

Secure Computation with Low Communication from Cross-checking

S. Dov Gordon¹, Samuel Ranellucci², and Xiao Wang³

¹ George Mason University

`gordon@gmu.edu`

² Unbound Tech

`samuel_ran@hotmail.com`

³ University of Maryland

`wangxiao@cs.umd.edu`

Abstract. We construct new four-party protocols for secure computation that are secure against a single malicious corruption. Our protocols can perform computations over a binary ring, and require sending just 1.5 ring elements per party, per gate. In the special case of Boolean circuits, this amounts to sending 1.5 bits per party, per gate. One of our protocols is robust, yet requires almost no additional communication. Our key technique can be viewed as a variant of the “dual execution” approach, but, because we rely on four parties instead of two, we can avoid any leakage, achieving the standard notion of security.

1 Introduction

As secure multi-party computation (MPC) is transitioning to practice, one setting that has motivated multiple deployments is that of outsourced computation, in which hundreds of thousands, or millions of users secret share their input among some small number of computational servers. In this setting, the datasets can be extremely large, while the number of computing parties is small. The use of secure computation in such settings is often viewed as a safeguard that helps to reduce risk and liability. While companies and government agencies are increasingly choosing to deploy this safeguard, it is a security / performance tradeoff that many are not yet willing to make.

One important notion related to the security of an MPC protocol is the choice of adversarial threshold: a higher threshold means that the protocol can tolerate more corrupted parties. However, requiring a higher threshold usually results in feasibility and efficiency obstacles. For example, the earliest results in the field demonstrated key distinctions between $t \geq n/2$, $t < n/2$, and $t < n/3$ corruptions [32,3,22,5], including whether fairness could be guaranteed ($t < n/2$), whether a broadcast channel is needed ($t > n/3$), and whether cryptographic assumptions are necessary ($t > n/3$). More recently, when $t > n/2$, there are results showing how to reduce the bandwidth to just a constant number of field elements per party, per gate [18,14,13]. In contrast, when $t \geq n/2$, our best

protocols require expensive preprocessing, with communication cost that grows quadratically in n .

In this work, we develop a new protocol in the honest majority setting, tailored to the case where $n = 4$. Our protocol is secure against a single malicious corruption, consistent with the requirement that $t < n/2$. Focusing on this domain, we are able to construct extremely efficient protocols.

Looking at concrete costs, the most efficient secure two-party computation protocol (in terms of communication) requires roughly 290 bytes of communication per party per gate [33,31]. If we are willing to relax the setting by assuming that a malicious adversary can only corrupt one out of three parties, then we can further reduce the cost to 7 bits per party per gate [1]. Our protocol further reduces the cost significantly: our four-party protocol requires only 1.5-bits of communication per party. Furthermore, the results just cited for the two-party and three-party settings are for 40-bit statistical security, and their costs per gate increase for higher statistical security. Our protocol has no dependence on a statistical security parameter, and has only an additive $O(\kappa)$ term (where κ is a computational security parameter).

We also note that we can achieve 1-bit communication per party in the six-party setting. For these previous works as well as the protocol in this paper, all computation can be hardware accelerated and thus communication complexity is the most suitable indicator of real performance.

Contributions. We now summarize our contributions. Our main result is summarized in the theorem below. The construction and proof of security appear in Sections 3 and 4. An additional improvement appears in Section 5.

Theorem 1. *In the four party setting, it is possible to construct a protocol for securely computing a circuit of size $|C|$ over a finite field F whose total communication complexity is $6|C| \log |F| + O(\kappa)$. In particular, for a Boolean circuit, this amounts to 1.5 bits per player, per gate.*

Binary Rings. An interesting result of our work is that we can securely evaluate an arithmetic function over binary rings, such as $(\mathbb{Z}_{2^{32}}, +, *)$, where $(+, *)$ denotes modular addition and multiplication. Note that most MPC protocols do not work over rings that are not fields. In particular, MAC-based protocols based on SPDZ [17] do not work over $\mathbb{Z}_{2^{32}}$, as the multiplicative inverse is necessary for constructing linear MAC schemes. The security of our protocol only relies on additive maskings, so we do not need a multiplicative inverse. The correctness of our protocol when computing over a binary ring follows from the distributivity property of rings. The distributivity property allows us to use beaver triples. A similar observation, in the semi-honest setting, was recently made by Mohassel and Zhang [30].

Robustness. We construct a robust variant of our protocol, guaranteeing that the honest parties always receive correct output. The cost of adding robustness is free if all the players act honestly, and only requires an additional $O(\kappa \log |C| \log |\mathbb{F}|)$ overhead when a player misbehaves (Section 6).

1.1 Technical Overview

From a high-level view, the construction of our protocol starts with a semi-honest protocol, π_1 , for two-party computation in the preprocessing model. We would like two participants in the protocol to execute π_1 . There are two main tasks towards our final goal:

1. Generating the preprocessing data for π_1 with malicious security.
2. Strengthening the security of π_1 in the online phase from semi-honest to malicious security.

Our solutions to these challenges rely heavily on the fact that we work in the four-party setting with only one corruption. In order to generate maliciously secure preprocessing, we ask the other two parties to locally emulate the preprocessing ideal functionality, both using the same randomness. To ensure that the computation of the preprocessing is done correctly, each of the parties executing π_1 verifies that he was sent two identical copies of the preprocessing.

The second challenge is trickier. Existing work that compiles semi-honest security to malicious security are not suitable for our use. The techniques can be broadly described as follows: 1) Using generic zero-knowledge proof, which is impractical for most cases; 2) Using certain forms of MACs on each party's share to ensure honest behavior. This approach has been made practical, but it requires preprocessing data of size (at least) $\Omega(\rho)$ bits per gate, to achieve $2^{-\rho}$ statistical security. 3) In the honest majority setting, one can use Shamir secret sharing, but our π_1 is a two-party protocol, where one can be malicious. Instead, our approach is based on a technique called “dual-execution” [28,24], which is known to have one-bit leakage in general. However, we show that in the four-party setting, by performing a special cross-checking protocol at the end, we are able to eliminate the leakage without any penalty to the performance. Details follow below.

Dual execution without leakage. In order to accommodate dual execution, we require that π_1 has certain special properties. Intuitively, the outcome of π_1 should leave both parties with “masked wire values” for all wires in the circuit, together with a secret sharing of the masks. This property can be satisfied by many protocols, e.g. the modified Beaver triple protocol [2] that we use in this paper, as well as the semi-honest version of TinyTable [15].

Now we are in the setting, where, say, P_1 and P_2 have generated the preprocessing, and each hold the full set of wire masks, namely λ^1 . P_3 and P_4 have executed π_1 , and recovered masked values, namely m^1 . Our dual execution is done by letting P_1 and P_3 switch roles with P_2 and P_4 . As a result, P_1 and P_2 will obtain m^2 , while P_3 and P_4 will obtain λ^2 in the second execution. Conceptually, our cross-checking compares, for all wire values in the circuit, whether

$$\lambda^1 + m^2 = \lambda^2 + m^1.$$

Note that the above holds if both executions are honest, since both sides of the equation are equal to the true wire values, masked by both masks (λ^1 and λ^2). For details of the protocol, see Section 3.

Readers that are familiar with the dual execution paradigm in the two-party setting, from garbled circuits, might wonder how we remove the bit of leakage. There are two key insights here. First, when using garbled circuits, it seems difficult to check the consistency of internal wires, whereas the masked wires of the form just described allows us to easily check the consistency of *all* wires in the two evaluations. This eliminates the possibility of input inconsistency, and also prevents the adversary from flipping a wire value to see if it has any impact on the output. Second, in a garbled circuit implementation, the adversary can fix the output of a particular gate arbitrarily, creating a “selective failure attack”: the change goes undetected if the output he chooses is consistent with the true output on that wire, and would otherwise cause an abort. With these masked wire evaluations, the adversary cannot fix a wire value arbitrarily; he is limited to adding some value to the wire, changing it in all cases, and always causing an abort. In particular, then, whether he is caught cheating no longer depends on any private value. By exploiting the structure of masks and masked values, checking for inconsistencies requires only $O(\kappa)$ bits of communication.

Reducing communication. The protocol described until this point is already extremely efficient, but we further reduce the communication in several interesting ways. In the preprocessing, we do this in a fairly straightforward way, using PRG seeds and hash functions to compress the material. In the cross checking, recall that we need the parties to verify, twice, whether $\lambda^1 + \mathbf{m}^2 = \lambda^2 + \mathbf{m}^1$, where these values have size $|C|$. (They verify twice because each member of one evaluation compares with one member of the other evaluation.) A naive way here would be to twice compare the hash of these values, but this is in fact insecure. If an adversary changes a value on one of the wires in his evaluation, as we have already noted, he will always be caught, because his partner will compare the hash of his modified masked wire values with an honest party from the other evaluation. However, the adversary can still learn sensitive information from the result of his *own* comparison with a member of the other evaluation. Instead, we can use any honest-majority, four-party protocol for comparing these two hash values. The circuit for this comparison has only $O(\kappa)$ gates, so this introduces very little overhead. Nevertheless, in Section 5 we show how to bootstrap this comparison, removing the reliance on other protocols.

Related work. Maliciously secure protocols, tailored for the three-party setting, have been studied in many works. Choi et al. [8] studied the dishonest majority setting based on garbled circuits. Araki et al. [1], Mohassel et al. [29], Furukawa et al. [19] studied the honest majority setting. However, we are not aware of any MPC protocol tailored for the four-party setting.

Other protocols that work in the four-party setting include honest majority protocols [4,10,11,12,26,16] and dishonest majority protocols [25,6,17,31,27,34]. These protocols can be used for MPC with more parties, but when applied in the four-party setting, their concrete performances are worse than our protocol. In particular, Damgård et al. [16] designed an efficient protocol in the honest majority setting. Their techniques appear to be different from ours and can be

extended to more parties; however, our protocol is much more efficient in terms of the total cost than theirs in the four-party setting.

2 Preliminaries

In this paper, we mainly consider arithmetic circuits C with addition gates and multiplication gates. Each gate in the circuit is represented as (a, b, c, T) , where $T \in \{+, \times\}$ is the operation; a and b are the input wire indices; and c is the output wire index.

We denote the set of wires as \mathcal{W} , the set of input wires as $\mathcal{W}_{\text{input}}$, the set of output wires of all addition gates as $\mathcal{W}_{\text{plus}}$, the set of output wires of all multiplication gates as $\mathcal{W}_{\text{mult}}$.

Masked evaluation. One important concept that we use in the paper is masked evaluation. Intuitively, every wire w in the circuit, including each input and output wire, is associated with a random mask, namely λ_w . The masked evaluation procedure works in a way such that for each gate two parties, holding masked input and some helper information, are able to obtain the masked output. All parties hold only secret shares of λ_w , namely $\langle \lambda_w \rangle$, therefore obtaining masked wire values does not reveal any information. We will use m_w to denote the masked wire value on wire w . That is, $m_w = \lambda_w + x$, assuming that the underlying wire value on wire w is x .

Secure evaluation-masking two-party protocol. A secure two-party protocol for computing circuit C is an evaluation-masking scheme if (1) the protocol uses preprocessing, (2) the preprocessing assigns to the circuit C a masking λ (3) the players evaluate the gates of the circuit layer by layer; if a gate g is in layer L , then the evaluation of L allows both players to learn the masked values for the given layer, (4) if an adversary starts deviating from the protocol, the adversary should not learn any information about the computation unless the output is revealed. (5) any misbehavior from the adversary for a given wire is equivalent to him adding a fixed value to the wire that can be computed from his misbehavior. This type of attack is described as an additive attack in the work of [20]. They showed that certain MPC protocols have this property.

In this paper, we build upon a variant of Beaver’s scheme [2] which is an evaluation-masking scheme. The main modification of Beaver is that the players will hold for each wire, secret shares of masks and both players will learn the sum of the mask and the actual underlying value. We denote the sum of a mask and a value as either a masking or a masked value.

Committing encryption. A public-key encryption scheme is committing if the ciphertexts serve as commitments.

- **Completeness** A person who encrypts a message m resulting in a ciphertext c needs to be able to prove that c is indeed an encryption of m .

- **Soundness** If the player who generated c can prove that c is an encryption of m then $\text{dec}(sk, c) = m$.
- **Verifiability** Given the public-key, it is easy to determine if a ciphertext is valid.

ElGammal encryption is committing due to the fact that (1) every element of the ciphertext space is a valid ciphertext, (2) it is easy to check that an element is in the ciphertext space, and (3) a player can unambiguously prove that a ciphertext maps to a particular plaintext by simply providing the randomness that was used to encrypt a message m .

Theorem 2. *(Informal) From any secure evaluation-masking two-party protocol π_1 , secure against a semi-honest adversary, we can construct a protocol π_2 for four parties that is secure against a malicious adversary corrupting at most one player.*

Security definitions. We use the standard security definition for stand-alone security, as defined, for example, in Goldreich’s textbook [21]. In Section 4, we claim *security with abort*, in which the ideal-world adversary receives output first, and then dictates whether the functionality should provide output to the honest parties. In our protocol of Section 6, we achieve security with *robustness*, or guaranteed output delivery, as defined by Goldwasser and Lindell [23]. Under this definition, the honest parties always receive output, regardless of how the adversary behaves.⁴

As is standard, we prove security in a modular way. We start by proving security of our preprocessing protocol, and we assume we have a secure protocol for comparing two strings. We then prove security of our main protocol in a *hybrid model* in which the participants in the real-world protocol are assumed to have access to the ideal functionalities achieved by these protocols. The seminal work of Canetti proves that this suffices for achieving standard, stand-alone security; the functionalities used in the hybrid world can be instantiated using any protocols that securely realize them [7].

3 Our Main Construction

A quick summary of our idea is that we run two executions of a two-party, semi-honest protocol in the preprocessing model, and verify consistency between these two executions through a strategy that we call cross-checking. We start by partitioning the players into two evaluation groups with two players in each group. Each group prepares preprocessing for the other. They leverage the fact that there is at most one corruption to verify that the preprocessing was done correctly. Then, each group evaluates the circuit using that preprocessing. As the

⁴ Of course, if the adversary refuses to participate, we cannot hope to include his input in the computation. In this case, some default value can be used, or his input can be excluded entirely if the computation is well defined over 3 inputs.

The main protocol takes input from all 4 parties, and outputs the evaluation of C on those inputs. It makes use of the 3 components: \mathcal{F}_{pre} , π_{eval} , and π_{cross} .

Pre-processing

1. The four parties make two calls to \mathcal{F}_{pre} (Fig. 2). In the first call, P_1 and P_2 receive the output $\langle \lambda^2 \rangle, \langle \gamma^2 \rangle$ for E_1 and E_2 respectively, while P_3 and P_4 receive λ^2 , the output of D_1, D_2 . In the second call, they reverse their roles, with P_1 and P_2 receiving λ^1 , the output of D_1 and D_2 , and P_3 and P_4 receive $\langle \lambda^1 \rangle, \langle \gamma^1 \rangle$, the output of E_1 and E_2 .

Evaluation

1. The four parties run two instances of π_{eval} (Fig. 4). In the first instance, the players P_1 and P_2 take the role of evaluators E_1 and E_2 using $\langle \lambda^2 \rangle, \langle \gamma^2 \rangle$. Let \mathbf{m}^1 denote the resulting masked wire values. In the second instance P_3 and P_4 take the role of evaluators E_1 and E_2 using $\langle \lambda^1 \rangle, \langle \gamma^1 \rangle$. Let \mathbf{m}^2 denote the resulting masked evaluation.

Cross Checking

1. The four parties run π_{cross} (Fig. 5) where P_1, P_2 each input \mathbf{m}^1, λ^2 while P_3, P_4 both input \mathbf{m}^2, λ^1 .
2. If π_{cross} outputs 0, then abort.
3. We define λ_{out}^1 to be output masks for the first evaluation.
4. We define $\mathbf{m}_{\text{out}}^1$ to be the masked output wires for the first evaluation.
5. Player P_1, P_2 broadcast $\mathbf{m}_{\text{out}}^1$, if their broadcasts disagree then all players abort.
6. Player P_3, P_4 broadcast λ_{out}^1 , if their broadcasts disagree then all players abort.
7. Players compute the output by using $\mathbf{m}_{\text{out}}^1$ and λ_{out}^1 .

Fig. 1. Main protocol in the hybrid model

outcome of the evaluation, each party holds masked wire values for all wires in the circuit. Finally, the two groups check the consistency of the two evaluations using their masked wire values and masks. Since one of the evaluations is guaranteed to be correct, any cheating will be caught in this step. Below we provide the details of each of these steps as well as why it is secure. A formal description of the protocol appears in Fig. 1, and in the other figures referenced from there.

preprocessing. Recall that we partition four parties into two equal-sized groups. We first let one group create preprocessing material, and distribute the preprocessing to the other group. This procedure is then repeated with the roles reversed; we describe it only for one group. We will often refer to the group that is performing the preprocessing step as D_1 and D_2 , and to the group that uses the preprocessing in the evaluation phase as E_1 and E_2 , recognizing that one party plays the role of (say) D_1 in one execution while playing E_1 in the other execution. An ideal functionality for the preprocessing appears in Fig. 2.

To generate the preprocessing material, D_2 chooses a random string and sends it to D_1 . They then each use this randomness to locally generate preprocessing, choosing mask values for every wire in the circuit as follows. They select a random field element for every wire $w \in \mathcal{W}_{\text{input}} \cup \mathcal{W}_{\text{mult}}$ (that is, for every input wire, and every wire that is the output of a multiplication gate). We refer

This 4-party, randomized functionality is called by two distributors and two evaluators. No parties contribute any input. The functionality generates a vector of random masks as output for the distributors, and a secret-sharing of these masks for the evaluators.

Input: None.

Computation

1. Sample \mathbf{seed}_1 and \mathbf{seed}_2 uniformly at random. If the adversary corrupts D_2 , allow him to specify the seeds.
2. For each wire $w \in \mathcal{W}_{\text{input}} \cup \mathcal{W}_{\text{mult}}$:
 - (a) $A_{1,w} \leftarrow \mathbf{G}(\mathbf{seed}_1)$,
 - (b) $A_{2,w} \leftarrow \mathbf{G}(\mathbf{seed}_2)$
 - (c) $\lambda_w \leftarrow A_{1,w} + A_{2,w}$.
3. For each addition gate $(a, b, c, +)$: compute $\lambda_c \leftarrow \lambda_a + \lambda_b$.
4. For each multiplication gate (a, b, c, \times)
 - (a) $\gamma_c \leftarrow \lambda_a \cdot \lambda_b$
 - (b) $\Gamma_{1,c} \leftarrow \mathbf{G}(\mathbf{seed}_1)$
 - (c) $\Gamma_{2,c} = \gamma_c + \Gamma_{1,c}$

Output

1. Output $\langle \gamma \rangle, \langle \lambda \rangle$ to E_1, E_2 (by sending \mathbf{seed}_1 to E_1 and $(\mathbf{seed}_2, \{\Gamma_{2,w}\}_{w \in \mathcal{W}_{\text{mult}}})$ to E_2)
2. Output both \mathbf{seed}_1 and \mathbf{seed}_2 to both D_1, D_2

Malicious party: A malicious D_2 can choose the randomness.

Fig. 2. \mathcal{F}_{pre} : Ideal functionality for preprocessing

to these mask values as λ^1 , and the ones generated by the other 2 parties, in the second preprocessing execution, are denoted by λ^2 . For the output wire of addition gate $(a, b, c, +)$, suppose the input wires a and b have already been assigned mask values λ_a and λ_b . Then the output wire of the gate is assigned the mask value $\lambda_a + \lambda_b$. Note that all circuit wires now have well defined masks. For each multiplication gate (a, b, c, \times) , the two parties additionally compute $\gamma_c = \lambda_a \cdot \lambda_b$. We let $\gamma^1 = \{\gamma_c\}_{c \in \mathcal{W}_{\text{mult}}}$. D_1 and D_2 use their shared random string to construct secret sharings $\lambda^1 = A_1 + A_2$ and $\gamma^1 = \Gamma_1 + \Gamma_2$. That is, they create two identical copies of the secret sharing. They both send A_1 and Γ_1 to E_1 , and they both send A_2 and Γ_2 to E_2 . E_1 and E_2 each verify the equality of the two values he received before proceeding to the evaluation phase. Note that after agreeing on the random string at the beginning of the procedure described above, D_1 and D_2 require no further communication with each other. Because one of the parties must be honest, the equality checks performed by E_1 and E_2 suffice to catch any malicious behavior. Note that this idea shares some similarity with the one by Mohassel et al. [29] in the three-party setting based on garbled circuit.

We do not present the preprocessing protocol in quite the way that was just described. Instead, an optimized variant with reduced communication complexity is presented in Fig. 3. First, instead of choosing and sending random strings of length $O(|C|)$, the two parties choose two short seeds for a PRG: we

Two distributors D_1, D_2 want to generate preprocessing for players E_1, E_2 .

Creation

1. D_2 chooses two random seeds, \mathbf{seed}_1 and \mathbf{seed}_2 , and sends them to D_1 .
2. For each wire $w \in \mathcal{W}_{\text{input}} \cup \mathcal{W}_{\text{mult}}$:
 - (a) $A_{1,w} \leftarrow \mathbf{G}(\mathbf{seed}_1)$,
 - (b) $A_{2,w} \leftarrow \mathbf{G}(\mathbf{seed}_2)$
 - (c) $\lambda_w \leftarrow A_{1,w} + A_{2,w}$.
3. For each addition gate $(a, b, c, +)$: compute $\lambda_c \leftarrow \lambda_a + \lambda_b$.
4. For each multiplication gate (a, b, c, \times)
 - (a) $\gamma_c \leftarrow \lambda_a \cdot \lambda_b$
 - (b) $\Gamma_{1,c} \leftarrow \mathbf{G}(\mathbf{seed}_1)$
 - (c) $\Gamma_{2,c} = \gamma_c + \Gamma_{1,c}$

Distribution

1. D_1 sends \mathbf{seed}_1 to E_1 and $(\mathbf{seed}_2, \Gamma_2)$ to E_2 .
2. D_2 sends \mathbf{seed}_1 to E_1 and $\mathbf{H}(\mathbf{seed}_2 || \Gamma_2)$ to E_2 .
3. D_1 and D_2 output λ

E_1 **Reconstruction**

1. Receive \mathbf{seed}_1 from D_1 and D_2 and check they are the same. If not, abort.
2. $\{A_{1,w}\}_{w \in \mathcal{W}_{\text{input}} \cup \mathcal{W}_{\text{mult}}}, \{\Gamma_{1,w}\}_{w \in \mathcal{W}_{\text{mult}}} \leftarrow \mathbf{G}(\mathbf{seed}_1)$
3. Output $(\{A_{1,w}\}_{w \in \mathcal{W}_{\text{input}} \cup \mathcal{W}_{\text{mult}}}, \{\Gamma_{1,w}\}_{w \in \mathcal{W}_{\text{mult}}})$.

E_2 **Reconstruction**

1. Receive $(\mathbf{seed}_2, \Gamma_2)$ from D_1 and $\mathbf{H}(\mathbf{seed}_2 || \Gamma_2)$ from D_2 and check they are consistent. If not, abort.
2. $\{A_{2,w}\}_{w \in \mathcal{W}_{\text{input}} \cup \mathcal{W}_{\text{mult}}} \leftarrow \mathbf{G}(\mathbf{seed}_2)$
3. Output $(\{A_{2,w}\}_{w \in \mathcal{W}_{\text{input}} \cup \mathcal{W}_{\text{mult}}}, \{\Gamma_{2,w}\}_{w \in \mathcal{W}_{\text{mult}}})$.

Notation

1. For each wire $w \in \mathcal{W}_{\text{input}} \cup \mathcal{W}_{\text{mult}}$: $\langle \lambda_w \rangle \leftarrow (A_{1,w}, A_{2,w})$
2. For each multiplication gate (a, b, c, \times) : $\langle \gamma_c \rangle \leftarrow (\Gamma_{1,c}, \Gamma_{2,c})$

Fig. 3. Distributed Preprocessing of masked beaver triples

let $A_1 = \mathbf{G}(\mathbf{seed}_1)$, and $A_2 = \mathbf{G}(\mathbf{seed}_2)$. As before, $\lambda^1 = A_1 + A_2$. Since the value of γ^1 depends on λ^1 , we cannot do the same thing there, but we can generate the shares Γ_1 from the same \mathbf{seed}_1 , and then fix Γ_2 appropriately, using $O(|\mathcal{W}_{\text{mult}}|)$ bits. This reduces the communication cost for each of the parties from $(2|\mathbb{F}| + 1) \cdot |\mathcal{W}_{\text{mult}}|$ to $2\kappa + |\mathbb{F}| \cdot |\mathcal{W}_{\text{mult}}|$. Recall that D_1 and D_2 send identical copies of these values to an evaluator; we further reduce the communication by having one party send only a single hash of the preprocessing, which suffices for allowing each evaluator to verify the consistency of what he has received. Finally, note that this last optimization causes the communication costs to become unbalanced. Although we do not present it, note that we can re-balance the cost by having one party send the first half of Γ_2 together with a hash of the second half, while the other party sends the second half of Γ_2 together with a hash of the first half.

Evaluation. After receiving and verifying the consistency of the preprocessing, E_1 and E_2 proceed to perform a mask-evaluation of the circuit, layer by layer.

There are two evaluators E_1, E_2 who want to evaluate a circuit C using preprocessing provided by distributors D_1, D_2 . Each of the four players is assigned a set of input wires corresponding to his input to C .

Input

1. For each input wire w , one party holds input x_w .
2. For each input wire w , D_1 and D_2 hold λ_w .
3. For each $w \in \mathcal{W}$, E_1 and E_2 hold $\langle \lambda_w \rangle$.
4. For each multiplication gate (a, b, c, \times) , E_1 and E_2 hold $\langle \gamma_c \rangle \leftarrow \langle \lambda_a \cdot \lambda_b \rangle$.

Sharing Input Values

For each input wire w belonging to E_1 with value x_w :

1. D_1, D_2 both send λ_w to E_1 . E_1 aborts if they are different.
2. E_1 sends $\lambda_w + x_w$ to E_2

For each input wire w belonging to D_1 with value x_w :

1. D_1 sends $\mathbf{m}_w \leftarrow x_w + \lambda_w$ to E_1 and E_2 .
2. E_1 and E_2 verify that they each received the same value and abort if it is not the case.

The input of E_2 is processed similarly to the input of E_1 .

The input of D_2 is processed similarly to the input of D_1 .

Evaluation

For each gate (a, b, c, T) following topological order:

1. if $T = +$
 - (a) $\mathbf{m}_c \leftarrow \mathbf{m}_a + \mathbf{m}_b$.
2. if $T = \times$
 - (a) $\langle \mathbf{m}_c \rangle \leftarrow \mathbf{m}_a \cdot \mathbf{m}_b - \mathbf{m}_a \cdot \langle \lambda_b \rangle - \mathbf{m}_b \cdot \langle \lambda_a \rangle + \langle \lambda_c \rangle + \langle \lambda_a \cdot \lambda_b \rangle$
 - (b) $\mathbf{m}_c \leftarrow \text{open}(\langle \mathbf{m}_c \rangle)$

Fig. 4. π_{eval} : Two-party Masked Evaluation

To begin, they first need masked input values for every input wire; these are of the form $\mathbf{m}_w \leftarrow \lambda_w + x_w$. For an input wire w held by $E \in \{E_1, E_2\}$, D_1 and D_2 send λ_w to E . E verifies that they each sent the same value: if not, he aborts. Otherwise, he computes $\lambda_w + x_w$ and sends it to the other evaluator. For input wire w belonging to $D \in \{D_1, D_2\}$, D sends $\lambda_w + x_w$ to E_1 and E_2 . The evaluators compare values and abort if they don't agree.

For every gate $(a, b, c, +)$, E_1 and E_2 both locally compute $\mathbf{m}_c = \mathbf{m}_a + \mathbf{m}_b$. For every gate (a, b, c, \times) , they locally compute $\langle \mathbf{m}_c \rangle \leftarrow \mathbf{m}_a \cdot \mathbf{m}_b - \mathbf{m}_a \cdot \langle \lambda_b \rangle - \mathbf{m}_b \cdot \langle \lambda_a \rangle + \langle \lambda_c \rangle + \langle \lambda_a \cdot \lambda_b \rangle$. (Recall, they can compute the last term using $\langle \gamma_c \rangle$.) They then compute $\mathbf{m}_c \leftarrow \text{open}(\langle \mathbf{m}_c \rangle)$ by exchanging their shares of \mathbf{m}_c . At the conclusion of evaluation phase, one set of evaluators holds \mathbf{m}^1 , which is the set of masked values of all wires in the circuit, and the other group of parties hold \mathbf{m}^2 after their evaluation phase.

Cross-checking. Note that during the evaluation phase, a malicious evaluator can modify the value on any $w \in \mathcal{W}_{\text{mult}}$ simply by changing his share of \mathbf{m}_w before reconstructing the value. Therefore, before either group recovers output from their computation, they first compare their masking with the masking of

Input

1. P_1 has input $(\mathbf{m}_1^1, \lambda_1^2)$.
2. P_2 has input $(\mathbf{m}_2^1, \lambda_2^2)$.
3. P_3 has input $(\mathbf{m}_3^2, \lambda_3^1)$.
4. P_4 has input $(\mathbf{m}_4^2, \lambda_4^1)$.

Computation

1. P_1 computes $h_1 = H(\mathbf{m}_1^1 + \lambda_1^2)$,
 P_2 computes $h_2 = H(\mathbf{m}_2^1 + \lambda_2^2)$,
 P_3 computes $h_3 = H(\mathbf{m}_3^2 + \lambda_3^1)$,
 P_4 computes $h_4 = H(\mathbf{m}_4^2 + \lambda_4^1)$.
2. All players send h_i to \mathcal{F}_{eq} , which outputs 1 if and only if $h_1 = h_3$ and $h_2 = h_4$.

Fig. 5. π_{cross} : Cross Checking

the other evaluation. Of course, they cannot reveal the values on any wires while doing this check. Instead, for wire w that carries value x , each set of evaluators uses the masking from their evaluation, together with the masks that they generated for the other group during preprocessing, to compute

$$x + \lambda_w^1 + \lambda_w^2 = \mathbf{m}_w^1 + \lambda_w^2 = \mathbf{m}_w^2 + \lambda_w^1.$$

They then compare these “doubly masked” values for consistency.

As in the case of preprocessing, we use a hash function where possible, in order to reduce the communication cost. Each party begins by computing a hash of the doubly masked wire values described above; for P_i , we denote this hash by h_i . The four parties then call an ideal functionality, \mathcal{F}_{eq} , which takes input h_i , and outputs 1 if and only if $h_1 = h_3$, and $h_2 = h_4$.

Taking P_1 as example, he obtains \mathbf{m}_1^1 during evaluation and λ_1^2 when acting as a D . He will then compute $h_1 = H(\mathbf{m}_1^1 + \lambda_1^2)$. For the other three parties, it is defined similarly as follows: superscripts denote the index of the masked evaluation and subscripts denote the identity of the party.

$$h_2 = H(\mathbf{m}_2^1 + \lambda_2^2), \quad h_3 = H(\mathbf{m}_3^2 + \lambda_3^1), \quad h_4 = H(\mathbf{m}_4^2 + \lambda_4^1)$$

To see why this suffices for providing security, suppose P_1 changes some masking during evaluation, effectively changing a wire value for him and P_2 . In this case, the doubly masked evaluations of P_2 and P_4 are inconsistent, and \mathcal{F}_{eq} will return 0; intuitively, comparing these hash values is equivalent to checking the masked values wire by wire.

3.1 Concrete Performance

Here we briefly discuss the concrete performance of our protocol against the most related state-of-the-art protocol by Araki et al. [1]. As mentioned previously, our protocol requires 1.5 bits of communication per gate per party, a $4.5\times$ improvement over their protocol. Let’s see if the same applies to the computation cost.

Note that in the protocol by Araki et al., the heaviest part of the computation is random shuffling, due to the use of the random bucketing technique in their paper. The rest are AES and hash computation, which can be hardware accelerated or very fast. Compared to their protocol, our protocol is much simpler and more efficient in terms of computation cost. The bulk of our computation is in the evaluation phase, where we do not need any random shuffling. For each 128 AND gates, each party only needs 6 calls to fix-key AES to implement the PRG, and roughly one call to a hash function. Araki et al. have a higher computational cost than we do, because of their random shuffle; since they are able to fill a 10Gbps LAN, our protocol will certainly have no problem filling the same pipe. We believe the computation cost will not be the bottleneck for any reasonable hardware configuration.

Our protocol for the cross checking appears in Fig. 5. It is in a hybrid world where the parties have access to a functionality, \mathcal{F}_{eq} . We note that this functionality can be realized using any secure four party computation. The circuit needed to realize this functionality is small: it only performs two equality computations on strings of length $O(\kappa)$. Nevertheless, in Section 5, we also demonstrate how we can bootstrap this functionality, communicating just a small constant number of bits, and using almost no computation.

3.2 Multiplayer Extensions

Achieving one bit of communication using six parties. We note that if we use six players, we can maintain of the overhead of 6 bits communicated in total, thereby requiring each player to communicate just one bit per wire (on average). The idea of the six-party computation protocol is fairly straightforward given the four-party protocol. Two people agree on randomness for the preprocessing, and then each communicates the preprocessing material to two of the remaining four players. Those four parties now carry out two identical evaluations, in parallel, and cross check them with one another at the end. The communication overhead is still six bits per gate, but it is now divided among all six players.

Efficient Multiparty Protocols via Log-Depth Threshold Formulae. We note that our result can be used to construct MPC protocols tolerating up to a third of players being corrupted via the result of [9]. At a high level, their protocol employs player emulation and works by recursively composing a protocol for a small number of parties with itself via a log-depth threshold formulae.

4 Security Proof

4.1 Proof of Security for Preprocessing

Lemma 1. *The protocol in Fig. 3 for distributed preprocessing securely realizes the functionality of Fig. 2, with abort.*

Proof. Due to symmetry, we only prove the lemma for the following two cases: 1) D_1 is corrupt and 2) E_1 is corrupt.

Corrupted D_1 . We will first describe our simulator \mathcal{S} .

1. \mathcal{S} queries \mathcal{F}_{pre} and obtains $\text{seed}_1, \text{seed}_2$. If the \mathcal{A} chooses to input randomness, use \mathcal{A} 's choice.
2. \mathcal{S} acts as honest D_2, E_1 and E_2 for the rest of the protocol using the seeds obtained above. If an honest E_1 or E_2 would abort, \mathcal{S} sends **abort** to \mathcal{F}_{pre} .

Note that none of the parties in the protocol have input. Therefore the indistinguishability of the ideal-world protocol and the real-world protocol is immediate, given the observation that the protocol aborts in the real world protocol if and only if it aborts in the ideal world protocol.

Corrupted E_1 . Note that E_1 performs only local computation after receiving messages from other parties. The simulator queries \mathcal{F}_{pre} and receives the seeds. He then simulates honest D_1 and D_2 , sending seed_1 on their behalf. If E_1 aborts, the simulator will send **abort** to \mathcal{F}_{pre} and aborts. Indistinguishability from the real-world protocol is immediate.

4.2 Proof of Security of the Main Protocol

Theorem 3. *Assuming H is drawn from a family of collision resistant hash functions, our main protocol, in Fig. 1, securely realizes $\mathcal{F}_{4\text{pc}}$ in the $(\mathcal{F}_{\text{pre}}, \mathcal{F}_{\text{eq}})$ -hybrid model.*

Proof. In the following, we will prove the security of our main protocol assuming that P_1 is corrupted by \mathcal{A} . The simulator is as follows:

1. \mathcal{S} honestly simulates the execution of \mathcal{F}_{pre} . He sends P_1 his resulting output, and records the simulated mask values: λ^1 , which will mask the wire values in the evaluation of P_3 and P_4 , and λ^2 , which will mask the wire values in the evaluation of P_1 and P_2 .
2. \mathcal{S} simulates the masking of input values 0 from P_2, P_3 and P_4 for use in P_1 's evaluation with P_2 , using mask values from λ^2 . He receives three maskings of P_1 's input: one for each of P_3 and P_4 for use in their evaluation, using mask values from λ^1 , and one using values from λ^2 , sent to P_2 for his own evaluation with P_1 . If the values sent to P_3 and P_4 are not equal, \mathcal{S} sends abort to $\mathcal{F}_{4\text{pc}}$ and terminates the simulation. Otherwise, \mathcal{S} extracts the input sent to P_2 , and the one sent to P_3 and P_4 , using his knowledge of the masks; he notes if P_1 misbehaves by using inconsistent values in the two evaluations.
3. \mathcal{S} acts honestly as P_2, P_3 and P_4 in both executions of the masked evaluations. \mathcal{S} obtains m^1 by interacting honestly with P_1 on behalf of P_2 for the remainder of their evaluation. He obtains m^2 by simulating (internally) the remainder of the evaluation of P_3 and P_4 .
4. \mathcal{S} collects P_1 's input to \mathcal{F}_{eq} , and calculates the inputs of P_2, P_3 , and P_4 to \mathcal{F}_{eq} according to the honest execution using input values of 0. \mathcal{S} executes the code of \mathcal{F}_{eq} locally on these 4 values, and outputs the resulting value.

Comment: In Section 5, we describe a more efficient, interactive protocol, π_{vcc} , which replaces the use of \mathcal{F}_{eq} . To simulate our protocol when using π_{vcc} ,

we would proceed as follows, in place of the previous step. If \mathcal{S} noted that that P_1 misbehaved during evaluation, or when sending his masked input, then \mathcal{S} runs π_{vcc} , simulating the messages of P_2 and P_4 when using different (random) inputs from one another. Otherwise, he runs π_{vcc} as though P_2 , P_3 and P_4 all use input $m^1 \oplus \lambda^2$. If π_{vcc} outputs 0, \mathcal{S} sends **abort** to $\mathcal{F}_{4\text{pc}}$.

5. \mathcal{S} uses the input extracted in Step 2 and sends it to $\mathcal{F}_{4\text{pc}}$. He receives y and computes $\lambda^* = m_{\text{out}}^1 + y$. \mathcal{S} acts as P_3 and P_4 sending λ^* to \mathcal{A} .

Now we will show that the joint distribution of the output from \mathcal{A} and honest parties in the ideal world are indistinguishable from these in the real world protocol.

1. **Hybrid₁**: Same as the hybrid protocol, with \mathcal{S} playing the role of honest players, using their true input. That is, in this hybrid, we give the input of the honest players to the simulator. (The resulting distribution is equivalent to that of the real world execution.)
2. **Hybrid₂**: Same as **Hybrid₁**, with messages simulated based on the true inputs of the honest parties. However, now \mathcal{S} uses m^1 and λ^1 to compute $x_1 = m_{\text{in}}^1 + \lambda_{\text{in}}^1$. \mathcal{S} sends x_1 to $\mathcal{F}_{4\text{pc}}$, which returns y . In step 5, \mathcal{S} acts as P_3 and P_4 and broadcasts $\lambda^* = m_{\text{out}}^1 + y$.
3. **Hybrid₃**: Same as the **Hybrid₂** except that \mathcal{S} uses input values of 0 for all honest parties, instead of their true input.

It is fairly easy to see that **Hybrid₁** is indistinguishable from **Hybrid₂**, as long as the function H is drawn from a family of collision resistant hash functions. Up until the simulation of \mathcal{F}_{eq} , the view of the adversary is identically generated in the two hybrids. If H is collision resistant, then the output of \mathcal{F}_{eq} is 1, if and only if the adversary behaves honestly in his execution. When he acts honestly, the output generated in **Hybrid₂** by the ideal functionality is the same as the output computed by \mathcal{S} in **Hybrid₁**. The distribution characterizing the adversarial view in **Hybrid₃** is identical to that of **Hybrid₂**, by the one-time-pad security of the random masking. It is easy to verify that the joint distribution described by the adversarial view and the honest output is identical in these two hybrids as well.

5 Cross Check from Veto

In this section, we will demonstrate how to construct an efficient cross checking protocol based on a functionality for 4-party, logical OR, \mathcal{F}_{or} . We sometimes call this a *veto functionality*, as the parties use the OR to “veto” the execution, by submitting a value of 1 (veto). The cross checking protocol from Section 3 required a 4-party computation of \mathcal{F}_{eq} , which compared 2 pairs of strings, each κ bits long. The improved cross checking protocol based on veto requires each party to compare two hashes locally, and then input a single bit to the veto functionality. While the cost of either of these protocols is small compared to the evaluation phase, the simplicity of the protocol here makes it hard to pass up. We also describe how to bootstrap \mathcal{F}_{or} , using a variant of the protocol from

The protocol assumes access to an ideal functionality, \mathcal{F}_{or} , for computing the logical OR of 4 input bits, each provided by one of the parties.

Input

1. P_1 has input $d_1 = m_1^1 + \lambda_1^2$.
2. P_2 has input $d_2 = m_2^1 + \lambda_2^2$.
3. P_3 has input $d_3 = m_3^2 + \lambda_3^1$.
4. P_4 has input $d_4 = m_4^2 + \lambda_4^1$.

Checking

1. P_1 samples a random seed and sends it to P_3 .
2. P_1 (resp. P_3) send $H(d_1\|\text{seed})$ (resp. $H(d_3\|\text{seed})$) to P_2 and P_4 .
3. P_2 (resp. P_4) determines if it received the same value from P_1 and P_3 . If it did, it will provide 0 to the \mathcal{F}_{or} functionality, and otherwise it will provide 1.
4. Repeat the previous instructions with the variable exchanged as follows: P_1 is switched with P_2 , and P_3 is switched with P_4 , d_1 is switched with d_3 , and d_2 is switched with d_4 .
5. Players call the \mathcal{F}_{or} functionality with the input that they were instructed to use in step 3.

Fig. 6. Cross check protocol from veto

Section 3, and requiring just 6 bytes of communication per party. Perhaps one of the nicest features of this bootstrapping, from a practical standpoint, is that it allows us to avoid any dependence on other MPC implementations.

Naive implementation of cross checking. A naive way of implementing cross checking is to have the two verifiers exchange their doubly masked evaluations, and compare them for inconsistencies. Unfortunately, this approach fails because the adversary can modify the values carried on any of the wires in his own evaluation, and determine precisely how the change impacted the evaluation of the circuit by subtracting his doubly masked evaluation from the other. The differences between these two doubly masked evaluations reveals the differences in the values carried on each wire in the two evaluations of the circuit.

Achieving secure and efficient cross checking. Our main observation for simplifying the cross check protocol is that, in the attack just described, P_1 will always cause the verification run by P_2 and P_4 to fail. This is because the evaluation of P_2 was also modified on wire w , but he will not modify λ_w^1 the way P_1 did. If the output of the equality test between P_1 and P_3 were hidden from P_1 , shown only to P_2 and P_4 , and, symmetrically, if P_1 only saw the result of their verification (which he already knows), then we can remove the bit of leakage. Specifically, each P_i learns a single bit, b_i , indicating whether the *other* verifying set passed the equality test. The four parties then run a secure protocol that computes the logical OR of these 4 bits. They can do this using any existing 4-party protocol.

One verification group reveals the equality of their masked evaluations to the other verification group as follows. (1) They agree on a random seed, (2)

they hash it together with their doubly masked evaluation, and (3) they send the hash output to the players of the other verification group. The players in the other verification group can compute equality by simply checking that the hashes they receive are the same.

Note that \mathcal{F}_{or} is a constant size circuit, and it likely does not matter which four party secure computation we use to realize it. Still, it is interesting to note that we can actually bootstrap this computation with another variant of our own protocol. In the protocol just previously described, letting $d_{i,w}$ denote the doubly masked value held by P_i for wire w , the parties effectively compute $\bigvee_{w \in \mathcal{W}} (d_{1,w} \neq d_{3,w}) \vee \bigvee_{w \in \mathcal{W}} (d_{2,w} \neq d_{4,w})$, where the hash value received by P_1 and P_3 (resp. P_2 and P_4) reveals the first (resp. second) disjunction of size $|\mathcal{W}|$ to P_1 and P_3 (resp. P_2 and P_4). The disjunction in the middle is where we use \mathcal{F}_{or} . Following the same discussion above, the reader can verify that it is also secure to compute $\bigvee_{w \in \mathcal{W}} ((d_{1,w} \neq d_{3,w}) \vee (d_{2,w} \neq d_{4,w}))$. This can be achieved by having the four parties check the equality of gates in topological order by immediately exchanging the results of every equality check, rather than “batching them” with a hash function at the end of the evaluation. Removing the hash function in this way increases the communication to $O(|C|)$, so we would not prefer to use this as our cross-checking protocol. However, since \mathcal{F}_{or} only has three gates, it is inefficient to bootstrap \mathcal{F}_{or} using a hash function.

Security of Veto Cross Check

(Sketch.) Assuming H is a non-programmable random oracle, our main protocol is secure if we replace the cross checking in the main protocol with the cross checking described in this section.

If the adversary acted maliciously during the masked evaluation, then it is clear that the verification group that does not contain the corrupt player (i.e. the honest verification group) will have inconsistent evaluations. As a result, the simulator can run the cross checking on behalf of the honest players as though the player in the honest verification group had inconsistent evaluations. In this case, the honest player in the same validation group as the corrupt player will always provide a veto. As a result, the simulator can safely always provide a simulated output of veto from \mathcal{F}_{or} , sends **abort** to $\mathcal{F}_{4\text{pc}}$, and the result is indistinguishable from a real execution.

If instead the corrupt player only misbehaves in the cross checking, the only possible deviation is to send the wrong hash value. In this case, the simulator can compute whether the corrupt player misbehaved by analyzing the hash value that he sent, together with the seed. The simulator knows that both players in the honest verification group will veto. As a result, the simulator can simply provide a simulated output of veto from \mathcal{F}_{or} , submit **abort** to $\mathcal{F}_{4\text{pc}}$, and the result is indistinguishable from a real execution.

Finally, if the adversary never deviates from the protocol, the simulator accepts the adversary’s input to \mathcal{F}_{or} and sends it back to him as the output of \mathcal{F}_{or} . If this value is a veto, the simulator sends **abort** to $\mathcal{F}_{4\text{pc}}$, and otherwise, he

Two distributors D_1, D_2 want to generate preprocessing for players E_1, E_2 . We assume D_1 holds a key pair for a public key, committing encryption scheme, and that both hold key pairs for a digital signature scheme. The 3 public keys are known by all parties. We let (pk, sk) denote the encryption/decryption keys of D_1 , and (vk_i, sk_i) denote the verifying/signing keys of D_i .

Protocol

1. D_2 chooses $(seed_1, seed_2, r_{com})$ at random and broadcasts $enc(pk, seed_1 || seed_2 || r_{com})$. The other players then verify that D_2 broadcasted a valid ciphertext (by using the verifiability property of the committing encryption scheme). If D_2 did not broadcast a ciphertext, or if the ciphertext was not valid, then the other 3 parties halt this protocol and instead run a semi-honest 3-party protocol amongst themselves.
2. D_1 recovers $(seed_1, seed_2, r_{com})$.
3. Each D_i computes the preprocessing that was described in Fig. 3.
4. Each D_i computes $commit(\{\lambda_w\}_{w \in \mathcal{W}_{output}}; r_{com})$. He includes these commitments in the preprocessing material.
5. Each D_i signs the preprocessing material: $\sigma_{1,i} = sign(sk_i; seed_1)$ and $\sigma_{2,i} = sign(sk_i; seed_2 || r_{com})$. He sends $(seed_1, \sigma_{1,i})$ to E_1 , and $(seed_2, r_{com}, \sigma_{2,i})$ to E_2 .
6. E_j receives $(m_{j,1}, \sigma_{j,1})$ and $(m_{j,2}, \sigma_{j,2})$. He checks whether $verify(pk_1, \sigma_{j,1}) = verify(pk_2, \sigma_{j,2}) = 1$, and whether $m_{j,1} = m_{j,2}$.
 - If one of the signatures does not verify, E_j continues the protocol using only the preprocessing material that was validly signed.
 - If both signatures verify, but $m_{j,1} \neq m_{j,2}$,
 - E_j broadcasts the two signed messages.
 - D_2 broadcasts $(seed_1, seed_2)$, together with the encryption randomness used in Step 1. All honest parties can now determine whether D_1 or D_2 misbehaved. They eliminate the guilty party and execute a 3-party, semi-honest protocol.
 - E_j outputs $m_{j,1}$.

Fig. 7. Robust Preprocessing

submits the adversary’s input to \mathcal{F}_{4pc} , and simulates the opening of the output just as in Section 4.

The view in the real and ideal world are indistinguishable since (1) the simulator can always determine if there is a veto or not based on the behavior of the adversary and (2) the random oracles hides inputs from the other verification group.

6 Adding Robustness

We can make our protocol robust against a single cheater. We note that it is quite simple to strengthen our original protocol so that it is *fair*. If the malicious party aborts before anyone sends the output wire masks, then nothing is learned, and all parties can safely abort. If the adversary aborts after learning the output masks, his partner can still reveal the output for the other two evaluators. The

only necessary modification is to prevent the malicious distributor from changing his output masks, revealing output values that conflict with what his partner reveals. This is easily handled by having all parties commit to their output masks prior to the evaluation: if the two distributors use the same randomness in their commitments, the evaluators can verify that they have both committed to the same mask value.

The main challenge in achieving robustness is that we cannot simply abort when we detect improper behavior, even if the output has not been revealed yet. Instead, we have to ensure that all honest parties correctly identify a misbehaving party, or at least a pair of parties that contains the adversary. To facilitate this, we make several adjustments. First, we modify the preprocessing protocol so that it either allows everyone to identify the adversary, or it ensures that both evaluators receive good preprocessing material. The robust preprocessing appears in Fig. 7. We then modify the input sharing to make it robust; the input sharing in Section 3 would trigger an abort if any party used different inputs in the two executions, but it would not allow the others to determine who cheated. After receiving the preprocessing material and the masked inputs, the evaluators continue the evaluation protocol from Section 3 until each party has a masking of the circuit. They then perform a robust variant of the cross checking protocol. In this variant, the parties cross check gate by gate, and if they ever find an inconsistency, they run a sub-routine to identify a pair of parties that contains the adversary.⁵ Input sharing, evaluation, and robust cross checking are fully described in Fig. 8. We give a detailed overview of these changes below.

We will employ a broadcast channel throughout our protocol. We note that unconditionally-secure broadcast is possible in the four player model with one corruption due to the fact that less than $n/3$ players are corrupt.

Robust preprocessing: To make the preprocessing robust, one of the two distributors, D_2 , starts by committing to the randomness that will be used in the preprocessing. This commitment is constructed by broadcasting a committing encryption under the public key of D_1 . The randomness used in the preprocessing is denoted by $(\mathbf{seed}_1, \mathbf{seed}_2, r_{\text{com}})$: \mathbf{seed}_1 and \mathbf{seed}_2 are used to create masks, just as in Section 3. r_{com} is used to construct a commitment to the output masks, which is then included in the preprocessing output.

After generating the preprocessing material, D_1 and D_2 each sign a copy of the output before sending it to E_1 and E_2 . If they send conflicting values to E_1 , the signatures allow E_1 to convince the other honest parties that one of D_1 or D_2 is malicious. The honest one of the two can now be exonerated: D_2 broadcasts the randomness used to encrypt the preprocessing randomness. E_1 broadcasts their view, and the honest parties can check the validity of the messages sent

⁵ To reduce communication of the robust cross checking, we can iteratively apply our cross check protocol from Section 5, performing a binary search on the masked circuit layers until we find the problematic layer. We then repeat that, performing a binary search within the problematic layer to find the problematic gate. This would yield a worst-case communication cost of $O(\kappa \log |C|)$. For simplicity, we describe the protocol as operating gate per gate.

Input phase

1. For each input wire w
 - (a) Suppose $P_1 = E_1$ is the player who provides input x_w for wire w (we can generalize this to the other parties)
 - (b) E_1 awaits the mask λ_w^2 from D_1, D_2 , as well as signatures on λ_w^2 .
 - i. If E_1 receives a value for λ_w^2 without a signature from D_i then he ignores the mask that D_i sent him.
 - ii. Otherwise, he received inconsistent masks, E_1 broadcasts the signed masks (thus identifying which evaluation group contains a cheater). If the players receive two different masks with valid signatures, they run the protocol, using only the masked evaluation of E_1, E_2 .
 - (c) E_1 broadcasts $\mathbf{m}_w = x_w + \lambda_w^1 + \lambda_w^2$.
 - (d) E_1, E_2 set $\mathbf{m}_w^1 = \mathbf{m}_w - \lambda_w^1$ while D_1, D_2 set $\mathbf{m}_w^2 = \mathbf{m}_w - \lambda_w^2$.

Evaluation

Each evaluation group, using their own masked evaluation, as well as the share of the masks they received from distributors do the following:

For each gate (a, b, c, T) following topological order:

- (a) if $T = +$
 - i. $\mathbf{m}_c \leftarrow \mathbf{m}_a + \mathbf{m}_b$.
- (b) if $T = \times$
 - i. $\langle \mathbf{m}_c \rangle \leftarrow \mathbf{m}_a \cdot \mathbf{m}_b - \mathbf{m}_a \cdot \langle \lambda_b \rangle - \mathbf{m}_b \cdot \langle \lambda_a \rangle + \langle \lambda_c \rangle + \langle \lambda_a \cdot \lambda_b \rangle$
 - ii. $\mathbf{m}_c \leftarrow \text{open}(\langle \mathbf{m}_c \rangle)$

Cross Check

For every wire $w \in \mathcal{W}_{\text{mult}} \cup \mathcal{W}_{\text{input}}$, ordered by depth in the circuit.

For each verification group $V \in \mathcal{V}$

1. $(V_1, V_2) \leftarrow V$
2. V_1 send $d \leftarrow \mathbf{m}_w^1 + \lambda_w^1$ to V_2 .
3. V_2 broadcasts (**error**) if $d \neq \mathbf{m}_w^2 + \lambda_w^2$.
4. If a player in V broadcasts (**error**), run $\text{complaint}(w)$.
 - (a) If the complaint phase returns (**corrupt**, P_1, P_2), P_1, P_2 broadcast decommitments to λ_w^1 for each output wire w . P_3, P_4 compute $\mathbf{m}_w^2 - \lambda_w^1$, broadcast the result, and the protocol terminates.
 - (b) If the complaint phase returns (**corrupt**, P_3, P_4), P_3, P_4 broadcast decommitments to λ_w^2 for each output wire w . P_1, P_2 compute $\mathbf{m}_w^1 - \lambda_w^2$, broadcast the result, and the protocol terminates.
 - (c) If the complaint phase returns (**corrupt**, verifier), set $\mathcal{V} \leftarrow \mathcal{V} \setminus V$ and restart the protocol with the updated \mathcal{V} .

Output

For each output wire w ,

1. Players P_1, P_2 broadcast the decommitment to λ_w^1 .
2. Players P_3, P_4 broadcast the decommitment to λ_w^2 .
3. P_1 and P_2 broadcast $\mathbf{m}_w^1 - \lambda_w^2$. Denote these values by $(\text{out}_1, \text{out}_2)$.
4. P_3 and P_4 broadcast $\mathbf{m}_w^2 - \lambda_w^1$. Denote these values by $(\text{out}_3, \text{out}_4)$.
5. All parties output $\text{Majority}(\text{out}_1, \text{out}_2, \text{out}_3, \text{out}_4)$.

Fig. 8. Robust Evaluation

Complaint Subprotocol

The *complaint* subprotocol is initiated for output wire w for a multiplication gate g_w when a player C has complained that the two masked evaluations were inconsistent on that wire. We denote by $V \in \mathcal{V}$ the verification group that contains C . This subprotocol allows the parties to identify either an evaluation group or a verification group that contains a cheater.

Complaint

1. $\mathcal{E}_1 \leftarrow \{P_1, P_2\}$, $\mathcal{E}_2 \leftarrow \{P_3, P_4\}$, $\mathcal{D}_1 \leftarrow \{P_3, P_4\}$, $\mathcal{D}_2 \leftarrow \{P_1, P_2\}$
2. Players run the validation functionality for wire w using \mathcal{E}_1 as the evaluators and \mathcal{D}_1 as the distributors.
3. Players run the validation functionality for wire w using \mathcal{E}_2 as the evaluators and \mathcal{D}_2 as the distributors.
4. If any of the calls to the validation functionality results in the functionality returning $(\text{corrupt}, P_i, P_j)$, return the same. Otherwise, return $(\text{corrupt}, \text{verifier})$.

Validation

E_1, E_2, D_1 and D_2 want to verify that the masked evaluation of E_1, E_2 was done correctly for the gate g_w with output wire w .

Input

1. E_1, E_2 each input their masked evaluation for the 3 wires of g_w : $(\mathbf{m}_a, \mathbf{m}_b, \mathbf{m}_c)$.
2. D_1, D_2 each input the masks that they generated g_w : $(\lambda_a, \lambda_b, \lambda_c)$.

Functionality

1. If E_1 and E_2 provided distinct inputs, return $(\text{corrupt}, E_1, E_2)$ and halt.
2. If D_1 and D_2 provided distinct inputs, return $(\text{corrupt}, D_1, D_2)$ and halt.
3. If $(\mathbf{m}_a - \lambda_a) \cdot (\mathbf{m}_b - \lambda_b) + \lambda_c = \mathbf{m}_c$, then output (valid) . Otherwise output $(\text{corrupt}, E_1, E_2)$.

Fig. 9. Complaint

by D_1 and D_2 . After removing the malicious party, the remaining three parties can run a semi-honest protocol in which one party supplies the preprocessing, the other two perform the evaluation, and no checking needs to be performed.

One other case of note deserves mention: suppose E_1 receives nothing⁶ from, say, D_1 . In this case, because there is no signature, E_1 cannot prove that D_1 or D_2 is malicious: it is equally possible that E_1 is himself malicious, and that he made the problem up. In this case, though, E_1 does need to persuade anybody. Because E_1 knows that D_1 is malicious, E_1 can simply continue the protocol using the preprocessing he received from D_2 .

Robust input sharing: Let P_1 and P_2 perform distribution for P_3, P_4 , and vice versa. Recall that Section 3, P_1 shares input x_w on wire w with P_3, P_4 by using the mask λ_w^1 that he and P_2 generated together. He shares his input with P_1 , for their own evaluation, by using λ_w^2 , which he receives from P_3, P_4 . As written, nothing prevents him from sharing inconsistent values among the

⁶ equivalently, something that is not validly signed

parties, and nothing prevents those parties from pretending he did so. To fix this, we first require P_3, P_4 to each sign λ_w^2 , which allows P_1 to broadcast a proof of inconsistency when necessary. Then, P_1 signs and broadcasts his doubly masked input: $m_w = x_w + \lambda_w^1 + \lambda_w^2$. P_2 computes $m_w - \lambda_w^1$ for use in his evaluation with P_2 . P_3, P_4 each compute $m_w - m_w^2$ for use in their evaluation.

Robust cross checking: Instead of cross checking the hashes of the full circuit maskings, the parties instead cross check gate by gate, starting at the input layer, and proceeding topologically through the circuit. This protocol begins with a pass over the circuit, one layer at a time, with the parties comparing their doubly masked values to locate the first gate at which the two evaluations depart from one another. Consider the case where P_3 decides that the two masked evaluations of some gate are inconsistent, and initiates a complaint. This can be due to one of the following cases:

1. The masked evaluation performed by P_1 and P_2 is invalid.
2. The masked evaluation performed by P_3 and P_4 is invalid.
3. Both evaluations were executed correctly, but either P_1 modified his input to cross-checking (i.e. his reported masked evaluation), or P_3 complained for no valid reason.

If the honest players know that the first case holds, then the corrupt player is either P_1 or P_2 . They can therefore use the evaluation of P_3 and P_4 to determine their output. By the same argument, if the players know that the second case holds, they can all safely use the evaluation of P_1 and P_2 to produce the output. Finally, if the players know they are in the third case, they know that the malicious party is either P_1 or P_3 . In this case, they do not dismiss either evaluation, but they continue the cross checking using only between P_2 and P_4 ; since P_2 and P_4 are honest, their cross-checking suffices for ensuring a valid computation.

When someone detects an inconsistency in the cross checking of a gate, the parties execute a *complaint* subprotocol (See Fig. 9) to determine which of the above cases hold. In this subprotocol, the parties use an ideal functionality, which can later be bootstrapped generically using any MPC with identifiable abort. We stress that the circuit implementing this functionality is small: it only needs to be executed on a single gate, and it used at most twice in a computation. The functionality is called once for each of the two evaluations. In each instance, the two evaluators provide their masked input and masked output for the gate, while the two distributors provide the masks that they created for the gate. If the evaluators do not provide the same masked values then the functionality indicates that the evaluation set contains the cheater. If the distributors do not provide the same masks then the functionality indicates that the cheater is in the distribution set. Otherwise, the functionality uses the masked wire values and the mask values to check whether the gate evaluation was performed correctly. If the masked evaluation was invalid, the ideal functionality indicates that the evaluation set contains the cheater. Finally, if no error is detected, then the functionality indicates this, and the parties conclude that either the party that raised the alarm is malicious, or his partner in the cross-checking is malicious (case 3 above).

6.1 Robust Evaluation Simulator

Theorem 4. *If the robust evaluation protocol is instantiated using a CCA-Secure public-key committing encryption scheme, and a EU-CMA signature scheme, then it securely realizes $\mathcal{F}_{4\text{pc}}$ in the random oracle model. In addition, the protocol is robust.*

Simulator for robust preprocessing when D_2 is corrupt.

1. Wait that D_2 broadcasts the ciphertext c . If the ciphertext is not valid (which can be efficiently checked by the verifiability property of committing encryption) then the simulator submits a default input value to $\mathcal{F}_{4\text{pc}}$ on behalf of the adversary, and terminates. (This corresponds to the honest parties removing the adversary from the computation, upon agreeing that he is malicious.)
2. Recover $(\text{seed}_1 \parallel \text{seed}_2 \parallel r_{\text{com}}) \leftarrow \text{Dec}(sk, c)$.
3. The simulator computes the preprocessing and broadcasts $\text{commit}(\{\lambda_w\}_{w \in \mathcal{W}_{\text{output}}}; r_{\text{com}})$ on behalf of D_1 .
4. Simulator awaits that D_2 sends the preprocessing material and signatures on the preprocessing material to each player. Then,
 - (a) For each evaluator, if D_2 sent an invalid signature to the given evaluator, the simulator ignores what D_2 sent.
 - (b) Otherwise, if D_2 sent invalid preprocessing to either evaluator, then simulate the broadcast from the given evaluator of the the signed preprocessing and determining that D_2 misbehaved. The simulator notes that D_2 was identified as a cheater.

Simulator for robust preprocessing when D_1 is corrupt. Same as the simulation for D_2 except that the simulator broadcasts the encryption of the randomness to D_1 .

Simulator for robust preprocessing when an evaluator is corrupt. The simulator chooses randomness and simulates the three honest players. If an evaluator sends a message claiming he received inconsistent preprocessing, but the signed messages he forwards do not substantiate his claim, the simulator sends a default input to $\mathcal{F}_{4\text{pc}}$ and terminates. (Technically, we did not describe in our protocol that the other parties remove the evaluator when he does this, because we felt it would unnecessarily complicate the protocol description.)

Indistinguishability of robust preprocessing.

1. In the case where the distributor is corrupt, we claim the view in the real and ideal worlds are indistinguishable. If the distributor deviates from the protocol, it is either ignored (if it does not send a signature with the preprocessing it shares), or it is eliminated from the computation (if it sends bad preprocessing with a valid signature). The committing property of the

encryption scheme guarantees that he gets caught if he signs and sends a wrong value.

2. The only message sent by an evaluator is (possibly) to complain about inconsistent preprocessing. If the evaluator is corrupt, then in both the ideal and real world, the complaint would be ignored (due to the unforgeability of the underlying signature scheme).

Claim. Let $x = (x_1, x_2, x_3, x_4)$. Let $\text{view}[\text{Input}]_{\pi, \mathcal{A}} = \{\text{view}[\text{Input}]_{\pi, \mathcal{A}}\}_{x, \kappa}$ denote the random variable describing real-world view of \mathcal{A} for some fixed set of inputs, and some fixed security parameter. There exists a simulator \mathcal{S} such that $\text{view}[\text{Input}]_{\mathcal{F}, \mathcal{S}} \equiv$ denote the random variable resulting from the ideal world simulation described in the $\mathcal{F}_{4\text{pc}}$ (described below). There in the input sharing phase, and

1. If the corrupt player is providing input as an evaluator,
 - (a) The simulator provides the signed masks from the other distributors.
 - (b) The simulator awaits that the corrupt player broadcasts a double masking \mathbf{m}_w . The simulator then computes the input of the corrupt player from the masks that were produced in the preprocessing and the double masking that the corrupt player sent.
2. If the corrupt player is a distributor, and the input wire belongs to an evaluator,
 - (a) The simulator awaits that the corrupt player sends out a mask to the evaluator. If the mask is signed with the corrupt player's signature, and is not the value produced in the preprocessing, then the simulator produces a broadcast of the conflicting, signed masks. The simulator provides the default value to $\mathcal{F}_{4\text{pc}}$ on behalf of the corrupt player and terminates.
3. If the corrupt player is an evaluator, and the input wire belongs to the other evaluator, the simulator broadcasts the doubly masked input.

Indistinguishability of input phase. We argue that since the view until the end of the preprocessing phase in the ideal world is indistinguishable from the view until the end of the preprocessing phase in the real world, then the views are also indistinguishable up through the end of the input phase. In the real and ideal world, when the distributor is corrupt, any deviation would either be ignored, or would result in the dealer being caught and eliminated from the computation. If the evaluator is corrupt, and he broadcasts an invalid complaint, he is eliminated due to the unforgeability of the underlying signature scheme.

Simulator for evaluation. The simulator of the evaluation step follows the same steps as the simulator for the masked evaluation in the main protocol. In particular, the simulator stores if the corrupt player misbehaved during his evaluation. We argue that since the view until the end of the input phase in the ideal world is indistinguishable from the view until the end of the input phase in the real world, then the views are also indistinguishable up through the end

of the evaluation phase.. This holds from the fact that our main protocol (in particular the masked evaluation part) is secure.

Simulator for cross check. For every multiplication wire $w \in \mathcal{W}_{\text{mult}}$,

1. If the corrupt evaluator had previously sent a wrong value in the evaluation of wire w ,
 - (a) The simulator broadcasts (**error**) on behalf of the verifiers that are not in the same verification group as the corrupt player. (He might also do so with the player that is in the same verification group as him.)
 - (b) The simulator receives $(\lambda_a, \lambda_b, \lambda_c)$ and $(\mathbf{m}_a, \mathbf{m}_b, \mathbf{m}_c)$ from the adversary, intended for the first and second calls to the validation functionality, respectively (and without loss of generality). If $(\lambda_a, \lambda_b, \lambda_c)$ are inconsistent with the values simulated during preprocessing, the simulator implicates the adversary (and his partner) when simulating the output of the first call to the validation functionality. In either case he implicates the adversary (and his partner) in the simulated output of the second call to the validation functionality.

The simulator then runs the protocol on behalf of the honest players using the honest evaluation group's masked evaluation.

2. Otherwise:
 - (a) if simulating V_1 , the simulator checks to see if the adversary sends a wrong doubly masked value to his partner: $\mathbf{m}_w^1 + \lambda_w^1$.
 - (b) if simulating V_2 , the simulator checks to see if the adversary broadcasts (**error**).

The simulator receives $(\lambda_a, \lambda_b, \lambda_c)$ and $(\mathbf{m}_a, \mathbf{m}_b, \mathbf{m}_c)$ from the adversary, intended for the first and second calls to the validation functionality, respectively (and without loss of generality). If $(\lambda_a, \lambda_b, \lambda_c)$ are inconsistent with the values simulated during preprocessing, the simulator implicates the adversary (and his partner) when simulating the output of the first call to the validation functionality. If $(\mathbf{m}_a, \mathbf{m}_b, \mathbf{m}_c)$ are inconsistent with simulated masked values of the evaluation phase, the simulator implicates the adversary (and his partner) when simulating the output of the second call to the validation functionality. If he is not implicated in either instance, then any future messages he might send during cross checking are ignored.

Indistinguishability of cross check. We argue that since the view until the end of the evaluation phase in the ideal world is indistinguishable from the view until the end of the evaluation phase in the real world, then the views are also indistinguishable up through the end of the cross check. If the corrupt player's evaluation group is deemed corrupt, then the protocol in the real world would dictate that the corrupt player no longer receive messages during the cross check phase. Therefore, it is clear that after the elimination has taken place, the views in the real and ideal world are indistinguishable.

We now claim that the validation function eliminates the adversary's evaluation set in the real world, if and only if the simulator implicates the adversary's

evaluation set in the ideal world. Note that the simulator can detect if the adversary has modified any wire in the evaluation, as well as whether his input to the validation function is inconsistent with his partner’s input. The reader can verify by inspection that the claim holds. Since the complaint phase consists of just two calls to the validation functionality, it follows that the adversary’s view in the complaint phase is identically distributed in the two worlds. By the previous note, after this point, the cross check in the real and ideal worlds would be indistinguishable.

To complete the argument that the adversary’s view is correctly simulated through the end of the cross check phase, we argue that, prior to being eliminated, the simulated view in the cross check phase is sampled from the same distribution as his view in the real world. This follows because he only sees doubly masked wire values, which are computationally indistinguishable from uniformly distributed strings (because they are generated using a PRG).

Simulator for output phase. The output phase is the easiest to simulate.

1. First the simulator queries the ideal functionality with the adversary’s input and receives an output.
2. The simulator selects masks for the honest evaluation group so that the sum of the output and the masks of the honest evaluation group is equal to the masked evaluation of the corrupt player. The simulator then “broadcasts” decommitments to the masks of the honest evaluation group.
3. The simulator selects masked evaluation for the honest evaluation group so that the sum of the output and the masks of the corrupt player add up to the masked evaluation. The simulator then “broadcasts” the masked evaluations.

Indistinguishability of output phase. We now argue that the output distribution, conditioned on the adversary’s view, is indistinguishable in the two worlds. We have already argued that the adversary is caught if he ever manipulates his evaluation. The reader can verify that whenever a transcript results in the use of a default adversarial input in the real world, the simulator submits default input in the ideal world. If the adversary never changes the masked values, then the input used in both worlds is the one he committed to in the input sharing phase.

References

1. Araki, T., Barak, A., Furukawa, J., Lichter, T., Lindell, Y., Nof, A., Ohara, K., Watzman, A., Weinstein, O.: Optimized honest-majority MPC for malicious adversaries - breaking the 1 billion-gate per second barrier. In: 2017 IEEE Symposium on Security & Privacy. pp. 843–862 (May 2017)
2. Beaver, D.: Efficient multiparty protocols using circuit randomization. In: Crypto’91. LNCS, vol. 576, pp. 420–432 (Aug 1992)
3. Beaver, D., Micali, S., Rogaway, P.: The round complexity of secure protocols (extended abstract). In: ACM STOC. pp. 503–513 (May 1990)

4. Ben-David, A., Nisan, N., Pinkas, B.: FairplayMP: a system for secure multi-party computation. In: ACM CCS 2008. pp. 257–266 (Oct 2008)
5. Ben-Or, M., Goldwasser, S., Wigderson, A.: Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In: 20th ACM STOC. pp. 1–10 (May 1988)
6. Bendlin, R., Damgård, I., Orlandi, C., Zakarias, S.: Semi-homomorphic encryption and multiparty computation. In: Eurocrypt 2011. LNCS, vol. 6632, pp. 169–188 (May 2011)
7. Canetti, R.: Security and composition of multiparty cryptographic protocols. *Journal of Cryptology* 13(1), 143–202 (2000)
8. Choi, S.G., Katz, J., Malozemoff, A.J., Zikas, V.: Efficient three-party computation from cut-and-choose. In: Crypto 2014, Part II. LNCS, vol. 8617, pp. 513–530 (Aug 2014)
9. Cohen, G., Damgård, I.B., Ishai, Y., Kölker, J., Miltersen, P.B., Raz, R., Rothblum, R.D.: Efficient multiparty protocols via log-depth threshold formulae - (extended abstract). In: Crypto 2013, Part II. LNCS, vol. 8043, pp. 185–202 (Aug 2013)
10. Damgård, I., Geisler, M., Krøigaard, M., Nielsen, J.B.: Asynchronous multiparty computation: Theory and implementation. In: PKC 2009. LNCS, vol. 5443, pp. 160–179 (Mar 2009)
11. Damgård, I., Ishai, Y.: Constant-round multiparty computation using a black-box pseudorandom generator. In: Crypto 2005. LNCS, vol. 3621, pp. 378–394 (Aug 2005)
12. Damgård, I., Ishai, Y.: Scalable secure multiparty computation. In: Crypto 2006. LNCS, vol. 4117, pp. 501–520 (Aug 2006)
13. Damgård, I., Ishai, Y., Krøigaard, M.: Perfectly secure multiparty computation and the computational overhead of cryptography. In: Eurocrypt 2010. LNCS, vol. 6110, pp. 445–465 (May 2010)
14. Damgård, I., Ishai, Y., Krøigaard, M., Nielsen, J.B., Smith, A.: Scalable multiparty computation with nearly optimal work and resilience. In: Crypto 2008. LNCS, vol. 5157, pp. 241–261 (Aug 2008)
15. Damgård, I., Nielsen, J.B., Nielsen, M., Ranellucci, S.: The TinyTable protocol for 2-party secure computation, or: Gate-scrambling revisited. In: Crypto 2017, Part I. pp. 167–187. LNCS (2017)
16. Damgård, I., Orlandi, C., Simkin, M.: Yet another compiler for active security or: Efficient MPC over arbitrary rings. *Cryptology ePrint Archive*, Report 2017/908 (2017), <http://eprint.iacr.org/2017/908>
17. Damgård, I., Pastro, V., Smart, N.P., Zakarias, S.: Multiparty computation from somewhat homomorphic encryption. In: Crypto 2012. LNCS, vol. 7417, pp. 643–662 (Aug 2012)
18. Franklin, M.K., Yung, M.: Communication complexity of secure computation (extended abstract). In: 24th ACM STOC. pp. 699–710 (May 1992)
19. Furukawa, J., Lindell, Y., Nof, A., Weinstein, O.: High-throughput secure three-party computation for malicious adversaries and an honest majority. In: Eurocrypt 2017, Part II. LNCS, vol. 10211, pp. 225–255 (May 2017)
20. Genkin, D., Ishai, Y., Prabhakaran, M., Sahai, A., Tromer, E.: Circuits resilient to additive attacks with applications to secure computation. In: 46th ACM STOC. pp. 495–504 (May / Jun 2014)
21. Goldreich, O.: *Foundations of Cryptography: Basic Applications*, vol. 2. Cambridge University Press, Cambridge, UK (2004)

22. Goldreich, O., Micali, S., Wigderson, A.: How to play any mental game or A completeness theorem for protocols with honest majority. In: 19th ACM STOC. pp. 218–229 (May 1987)
23. Goldwasser, S., Lindell, Y.: Secure multi-party computation without agreement. *Journal of Cryptology* 18(3), 247–287 (Jul 2005)
24. Huang, Y., Katz, J., Evans, D.: Quid-Pro-Quo-tocols: Strengthening semi-honest protocols with dual execution. In: 2012 IEEE Symposium on Security & Privacy. pp. 272–284 (May 2012)
25. Ishai, Y., Prabhakaran, M., Sahai, A.: Founding cryptography on oblivious transfer - efficiently. In: *Crypto 2008*. LNCS, vol. 5157, pp. 572–591 (Aug 2008)
26. Lindell, Y., Nof, A.: A framework for constructing fast MPC over arithmetic circuits with malicious adversaries and an honest-majority. In: *ACM CCS 2017*. pp. 259–276 (2017)
27. Lindell, Y., Pinkas, B., Smart, N.P., Yanai, A.: Efficient constant round multi-party computation combining BMR and SPDZ. In: *Crypto 2015, Part II*. LNCS, vol. 9216, pp. 319–338 (Aug 2015)
28. Mohassel, P., Franklin, M.: Efficiency tradeoffs for malicious two-party computation. In: *PKC 2006*. LNCS, vol. 3958, pp. 458–473 (Apr 2006)
29. Mohassel, P., Rosulek, M., Zhang, Y.: Fast and secure three-party computation: The garbled circuit approach. In: *ACM CCS 2015*. pp. 591–602 (Oct 2015)
30. Mohassel, P., Zhang, Y.: SecureML: A system for scalable privacy-preserving machine learning. In: 2017 IEEE Symposium on Security & Privacy. pp. 19–38 (May 2017)
31. Nielsen, J.B., Nordholt, P.S., Orlandi, C., Burra, S.S.: A new approach to practical active-secure two-party computation. In: *Crypto 2012*. LNCS, vol. 7417, pp. 681–700 (Aug 2012)
32. Rabin, T., Ben-Or, M.: Verifiable secret sharing and multiparty protocols with honest majority (extended abstract). In: 21st ACM STOC. pp. 73–85 (May 1989)
33. Wang, X., Ranellucci, S., Katz, J.: Authenticated garbling and efficient maliciously secure two-party computation. In: *ACM CCS 2017*. pp. 21–37 (2017)
34. Wang, X., Ranellucci, S., Katz, J.: Global-scale secure multiparty computation. In: *ACM CCS 2017*. pp. 39–56 (2017)