we implement and benchmark the aforementioned application on top of Asterisk.

5.1. Building blocks

We realize several building blocks, which are essentially adaptations of Asterisk for certain specific (sub-)functionalities such as equality, comparison, dot product, and shuffle (among others). These sub-protocols are then used for the desired application. While these building blocks have been studied in the traditional MPC settings [11], [16], [33], [38], we adapt these building blocks in the HP-aided setting to attain efficient realizations.

Table 2 summarizes the various sub-protocols that we use as building blocks for our applications, as well as the corresponding overheads. Due to the lack of space in the body of the paper, we defer the full details of these building blocks to the full version of the paper. Conceptually, our realizations of these sub-protocols follow the same design paradigm as our multiplication protocol described in Section 4.3, where all of the communication happens via a designated P_{King} , and the correctness of the computation is verified before output reconstruction via a verification phase.

5.2. Prototype implementation

We now describe a prototype implementation of Asterisk, which we use for benchmarking and comparison.

Environment. Benchmarks are performed over LAN and WAN using n1-highmem-64 instance of Google Cloud. The machine is equipped with a 2.3GHz n1-highmem-64, 64 core, Intel® Xeon® E5-2696V3 processor, with 416GB RAM memory. Each party is run as a process on a single

Protocol	Input	Output	Preprocessing	Online Comm	Online Round
Π_{mult}	$[[a], [b]]$	$[[x]] : x = a \cdot b$	3	$2n$	2
$\Pi_{\text{mult}3}$	$[[a], [b], [c]]$	$[[y]] : y = a \cdot b \cdot c$	9	$2n$	2
$\Pi_{\text{mult}4}$	$[[a], [b], [c], [d]]$	$[[z]] : z = a \cdot b \cdot c \cdot d$	23	$2n$	2
$\Pi_{\text{PrefixAND}}$	$[[x_1]^B, \dots, [x_k]^B]$	$[[y_1]^B, \dots, [y_k]^B] : y_i = \bigwedge_{j=1}^i x_j$	$\frac{35}{4} \log_4 k$	$\frac{3}{2} n \log_4 k$	$2 \log_4 k$
$\Pi_{k\text{-mult}}$	$[[x_1]^B, \dots, [x_k]^B]$	$[[b]^B] : b = \bigwedge_{i=1}^k x_i$	8	$\frac{2n}{3}$	$2 \log_4 k$
Π_{EQZ}	$[[x]]$	$[[b]^B] : b = 1(x \stackrel{?}{=} 0)$	12	$\frac{2n}{3}$	$2 \log_4 k$
Π_{BitA}	$[[b]^B]$	$[[b]]$	3	$2n$	2
Π_{LTZ}	$[[x]]$	$[[b]] : b = 1(x \stackrel{?}{\leq} 0)$	$7 + \frac{35}{4} \log_4 k$	$2n + \frac{2n}{k} + \frac{3n}{2} \log_4 k$	$2 \log_4 k + 4$
Π_{DotP}	$\{[[x_s], [y_s]]\}_{s \in [N]}$	$[[\sum_{s \in [N]} x_s \cdot y_s]]$	3	$2n$	2
Π_{shuffle}	$[[x_1], \dots, [x_N]]$	$[[x_{\pi(1)}], \dots, [x_{\pi(N)}]]$	$2N^2 + 3N$	$2nN$	2
Π_{sel}	$[[b]^B, [x_0], [x_1]]$	$[[x_{1-b}]]$	6	$4n$	4

Note: k is the number of bits to represent an element of \mathbb{F} ; N is the vector-length in Π_{DotP} and Π_{shuffle} .

TABLE 2: Summary of the various building blocks that we design for our applications

machine, and we use interprocess communication for emulating the actual communication. To simulate the distributed environments, we use the Linux `tc` command from the network emulation package `netem` to modify the bandwidth and latency. To simulate the LAN, we consider a bandwidth of 1 Gbps and 0.5 ms of latency, and for the WAN, we consider a bandwidth of 100 Mbps and 50 ms of latency. We demonstrate the performance of **Asterisk** over both LAN and WAN, and we provide a comparison with other works as well as the performance for privacy-preserving dark pools over WAN as the applications are more suitable over WAN.

Implementation of Asterisk. We implement **Asterisk** in C++17 using the codebase of QuadSquad [26] and the EMP toolkit [41]. We use computational security parameter $\kappa = 128$ and statistical security of at least 2^{-40} . We use AES-NI in the CTR mode to realize the PRFs, and use the NTL library [40] for all \mathbb{Z}_p operations (p being a 64-bit prime)⁴.

Implementation of other protocols. For comparison with [31] and [32], we rely on publicly available implementations from the widely used MP-SPDZ [30] framework. We use our own implementation for [37], since it does not have a publicly available implementation.

5.3. Performance benchmarks for Asterisk

We evaluate **Asterisk** on synthetic circuits with a total 1M multiplication gates. We consider circuits of depth $d = 10$, and 100 (with 100000 and 10000 multiplication gates at each level, resp.), as well as varying number of parties n ranging from $n = 5$ to $n = 100$, in both LAN and

4. We will open-source our code upon acceptance. Our code is developed for benchmarking and is not optimized for industry-grade use. We believe that incorporating state-of-the-art code optimizations like GPU-assisted computing can further enhance the efficiency of our protocols.

Depth	n	Total comm. (MB)	LAN time (sec)		WAN time (sec)	
			Preproc.	Online	Preproc.	Online
10	5	89.60	16.20	4.34	18.65	12.20
	10	169.60	27.68	4.57	30.10	18.41
	25	409.60	60.31	6.39	62.89	32.11
	50	809.60	115.25	12.08	115.78	63.11
	100	1609.60	226.41	20.60	230.27	127.75
100	5	89.60	15.57	4.54	18.25	54.30
	10	169.60	26.50	5.12	29.03	104.71
	25	409.60	58.02	7.19	60.60	255.94
	50	809.60	111.67	13.57	112.23	508.40
	100	1609.60	218.69	24.45	218.94	1013.40

TABLE 3: **Asterisk**'s total communication cost and run time over LAN and WAN networks for circuits with 10^6 multiplication gates, for varying depth ($d \in \{10, 100\}$) and varying numbers of parties ($n \in \{5, 10, 25, 50, 100\}$).

WAN network settings. The corresponding run times and communication overheads are reported in Table 3.

Observe that the overheads for the preprocessing phase are independent of d and n . This is precisely as expected (recall that the preprocessing phase of **Asterisk** is a one-shot message of size 24MB from the HP to P_n) and validates the high practical efficiency of the preprocessing phase. For the online phase, the time taken and communication overheads increase with d and n resp., which is again as expected (since the number of communication rounds increases with d). The overall communication, however, only scales with n and the *total number* of gates (and hence, is independent of d).

# of parties	Assisted MPC [37]		Asterisk	
	Comm. (MB)	Time (sec)	Comm. (MB)	Time (sec)
5	161	48.77	24	18.25
10	322	90.67	24	29.03
25	804	219.09	24	60.60
50	1608	432.33	24	112.23
100	3216	866.99	24	218.94

TABLE 4: Offline phase: Asterisk vs Assisted MPC [37].

Protocol	Communication (KB)	Throughput (/sec $\times 10^3$)
Mascot	8.45	0.19
Overdrive	7.35	0.24
Assisted MPC	0.161	20.50
Asterisk	0.024	54.79

TABLE 5: Offline phase overheads (per multiplication triple generation for 5 parties): Mascot [31] vs Overdrive [32] vs Assisted MPC [37] vs Asterisk.

Comparison with preprocessing phase of [37]. We compare the preprocessing phase of Assisted MPC [37] with that of Asterisk in Table 4 (Fig. 2) over a WAN network for varying number of parties n and for the same synthetic circuit with 10^6 multiplication gates⁵ (the online phase of [37] is identical to that of [32], and the corresponding comparison with Asterisk is presented subsequently in (Table 6). Unlike Asterisk, Assisted MPC requires the HP to communicate with all the parties in the preprocessing phase. Hence, the communication cost of Assisted MPC increases linearly with the n , whereas it remains constant for Asterisk. This allows Asterisk to save up to $134\times$ and $4\times$ in communication and run time overheads, respectively.

Comparison with preprocessing phase of [31], [32]. In Table 5, we compare the preprocessing phases of Asterisk with that of [31], [32] (we also include a data-point for Assisted MPC [37] for completeness) in terms of the cost of generating a single multiplication triple coupled with the throughput (number of multiplication triples that can be generated per second). We opt for such a comparison since [31], [32] use function-independent preprocessing phases which generate batches of multiplication triples. As expected, Asterisk provides a much better throughput, where the gain is up to $288\times$, $228\times$ and $2.7\times$, respectively, in comparison to [31], [32], and [37], respectively.

Comparison with online phase of [31], [32], [37]. In Table 6, we compare the online phase of Asterisk with the other dishonest majority protocols [31], [32], [37] (since all of these have identical online phases, we only report the results for [31]) over a WAN for a circuit with 10^6 multiplication gates with depth 100, and for varying numbers of parties n . Asterisk achieves communication savings

5. For a fair comparison, we let the helper in [37] be honest because [37] achieves its best efficiency in the presence of an honest helper.

# of parties	Dishonest Majority		Atlas		Asterisk	
	Comm. (MB)	Time (sec)	Comm. (MB)	Time (sec)	Comm. (MB)	Time (sec)
5	256	29.2	128	31.2	80	54.30
10	576	54.0	320	42.8	160	104.71
25	1536	142.0	768	109.4	400	255.94
50	3137	291.9	1600	441.5	800	508.40
100	6338	982.4	-	-	1600	1013.40

TABLE 6: Online phase: Asterisk vs [37], [31], [32] and [24].

$N = M$	$n = M + N$	Preproc.		Online		Throughput (/min)
		Comm. (KB)	Time (sec)	Comm. (KB)	Time (sec)	
5	10	6.55	0.03	7.27	8.55	70.18
10	20	12.67	0.09	29.64	16.69	71.90
25	50	31.06	0.39	179.14	31.43	95.45
50	100	61.69	1.33	644.69	21.99	272.85

TABLE 7: Communication, run time, and throughput (per minute) for dark pool VM algorithm for varying buy list size (N), and sell list size (M) and number of parties ($n = M + N$).

upto $4\times$. The relatively poor running time does not follow the asymptotic analysis (see Table 1, where Asterisk is clearly better than [31], [32], [37]), and is a consequence of the vast gap in terms of optimizations between our implementation (which is purely for benchmarking and additional overhead of NTL library [40]) and that of [31] in the MP-SPDZ library (which is highly optimized but does not support running Asterisk for a fair comparison due to the inherently function-dependent of its preprocessing framework). For a fair comparison of the running time, one should use a similarly optimized implementation of Asterisk, which we leave as an interesting future work.

Comparison with online phase of [24]. Additionally, we also compare against the state-of-the-art honest majority protocol Atlas [24] implemented in [30]. We note that the implementation allows running only the online phase of [24], which is reported in Table 6. We observe that our protocol has up to $2\times$ improvement in online communication and a comparable run time in comparison to [24], despite the fact that the implementation of [24] is highly optimized.

5.4. Benchmarks for Darkpool

We now benchmark the secure dark pool protocol described in [33], which is adapted to build upon Asterisk. Here, we do not compare Asterisk with other MPC protocols, since we expect precisely the same performance savings as in Section 5.3. We consider two popular algorithms—volume matching (VM) and continuous double auctions (CDA)—that are used for matching buy requests with appropriate sell requests in dark pools.

Volume matching: Consider a list \mathcal{B} of N buy orders and a list \mathcal{S} of M sell orders that are present in the system at a given point in time, where each order is also associated with a (distinct) client’s name and the number of units to be bought or sold, respectively. A buy order is said to match a sell order (and vice-versa) if the units in the buy order are

greater than or equal to the units of a sell order (known as the volume criteria). To completely satisfy a buy order, it is matched with every order in \mathcal{S} until this volume criteria fails. Thus, the buy order may have more than one matching sell order. In this way, in volume matching, matching is performed based on the above volume criteria, and either all the buy orders or all the sell orders are guaranteed to be satisfied.

Table 7 discusses the performance of Asterisk for VM, where the buyers in \mathcal{B} and sellers in \mathcal{S} enact the role of parties in MPC and run the privacy-preserving protocol for VM [33] among themselves. Since the performance of VM varies with the number of buyers (N) and sellers (M), we vary N, M to analyze its performance. Observe that even when N, M is as large as 50, Asterisk executes within a few seconds (up to 24 seconds). The total amount of communication is less than 750 KB. We also report the throughput, which captures the number of incoming orders that can be processed in a minute, and is computed as [10] $\text{throughput (/min)} = (M + N) * 60 / \text{online run time}$.

Continuous double auction: Unlike volume matching, where orders are matched in one shot, in CDA, orders are processed in a continuous manner with \mathcal{B} and \mathcal{S} constantly being updated to keep track of partially satisfied orders. Moreover, each order additionally comprises a price, indicating the price at which the buyer (seller) is willing to buy (sell) the stated units. Hence, for a buy order to match a sell order (and vice-versa), it is required that not only the aforementioned volume criteria should be satisfied, but the price criteria should also be met, i.e., the buying price of the buy order should be greater than or equal to the selling price of the sell order. The algorithm maintains the invariant that orders in the buy list are sorted in descending order as per the buying price, while orders in the sell list are sorted in ascending order as per the selling price. When a new buy order arrives, the algorithm identifies all possible matching sell orders based on whether the price criteria and volume criteria are satisfied. When either one of the matching criteria fails, the buy order may either be fully satisfied, i.e., all of its units are exhausted by getting matched to sell orders or, it may be partially satisfied, i.e., some of its units are still unmatched. If the buy order is partially satisfied, then it is inserted in the right position as per its buying price in the sorted \mathcal{B} . Thus, the CDA algorithm comprises a matching phase followed by an insertion phase (see [33] for details). Observe that due to the constantly changing \mathcal{B} and \mathcal{S} with state information being reused across the changes, we execute this algorithm in the secure outsourced computation setting. Here, a small set ($n = 5$) of powerful external servers are hired to carry out the MPC computation. The incoming buy or sell request is secret shared among the set of servers, which then execute the protocol on the secret-shared \mathcal{B}, \mathcal{S} and the incoming request, to obtain the result in a secret-shared manner.

Table 8 demonstrates the performance of Asterisk for CDA. Since the performance of CDA varies with N, M , we vary these while analyzing the performance. For processing

$N = M$	Preproc.		Online		Throughput (/min)
	Comm. (KB)	Time (sec)	Comm. (KB)	Time (sec)	
50	57.36	0.18	48.86	7.68	7.81
100	114.12	0.37	96.86	7.67	7.82
250	284.39	0.91	240.86	8.06	7.44
500	568.17	1.81	480.86	9.26	6.48

TABLE 8: Communication, run time, and throughput (per minute) for dark pool CDA algorithm for $n = 5$ and varying buy list size (N), and sell list size (M).

an incoming trade request, Asterisk executes within a few seconds (up to 12 seconds), even when the buy list size (N) and sell list size (M) is as large as 500. The total amount of communication is approximately 1 MB. We also report the throughput, which is computed [10] as $\text{throughput (/min)} = 60 / \text{online run time}$.

6. Conclusion and open questions

This work studies a new model, HP-aided MPC, which considers an additional semi-honest party. The presence of this semi-honest party helps to attain a security guarantee (fairness) which is known to be impossible, in general, for a standard dishonest majority setting. In this model, we build a framework, Asterisk, that achieves the following: (a) fairness security guarantee, (b) much better efficiency in comparison to dishonest majority protocols, (c) comparable efficiency with honest majority protocols, (d) scalability for a large number of parties. Moreover, the existence of the HP is natural in various applications, such as the Securities and Exchange Commission in the dark pool, which makes the model practically relevant.

Our work gives rise to many interesting open questions. It is worthwhile to check if one can extend the security of Asterisk from fairness to GOD *while preserving privacy against the semi-honest HP*. Note that if we allow the HP to learn the output, then such an extension is immediate: if the protocol aborts (implying at least one malicious corruption in \mathcal{P}), the HP must be honest (this follows immediately from our model), and can be used by the parties as a trusted third party (TTP) to compute the output. It is interesting to explore alternative practically motivated models that would allow for circumventing and improving the de facto security and efficiency bottleneck of dishonest majority MPC. One can explore the use of Asterisk for building additional privacy-preserving applications (e.g., federated cloud computing, where the assumption of a semi-honest HP is naturally justified).

References

- [1] Bar Alon, Eran Omri, and Anat Paskin-Cherniavsky. MPC with friends and foes. In *CRYPTO*, 2020.

- [2] Toshinori Araki, Assi Barak, Jun Furukawa, Tamar Lichter, Yehuda Lindell, Ariel Nof, Kazuma Ohara, Adi Watzman, and Or Weinstein. Optimized honest-majority mpc for malicious adversaries—breaking the 1 billion-gate per second barrier. In *IEEE S&P*, 2017.
- [3] Toshinori Araki, Jun Furukawa, Yehuda Lindell, Ariel Nof, and Kazuma Ohara. High-throughput semi-honest secure three-party computation with an honest majority. In *ACM CCS*, 2016.
- [4] Saikrishna Badrinarayanan, Abhishek Jain, Rafail Ostrovsky, and Ivan Visconti. Non-interactive secure computation from one-way functions. In *ASIACRYPT*, 2018.
- [5] Aner Ben-Efraim, Michael Nielsen, and Eran Omri. Turbospeedz: Double your online spdz! improving spdz using function dependent preprocessing. In *ACNS*, 2019.
- [6] Dan Boneh, Elette Boyle, Henry Corrigan-Gibbs, Niv Gilboa, and Yuval Ishai. Zero-knowledge proofs on secret-shared data via fully linear peps. In *CRYPTO*, 2019.
- [7] Elette Boyle, Niv Gilboa, Yuval Ishai, and Ariel Nof. Efficient fully secure computation via distributed zero-knowledge proofs. In *ASIACRYPT*, 2020.
- [8] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *FOCS*, 2001.
- [9] Ran Canetti and Marc Fischlin. Universally composable commitments. In *CRYPTO*, 2001.
- [10] John Carlidge, Nigel P Smart, and Younes Talibi Alaoui. Mpc joins the dark side. In *ACM ASIACCS*, 2019.
- [11] Octavian Catrina and Sebastiaan De Hoogh. Improved primitives for secure multiparty integer computation. In *Security and Cryptography for Networks*, 2010.
- [12] Nishanth Chandran, Wutichai Chongchitmate, Rafail Ostrovsky, and Ivan Visconti. Universally composable secure computation with corrupted tokens. In *CRYPTO*, pages 432–461, 2019.
- [13] Nishanth Chandran, Vipul Goyal, and Amit Sahai. New constructions for uc secure computation using tamper-proof hardware. In *EUROCRYPT*, 2008.
- [14] Seung Geol Choi, Jonathan Katz, Dominique Schröder, Arkady Yerukhimovich, and Hong-Sheng Zhou. (efficient) universally composable oblivious transfer using a minimal number of stateless tokens. In *TCC*, 2014.
- [15] Richard Cleve. Limits on the security of coin flips when half the processors are faulty. In *ACM STOC*, 1986.
- [16] Geoffroy Couteau. New protocols for secure equality test and comparison. In *ACNS*, 2018.
- [17] Mariana Botelho da Gama, John Carlidge, Antigoni Polychroniadou, Nigel P Smart, and Younes Talibi Alaoui. Kicking-the-bucket: Fast privacy-preserving trading using buckets. *Cryptology ePrint Archive*, 2021.
- [18] Dana Dachman-Soled, Tal Malkin, Mariana Raykova, and Muthuramakrishnan Venkatasubramanian. Adaptive and concurrent secure computation from new adaptive, non-malleable commitments. In *ASIACRYPT*, 2013.
- [19] Ivan Damgård, Marcel Keller, Enrique Larraia, Valerio Pastro, Peter Scholl, and Nigel P Smart. Practical covertly secure mpc for dishonest majority—or: breaking the spdz limits. In *ESORICS*, 2013.
- [20] Ivan Damgård and Jesper Buus Nielsen. Scalable and unconditionally secure multiparty computation. In *CRYPTO*, 2007.
- [21] Matthias Fitzi, Juan A. Garay, Ueli M. Maurer, and Rafail Ostrovsky. Minimal complete primitives for secure multi-party computation. In *CRYPTO*, 2001.
- [22] S. Dov Gordon, Yuval Ishai, Tal Moran, Rafail Ostrovsky, and Amit Sahai. On complete primitives for fairness. In *TCC*, 2010.
- [23] S. Dov Gordon, Yuval Ishai, Tal Moran, Rafail Ostrovsky, and Amit Sahai. On complete primitives for fairness. In *TCC*, 2010.
- [24] Vipul Goyal, Hanjun Li, Rafail Ostrovsky, Antigoni Polychroniadou, and Yifan Song. Atlas: Efficient and scalable mpc in the honest majority setting. In *CRYPTO*, 2021.
- [25] Carmit Hazay, Antigoni Polychroniadou, and Muthuramakrishnan Venkatasubramanian. Composable security in the tamper-proof hardware model under minimal complexity. In *TCC*, 2016.
- [26] Aditya Hegde, Nishat Koti, Varsha Bhat Kukkala, Shravani Patil, Arpita Patra, and Protik Paul. Attaining god beyond honest majority with friends and foes. In *ASIACRYPT*, 2022.
- [27] Yuval Ishai, Rafail Ostrovsky, and Hakan Seyalioglu. Identifying cheaters without an honest majority. In *TCC*, 2012.
- [28] Yuval Ishai, Arpita Patra, Sikhar Patranabis, Divya Ravi, and Akshayaram Srinivasan. Fully-secure MPC with minimal trust. In *TCC*, 2022.
- [29] Jonathan Katz. Universally composable multi-party computation using tamper-proof hardware. In *EUROCRYPT*, 2007.
- [30] Marcel Keller. MP-SPDZ: A versatile framework for multi-party computation. In *ACM CCS*, 2020.
- [31] Marcel Keller, Emmanuela Orsini, and Peter Scholl. Mascot: faster malicious arithmetic secure computation with oblivious transfer. In *ACM CCS*, 2016.
- [32] Marcel Keller, Valerio Pastro, and Dragos Rotaru. Overdrive: making spdz great again. In *EUROCRYPT*, 2018.
- [33] Nishat Koti, Varsha Bhat Kukkala, Arpita Patra, and Bhavish Raj Gopal. Pentagod: Stepping beyond traditional god with five parties. In *ACM CCS*, 2022.
- [34] Nishat Koti, Shravani Patil, Arpita Patra, and Ajith Suresh. Mpplan: Protocol suite for privacy-conscious computations. *Journal of Cryptology*, 2023.
- [35] Nishat Koti, Arpita Patra, Rahul Rachuri, and Ajith Suresh. Tetrad: Actively Secure 4PC for Secure Training and Inference. In *NDSS*, 2022.
- [36] Nancy A Lynch. *Distributed algorithms*. Elsevier, 1996.
- [37] Philipp Muth and Stefan Katzenbeisser. Assisted mpc. *Cryptology ePrint Archive*, 2022.
- [38] Arpita Patra, Thomas Schneider, Ajith Suresh, and Hossein Yalame. Aby2. 0: Improved mixed-protocol secure two-party computation. In *USENIX Security*, 2021.
- [39] Antigoni Polychroniadou, Gilad Asharov, Benjamin Diamond, Tucker Balch, Hans Buehler, Richard Hua, Suwen Gu, Greg Gimler, and Manuela Veloso. Prime match: A privacy-preserving inventory matching system. *Cryptology ePrint Archive*, 2023.
- [40] Victor Shoup. NTL: A Library for doing Number Theory. <https://libntl.org/>, 2021.
- [41] Xiao Wang, Alex J. Malozemoff, and Jonathan Katz. EMP-toolkit: Efficient MultiParty computation toolkit. <https://github.com/emp-toolkit>, 2016.
- [42] Andrew C Yao. Protocols for secure computations. In *FOCS*, 1982.

Appendix A. Security model

The real world.. An n -party protocol Π with n parties $\mathcal{P} = \{P_1, \dots, P_n\}$ is an n -tuple of probabilistic polynomial-time (PPT) interactive Turing machines (ITMs),

We let \mathcal{A} denote a special PPT ITM that represents the adversary and that is initialized with input that contains the identities of the corrupt parties, their respective private inputs, and an auxiliary input. During the execution of the protocol, the maliciously corrupt parties (sometimes

referred to as ‘active’) receive arbitrary instructions from the adversary \mathcal{A} , while the honest parties and the semi-honestly corrupt (sometimes referred to as ‘passive’) parties faithfully follow the instructions of the protocol. We consider the adversary \mathcal{A} to be rushing, i.e., during every round the adversary can see the messages the honest parties sent before producing messages from actively corrupt parties.

At the end of the protocol execution, the honest parties produce output, the corrupt parties produce no output, and the adversary outputs an arbitrary function of its view. The view of a party during the execution consists of its input, random coins and the messages it sees during the execution.

Definition A.1 (Real-world execution). *Let Π be an n -party protocol amongst $\{P_1, \dots, P_n\}$ computing an n -party function $f : (\{0, 1\}^*)^n \rightarrow (\{0, 1\}^*)^n$ and let $\mathcal{C} \subset [1, n] \cup \{\text{HP}\}$ denote the set of indices of the corrupted parties. The execution of Π under $(\mathcal{A}, \mathcal{C})$ in the real world, on input vector $\vec{x} = (x_1, \dots, x_n)$, auxiliary input aux and security parameter κ , denoted $\text{real}_{\Pi, \mathcal{C}, \mathcal{A}(\text{aux})}(\vec{x}, \kappa)$, is defined as the output vector of P_1, \dots, P_n and $\mathcal{A}(\text{aux})$ resulting from the protocol interaction.*

The ideal world. We describe ideal world execution with the fairness security guarantee.

Definition A.2 (Ideal Computation). *Let $f : (\{0, 1\}^*)^n \rightarrow (\{0, 1\}^*)^n$ be an n -party function and let $\mathcal{C} \subset [1, n] \cup \{\text{HP}\}$ denotes the set of all corrupt parties. Then, the joint ideal execution of f under $(\mathcal{S}, \mathcal{C})$ on input vector $\vec{x} = (x_1, \dots, x_n)$, auxiliary input aux to \mathcal{S} and security parameter κ , denoted $\text{ideal}_{f, \mathcal{C}, \mathcal{S}(\text{aux})}(\vec{x}, \kappa)$, is defined as the output vector of P_1, \dots, P_n and \mathcal{S} resulting from the following ideal process.*

- 1) Parties send inputs to trusted party: *An honest party P_i sends its input x_i to the trusted party. The simulator \mathcal{S} may send to the trusted party arbitrary inputs for the actively corrupt parties. Let x'_i be the value actually sent as the input of party P_i .*
- 2) Trusted party speaks to simulator: *The trusted party computes $(y_1, \dots, y_n) = f(x'_1, \dots, x'_n)$. The trusted party sends ready to \mathcal{S} .*
- 3) Simulator \mathcal{S} responds to trusted party: *The simulator can send abort to the trusted party.*
- 4) Trusted party answers parties:
 - a) *If the trusted party got abort from the simulator \mathcal{S} , it sets the abort message $\text{abortmsg} = \perp$.*
 - b) *Otherwise, it sends y_j to every P_j , $j \in [1, n]$.*
- 5) *Outputs: Honest parties always output the message received from the trusted party, while corrupt parties output nothing. The simulator \mathcal{S} outputs an arbitrary function of the initial inputs $\{x_i\}_{i \in \mathcal{C}}$, the messages received by the corrupt parties from the trusted party and its auxiliary input.*

In this work, we primarily consider the non-colluding security notion of [28], which is described below.

Non-colluding security. Informally, protocols with this notion of security are secure against a non-colluding

adversary that corrupts *either* any subset of the n parties $\{P_1, \dots, P_n\}$ maliciously *or* the HP passively (i.e. in a semi-honest manner). The formal definition appears below.

Definition A.3 (Non-colluding security). *Let $f : (\{0, 1\}^*)^n \rightarrow (\{0, 1\}^*)^n$ be an n -party function. A protocol Π securely computes the function f in the non-colluding model with fairness security if for any adversary \mathcal{A} , there exists a simulator \mathcal{S} such that for every \mathcal{C} where either (a) $\mathcal{C} \subset [1, n]$ malicious corruptions or (b) $\mathcal{C} = \text{HP}$ semi-honest corruption, we have*

$$\{\text{real}_{\Pi, \mathcal{C}, \mathcal{A}(\text{aux})}(\vec{x}, \kappa)\}_{\vec{x} \in (\{0, 1\}^*)^n, \kappa \in \mathbb{N}} \equiv \{\text{ideal}_{f, \mathcal{C}, \mathcal{S}(\text{aux})}(\vec{x}, \kappa)\}_{\vec{x} \in (\{0, 1\}^*)^n, \kappa \in \mathbb{N}}.$$

Note that the corruption is non-simultaneous. Therefore we need the above indistinguishability to hold in both the corruption cases.

A protocol achieves computational security, if the above distributions are computationally close in the presence of the parties, \mathcal{A} , \mathcal{S} that are PPT. A protocol achieves statistical (resp. perfect) security if the distributions are statistically close (resp. identical).

A.1. Alternative corruption models

As mentioned above, we consider a non-colluding adversary who either corrupts up to $n - 1$ among the n parties maliciously or the HP in a semi-honest manner. One might also consider the following alternative corruption models.

A.1.1. Colluding security. Informally, in this corruption model, the adversary can *simultaneously* corrupt the majority of the parties maliciously as well as the HP in a semi-honest manner. More formally, security is defined similar to non-colluding security (defined above), except that the indistinguishability must hold for every $\mathcal{C} \subset [1, n] \cup \{\text{HP}\}$, where the corruptions in $[1, n]$ are malicious and corruption of HP is semi-honest.

A.1.2. Friends and foes (FaF) security with HP. When proving security in the FaF model (as discussed in Section 1), there is the additional requirement of simulating the view of any subset of uncorrupted (or semi-honest) parties, in addition to simulating the view of the maliciously corrupt parties. This necessitates the use of two simulators in the ideal world— one for the malicious adversary and one for the semi-honest adversary. Further, the malicious adversary is allowed to send its entire view to the semi-honest adversary in the ideal world to capture the behaviour where the malicious adversary may send non-protocol messages to uncorrupted parties in the real world. Elaborately, let \mathcal{A} denote the probabilistic polynomial time (PPT) malicious adversary in the real-world corrupting t parties in $\mathcal{I} \subset \mathcal{P}$, and $\mathcal{S}_{\mathcal{A}}$ denote the corresponding ideal-world simulator. Similarly, let $\mathcal{A}_{\mathcal{H}}$ denote the (PPT) semi-honest adversary corrupting h^* parties in $\mathcal{H} \subset \mathcal{P} \setminus \mathcal{I}$ in the real-world, and $\mathcal{S}_{\mathcal{A}_{\mathcal{H}}}$ be the ideal-world simulator. Note that in the classical definition of ideal-world, $\mathcal{H} = \emptyset$ and FaF-security with HP

definition, $\mathcal{H} = \{\text{HP}\}$ and $t = n - 1$. Unlike the FaF-security model, where security is required not only against t maliciously corrupt parties, but also against *any* subset of h^* semi-honest parties, FaF-security with HP is a special case. Here, we require semi-honest security to be provided only against the single party designated as the HP. Let \mathcal{A} be an adversary corrupting at most $n - 1$ parties maliciously, and let \mathcal{A}_{HP} be a semi-honest adversary that corrupts the HP. We say a protocol, Π , is FaF-secure with HP if Let \mathcal{F} be the ideal-world functionality. Let $\text{VIEW}_{\mathcal{A},\Pi}^{\text{REAL}}(1^\lambda, z_{\mathcal{A}})$ be the malicious adversary's (\mathcal{A}) view and $\text{OUT}_{\mathcal{A},\Pi}^{\text{REAL}}(1^\lambda, z_{\mathcal{A}})$ denote the output of the uncorrupted parties (in $\mathcal{P} \setminus \mathcal{I}$) during a random execution of Π , where $z_{\mathcal{A}}$ is the auxiliary input of \mathcal{A} . Similarly, let $\text{VIEW}_{\mathcal{A},\mathcal{A}_{\text{HP}},\Pi}^{\text{REAL}}(1^\lambda, z_{\mathcal{A}}, z_{\mathcal{A}_{\text{HP}}})$ be the semi-honest \mathcal{A}_{HP} 's view during an execution of Π running alongside \mathcal{A} , where $z_{\mathcal{A}_{\text{HP}}}$ is the auxiliary input of \mathcal{A}_{HP} . Note that $\text{VIEW}_{\mathcal{A},\mathcal{A}_{\text{HP}},\Pi}^{\text{REAL}}(1^\lambda, z_{\mathcal{A}}, z_{\mathcal{A}_{\text{HP}}})$ consists of the non-prescribed messages sent by the malicious adversary to \mathcal{A}_{HP} . Correspondingly, let $\text{VIEW}_{\mathcal{S}_{\mathcal{A}},\mathcal{F}}^{\text{IDEAL}}(1^\lambda, z_{\mathcal{A}})$ be the malicious adversary's simulated view with $\mathcal{S}_{\mathcal{A}}$ corrupting parties in \mathcal{I} and $\text{OUT}_{\mathcal{S}_{\mathcal{A}},\mathcal{F}}^{\text{IDEAL}}(1^\lambda, z_{\mathcal{A}})$ denote the output of the uncorrupted parties (in $\mathcal{P} \setminus \mathcal{I}$) during a random execution of ideal-world functionality \mathcal{F} . Similarly, let $\text{VIEW}_{\mathcal{S}_{\mathcal{A}},\mathcal{S}_{\mathcal{A}_{\text{HP}}},\mathcal{F}}^{\text{IDEAL}}(1^\lambda, z_{\mathcal{A}}, z_{\mathcal{A}_{\text{HP}}})$ be the semi-honest \mathcal{A}_{HP} 's simulated view with $\mathcal{S}_{\mathcal{A}_{\text{HP}}}$ corrupting \mathcal{A}_{HP} during an execution of \mathcal{F} running alongside \mathcal{A} . A protocol Π is said to compute \mathcal{F} with computational FaF-security with HP if

$$\begin{aligned} & (\text{VIEW}_{\mathcal{S}_{\mathcal{A}},\mathcal{F}}^{\text{IDEAL}}(1^\lambda, z_{\mathcal{A}}), \text{OUT}_{\mathcal{S}_{\mathcal{A}},\mathcal{F}}^{\text{IDEAL}}(1^\lambda, z_{\mathcal{A}})) \equiv \\ & \quad (\text{VIEW}_{\mathcal{A},\Pi}^{\text{REAL}}(1^\lambda, z_{\mathcal{A}}), \text{OUT}_{\mathcal{A},\Pi}^{\text{REAL}}(1^\lambda, z_{\mathcal{A}})), \\ & (\text{VIEW}_{\mathcal{S}_{\mathcal{A}},\mathcal{S}_{\mathcal{A}_{\text{HP}}},\mathcal{F}}^{\text{IDEAL}}(1^\lambda, z_{\mathcal{A}}, z_{\mathcal{A}_{\text{HP}}}), \text{OUT}_{\mathcal{S}_{\mathcal{A}},\mathcal{F}}^{\text{IDEAL}}(1^\lambda, z_{\mathcal{A}})) \equiv \\ & \quad (\text{VIEW}_{\mathcal{A},\mathcal{A}_{\text{HP}},\Pi}^{\text{REAL}}(1^\lambda, z_{\mathcal{A}}, z_{\mathcal{A}_{\text{HP}}}), \text{OUT}_{\mathcal{A},\Pi}^{\text{REAL}}(1^\lambda, z_{\mathcal{A}})). \end{aligned}$$

Here, $\text{VIEW}_{\mathcal{A},\mathcal{A}_{\text{HP}},\Pi}^{\text{REAL}}(1^\lambda, z_{\mathcal{A}}, z_{\mathcal{A}_{\text{HP}}})$ is \mathcal{A}_{HP} 's simulated view during an execution of \mathcal{F} alongside \mathcal{A} , and $\text{VIEW}_{\mathcal{S}_{\mathcal{A}},\mathcal{S}_{\mathcal{A}_{\text{HP}}},\mathcal{F}}^{\text{IDEAL}}(1^\lambda, z_{\mathcal{A}}, z_{\mathcal{A}_{\text{HP}}})$ consists of non-prescribed messages sent by \mathcal{A} to \mathcal{A}_{HP} .

Appendix B.

Proof of Theorem 3.1

We present the proof for a two-party functionality. The argument can be extended for any number of parties using the player-partitioning argument. We refer to the two parties as A, B .

Lemma B.1. *Assume that one-way permutation exists. Then there exists a 2-party functionality that no protocol computes with fairness in the FaF security model with HP (refer to precise the definition given in Section A.1.2).*

Proof of Lemma B.1. Let $f = \{f_\kappa : \{0, 1\}^\kappa \rightarrow \{0, 1\}^\kappa\}_{\kappa \in \mathbb{N}}$ be a one-way permutation. Define the symmetric 2-party functionality $\text{Swap} = \{\text{Swap}_\kappa : \{0, 1\}^{2\kappa} \times \{0, 1\}^{2\kappa} \rightarrow \{0, 1\}^{2\kappa}\}_{\kappa \in \mathbb{N}}$ as follows. Parties A and B each

hold string (a, y_B) and (b, y_A) respectively. The output is then defined to be

$$\text{Swap}_\kappa((a, y_B), (b, y_A)) = \begin{cases} (a, b) & \text{if } f_\kappa(a) = y_A \text{ and } f_\kappa(b) = y_B \\ 0^\kappa & \text{otherwise} \end{cases}$$

For the sake of contradiction, suppose a protocol Π securely computes the Swap functionality with fairness FaF security with HP model. We fix a security parameter κ and let r denote the number of rounds in Π . Consider an evaluation of Swap with the output being (a, b) . Formally, we consider the following distribution over the inputs.

- a, b are each selected from $\{0, 1\}^\kappa$ uniformly at random and independently.
- $y_A = f_\kappa(a)$ and $y_B = f_\kappa(b)$.

For $i \in \{0, \dots, r\}$ let a_i be the final output of A assuming that B aborted after sending i th round messages. Similarly, for $i \in \{0, \dots, r\}$ we define b_i to be the final output of B assuming that A aborted after sending i th round messages. Observe that a_r (b_r) is the output of A (B) when B (A) sends messages in all the rounds. We first claim that there exists a round where either A and HP jointly or B and HP jointly, gain an advantage in computing the correct output.

Claim B.2. *Either there exists $i \in [0, r]$ such that*

$$\Pr[a_i = (a, b)] - \Pr[b_i = (a, b)] \geq \frac{1 - \text{neg}(\kappa)}{2r + 1},$$

or there exists $i \in [1, r]$ such that

$$\Pr[b_i = (a, b)] - \Pr[a_{i-1} = (a, b)] \geq \frac{1 - \text{neg}(\kappa)}{2r + 1}.$$

The probabilities above are taken over the choice of inputs and of random coins for the parties.

Proof. The proof follows by the following averaging argument. By correctness and the fact that f_κ is one-way, it follows that

$$\begin{aligned} 1 - \text{neg}(\kappa) & \leq \Pr[a_r = (a, b)] - \Pr[b_0 = (a, b)] \\ & = \sum_{i=0}^r (\Pr[a_i = (a, b)] - \Pr[b_i = (a, b)]) \\ & \quad + \sum_{i=1}^r (\Pr[b_i = (a, b)] - \Pr[a_{i-1} = (a, b)]) \end{aligned}$$

Since there are $2r + 1$ summands, there must exist an i for which one of the differences is at least $\frac{1 - \text{neg}(\kappa)}{2r + 1}$. \square

Assume without loss of generality that there exists an $i \in [1, r]$ such that the former equality in Claim B.2 holds (the other case is done analogously). Define a malicious adversary \mathcal{A} as follows. For the security parameter κ , it receives the round i as auxiliary input. Now, \mathcal{A} corrupts A and makes it act honestly (using the party's original input a) up to and including round i . After receiving the i -th message, the adversary instructs A to abort. Finally, the adversary sends its entire view to HP. Note that this is allowed as per the FaF security notion as nothing stops

a malicious adversary from sending its view to the semi-honest HP in the real-world.

Next, we show that no pair of simulators \mathcal{S}_A and $\mathcal{S}_{A,HP}$ can produce views for A and HP so that the ideal view and the real view are indistinguishable. For that, we assume towards contradiction that such simulators do exist. Let $(a^*, y_B^*) \in \{0, 1\}^{2\kappa}$ be the input that \mathcal{S}_A sent to the ideal functionality. Additionally, denote $q = \Pr[b_i = (a, b)]$.

Now depending upon $\Pr[(a^*, y_B^*) = (a, y_B)]$, two scenarios are possible. We will analyze them case by case.

Case I: $\Pr[(a^*, y_B^*) = (a, y_B)] \geq q + \frac{1}{p(\kappa)}$ for some polynomial $p(\cdot)$ for infinitely many κ 's. Let $\text{OUT}_B^{\text{IDEAL}}$ be B 's output in the ideal world. Since f_κ is a permutation, we have that $\Pr[\text{OUT}_B^{\text{IDEAL}} = (a, b)] = \Pr[(a^*, y_B^*) = (a, y_B)] \geq q + \frac{1}{p(\kappa)}$. Thus, by comparing the output of B to (a, b) it is possible to distinguish the real from the ideal with advantage at least $1/p(\kappa)$.

Case II: $\Pr[(a^*, y_B^*) = (a, y_B)] \leq q + \text{neg}(\kappa)$. Here we show how to distinguish between the view of HP in the real world from its ideal world counterpart. Recall that in the real world A sent its view to HP. Let M be the algorithm specified by the protocol, that A and HP use to compute their output assuming B has aborted. Namely, M outputs a_i in the real world. By Claim B.2 it holds that $\Pr[a_i = (a, b)] \geq q + \frac{1 - \text{neg}(\kappa)}{2r+1}$. We next consider the ideal world. Let V be the view generated by $\mathcal{S}_{A,HP}$. We claim that $\Pr[M(V) = (a, b) \wedge (a^*, y_B^*) \neq (a, y_B)] \leq \text{neg}(\kappa)$.

Note that $\Pr[M(V) = (a, b) \wedge a^* \neq a] \leq \text{neg}(\kappa)$. Since f_κ is a permutation and B does not change the input it sends to the ideal functionality, (i.e. B sends its input (b, y_A)) the output computed by the ideal functionality will be 0^κ . Moreover, as f_κ is one-way, it follows that if $M(V)$ did output (a, b) , then it can be used to break the security of f_κ . At a high level, this is because M has computed $f_\kappa^{-1}(y_B)$. Elaborately, this can be done by sampling $a \leftarrow \{0, 1\}^\kappa$, computing $f_\kappa(a)$, and finally, computing a view V using the simulators and applying M to it (if a^* computed by \mathcal{S}_A equals to a then abort).

On the other hand, $\Pr[M(V) = (a, b) \wedge y_B^* \neq y_B] \leq \text{neg}(\kappa)$. Since B does not change its input b to the ideal functionality, thus $f_\kappa(b) = y_B \neq y_B^*$, due to the well-definedness of the function f_κ . Therefore, the output computed by the ideal functionality will be 0^κ .

Hence,

$$\begin{aligned} & \Pr[M(V) = (a, b) \wedge (a^*, y_B^*) \neq (a, y_B)] \\ \leq & \Pr[M(V) = (a, b) \wedge a^* \neq a] + \Pr[M(V) = (a, b) \wedge y_B^* \neq y_B]. \end{aligned}$$

Thus, $\Pr[M(V) = (a, b) \wedge (a^*, y_B^*) \neq (a, y_B)] \leq \text{neg}(\kappa)$
Hence, we conclude that

$$\begin{aligned} & \Pr[M(V) = (a, b)] \\ = & \Pr[M(V) = (a, b) \wedge (a^*, y_B^*) = (a, y_B)] \\ & + \Pr[M(V) = (a, b) \wedge (a^*, y_B^*) \neq (a, y_B)] \\ \leq & \Pr[(a^*, y_B^*) = (a, y_B)] + \text{neg}(\kappa) \leq q + \text{neg}(\kappa). \end{aligned}$$

Therefore, by applying M to the view it is possible to distinguish with advantage at least $\frac{1 - \text{neg}(\kappa)}{2r+1} - \text{neg}(\kappa)$.

Appendix C. Asterisk

C.1. Shared key setup

Here we provide details about how $\mathcal{F}_{\text{Setup}}$ (Fig. 3) can be instantiated. Note that $\{k_i\}_{i \in [1, n]}$ and k_{all} can be sampled by HP and sent to the corresponding parties. However, generating $k_{\mathcal{P}}$ is slightly challenging since this key should be generated while ensuring that all parties agree on the key despite misbehaviour of corrupt parties. We rely on the following abort secure protocol for generating this key. To generate $k_{\mathcal{P}}$, each party sends a randomly sampled value to all other parties. Upon receiving values from all parties, each party takes the sum of the received value and sets that as the common key. To verify that all parties have agreed upon a common key, they rely on the HP as follows. All parties evaluate the PRF on a common counter (ctr) value using the common key and send the output to HP. If all the received values at the HP match, then HP sends continue to all the parties; else, it sends abort. If parties receive abort, they terminate the protocol. Otherwise, the protocol continues with $k_{\mathcal{P}}$ as the common key among the parties. Note that, though the above-mentioned procedure provides only abort security, while running it inside Asterisk it does not impact the fairness guarantee of Asterisk. This is because aborting during key setup which is an input independent phase does not allow the adversary to learn any information about the output.

C.2. Security proof

The formal security proof of Asterisk is provided here.

Simulator \mathcal{S}_A

Malicious Let \mathcal{A} be a malicious adversary corrupting up to $n - 1$ parties among the computing parties $\mathcal{P} = \{P_1, \dots, P_n\}$. Therefore there is at least one honest party, say P_h . In this case, recall that HP is also honest. Let \mathcal{S}_A denote the simulator for this case.

- **Preprocessing:** \mathcal{S}_A simulates the preprocessing phase of the protocol by acting on behalf of HP as per the protocol specifications.
- **Input:** Let P_d be the input provider.
 - Case I: $d \neq h$. In the online phase, \mathcal{S}_A receives q_v from P_d (since $P_d \neq P_h$, thus corrupted by \mathcal{A}) and sends q_v to all the parties. \mathcal{S}_A extracts P_d 's input $v = q_v - r - \delta_v$. Note that r is the common random pad known to P_h and δ_v is known to HP from the preprocessing phase.
 - Case II: $d = h$. In the online phase, \mathcal{S}_A samples a random value q_v and sends it to all the parties on behalf of HP.
- **Multiplication:**
 - Case I: Let P_{King} be corrupt. In the online phase, \mathcal{S}_A computes $[m_z]_h$ according to the protocol. Since \mathcal{S}_A has $[\delta_x]_h, [\delta_y]_h, [\delta_{xy}]_h, [\delta_z]_h$ from the preprocessing step and m_x, m_y is available to all the parties (excluding HP), therefore it can compute $[m_z]_h$ correctly. \mathcal{S}_A sends it to P_{King} on behalf

of P_h . Then it receives m_z from P_{King} .

– If \mathcal{S}_A receives multiple m_z from P_{King} ^a, then \mathcal{S}_A sets adversary's message to \mathcal{F}_{MPC} as abort. \mathcal{S}_A continues the steps of the protocol on behalf of the honest parties with the received values.

– If $m_z \neq m_x m_y - m_x \delta_y - m_y \delta_x + \delta_{xy} + \delta_z$, then \mathcal{S}_A sets adversary's message to \mathcal{F}_{MPC} as abort. \mathcal{S}_A continues the steps of the protocol on behalf of the honest parties with the received values.

– If $m_z = m_x m_y - m_x \delta_y - m_y \delta_x + \delta_{xy} + \delta_z$ then \mathcal{S}_A , then \mathcal{S}_A continues the steps of the protocol on behalf of the honest parties with the received values.

• **Case II:** Let P_{King} be honest. In the online phase, \mathcal{S}_A receives shares of $[m_z]_i$ from the corrupt P_i . If $[m_z]_i \neq m_x m_y - m_x [\delta_y]_i - m_y [\delta_x]_i + [\delta_{xy}]_i + [\delta_z]_i$ then it sets \mathcal{F}_{MPC} message as abort. \mathcal{S}_A sends m_z to all the parties on behalf of P_{King} by following the protocol steps honestly.

• **Verification:** \mathcal{S}_A simulates $\mathcal{F}_{\text{Rand}}$ and obtains ρ . Further, it obtains a sharing of zero, $\{\alpha_1, \alpha_2, \dots, \alpha_n\}$ such that $\sum_{i=1}^n \alpha_i = 0$. For all parties P_i corrupted by \mathcal{A} , \mathcal{S}_A receives $[\omega_z]_i$. \mathcal{S}_A checks if $[\omega_z]_i = \rho_0 \cdot \alpha_i + \sum_{j \in [m]} \rho_j \cdot ([t_{m_{z_j}}]_i - m_{z_j}^* \cdot [\Delta]_i)$, where $m_{z_j}^*$ be the unique opened value. Note that $[t_{m_{z_j}}]_i = -m_{x_j} \cdot [t_{y_j}]_i - m_{y_j} \cdot [t_{x_j}]_i + [t_{x_j y_j}]_i + [t_{z_j}]_i + [\Delta]_i \cdot (m_{x_j} \cdot m_{y_j})$, where m_{x_j}, m_{y_j} for all $j \in [1, m]$ are held by the honest party P_h and $[t_{x_j}], [t_{y_j}], [t_{x_j y_j}], [t_{z_j}]$ are held by HP.

• **Output:** If the above verification fails, \mathcal{S}_A sets the adversary's message to \mathcal{F}_{MPC} as abort. \mathcal{S}_A invokes \mathcal{F}_{MPC} with the adversary's input (\perp or v) Recall that \mathcal{S}_A had extracted corrupt P_i 's input during the input stage. If the adversary's input is not abort, then \mathcal{F}_{MPC} gives y as the output, then it computes $\delta_y = m_y - y$ and sends on behalf of HP.

a. this corresponds to the case when there is more than one honest party (on behalf of whom the simulator is acting) and P_{King} sends inconsistent information to them.

Figure 13: Simulator \mathcal{S} for Π_{MPC} for a malicious adversary corrupting up to $n - 1$ parties among n parties

Indistinguishability.

Case I: Fig. 13 describes the simulation steps when a malicious adversary corrupts up to $n - 1$ parties among n parties. In this scenario HP and party P_h is honest. Consider \mathcal{A} be a malicious adversary corrupting up to $n - 1$ parties (i.e. all parties excluding P_h) among the party set \mathcal{P} .

For input gates, if P_d is corrupt, then \mathcal{S}_A does not send any messages. Thus, the view is indistinguishable trivially. If P_d is honest then \mathcal{S}_A sends a random value \tilde{q}_v to \mathcal{A} . In the real protocol, P_d sends q_v to HP, which is forwarded to all the parties. Here $q_v = m_v + r$, where $m_v = v + \delta_v$. Since δ_v is sampled uniformly, therefore q_v is distributed uniformly. Hence the distribution of q_v and \tilde{q}_v are identical.

For a multiplication gate, if P_{King} is corrupt, then in the protocol, \mathcal{A} receives $[m_z]_h$ from P_h and let \mathcal{S}_A sends the same on behalf of P_h . Thus, the simulated view and the real view are identically distributed. Let P_{King} sends m_z to \mathcal{S}_A . In the first two cases, \mathcal{S}_A sets the message to the functionality as abort. Let's consider P_{King} has introduced

an error ϵ_i towards an honest P_i 's m_z . Then during the verification step, P_i computes $[\omega_z]_i = [t_{m_z}]_i - (m_z + \epsilon_i) \cdot [\Delta]_i$. Then $\sum_{i \in [1, n]} [\omega_z]_i \neq 0$ with a very high probability, since $[t_{m_z}]_i, [\Delta]_i$ is unknown to P_{King} . Thus, a random linear combination of ω_z will remain non-zero. Therefore, the real protocol aborts with high probability, and in the simulated protocol, \mathcal{S}_A sets the adversary's message to \mathcal{F}_{MPC} as abort. If the m_z sent by P_{King} is correct, that is, the third case, then \mathcal{S}_A receives the correct value, and it follows the protocol steps correctly, therefore the generated view is indistinguishable. If P_{King} is honest, then \mathcal{S}_A receives shares of m_z from all the corrupt parties. If it receives an incorrect share, then \mathcal{S}_A sets the \mathcal{F}_{MPC} message as abort, and in the real protocol, the verification fails with a very high probability. \mathcal{S}_A sends m_z to all the parties by honestly following the protocol steps.

Next, we analyze the difference between the verification check carried out by the simulator \mathcal{S}_A and the verification check in the real-world protocol. In the verification step, \mathcal{A} can introduce an error in m_{z_j} . However, it can obtain a sharing of the corresponding tag with negligible probability $(\frac{1}{|\mathbb{F}|})$. Thus, consider a $j \in [1, m]$ such that $[t_{m_{z_j}}]_i - m_{z_j}^* \cdot [\Delta]_i \neq 0$, however

$\sum_{i=1}^n [\rho_0 \cdot \alpha_i + \sum_{j \in [m]} \rho_j \cdot ([t_{m_{z_j}}]_i - m_{z_j}^* \cdot [\Delta]_i)] = 0$. This is possible for some choice of ρ . To elaborate, consider $\xi_0 = \sum_{i=1}^n \alpha_i$ and $\xi_j = \sum_{i=1}^n ([t_{m_{z_j}}]_i - m_{z_j}^* \cdot [\Delta]_i)$ for all $j \in [1, m]$, set $\xi = (\xi_0, \xi_1, \dots, \xi_m)$. Then $T_\xi(\rho) = \sum_{j=0}^m \xi_j \cdot \rho_j$. T_ξ is a non-zero linear map, then $\dim(\ker(T_\xi)) \leq m$. Therefore, the probability that a randomly sampled ρ belongs to $\ker(T_\xi) \leq \frac{|\mathbb{F}|^m}{|\mathbb{F}|^{m+1}} = \frac{1}{|\mathbb{F}|}$. \mathcal{S}_A gets $[\omega_z]_i$ from all parties P_i corrupted by \mathcal{A} . In the real protocol, $[\omega_z]_i \neq \rho_0 \cdot \alpha_i + \sum_{j \in [m]} \rho_j \cdot ([t_{m_{z_j}}]_i - m_{z_j}^* \cdot [\Delta]_i)$. \mathcal{A} can introduce an error ϵ in the $[\omega_z]_i$ term such that $\sum_{i=1}^n [\omega_z]_i + \epsilon = 0$. In that case, in the real protocol, \mathcal{A} cheats successfully. In the simulated world, the simulator aborts. However, \mathcal{A} can find such an ϵ with probability $\frac{1}{|\mathbb{F}|}$. Therefore, \mathcal{A} can successfully cheat with probability at most $\frac{3}{|\mathbb{F}|}$ which is negligible.

Thus the real world view is indistinguishable from the simulated view.

Simulator \mathcal{S}_{HP}

Semi-Honest Let HP be semi-honest. Thus all the computational parties are honest. Let \mathcal{S}_{HP} be the simulator. Since HP does not have any input, \mathcal{S}_{HP} works as follows on input y where y is the output of the functionality.

- **Preprocessing:** \mathcal{S}_{HP} acts on behalf of the parties in \mathcal{P} and receives the values from HP.
- **Input:** \mathcal{S}_{HP} samples a random value \tilde{q}_v for an input wire on behalf of an input provider. Then it receives the same \tilde{q}_v from HP.
- **Addition and Multiplication:** For addition and multiplication gates, HP does not receive any messages, thus nothing to

simulate.

- **Verification:** \mathcal{S}_{HP} samples a random sharing of zero, $[\tilde{\omega}]_i$ and sends it to HP.

Figure 14: Simulator \mathcal{S}_{HP} for Π_{mpc} for a semi-honest HP

Case II: Fig. 14 describes the simulation steps when the HP is semi-honest. In this case, all parties in \mathcal{P} are honest.

Since HP does not have any input to the protocol, the simulated view of the preprocessing step of the protocol is identical to the real view of the preprocessing phase.

For an input gate, \mathcal{S}_{HP} samples and sends a randomly sampled \tilde{q}_v . In the real protocol, HP receives $q_v = m_v + r$, where r is uniformly sampled. Therefore, q_v is also uniformly distributed. Thus, the distribution of q_v and \tilde{q}_v are identical. Since HP does not receive any messages for addition and multiplication gates, nothing to simulate.

In the verification step, \mathcal{S}_{HP} samples a random sharing of zero $[\tilde{\omega}]_i$. In the real protocol, P_i sends $[\omega]_i$ such that $\sum_{i=1}^n [\omega]_i = 0$, since all the parties follow the protocol. Also, $[\omega]_i = \rho_0 \cdot \alpha_i + \sum_{j=1}^m \rho_j [\omega_{z_j}]_i$, where $\{\alpha_i\}_{i \in [1, n]}$ is a random sharing of zero. Therefore $[\omega]_i$ is uniformly distributed such that $\sum_{i=1}^n [\omega]_i = 0$. Therefore, $[\omega]_i$ and $[\tilde{\omega}]_i$ are identically distributed.