



Best of Both Worlds in Secure Computation, with Low Communication Overhead

Daniel Genkin^{1,3}, S. Dov Gordon^{2(✉)}, and Samuel Ranellucci^{2,3}

¹ University of Pennsylvania, Philadelphia, USA
danielg3@cis.upenn.edu

² George Mason University, Arlington, USA
gordon@gmu.edu

³ University of Maryland, College Park, USA
samuel@umd.edu

Abstract. When performing a secure multiparty computation with a few hundred parties, using the best protocols known today, bandwidth constraints are the primary bottleneck. A long line of work demonstrates that n parties can compute a circuit C of depth d while communicating $O(|C| \log |C| + \text{poly}(d, n))$ field elements per party, as long as a majority of parties are honest. However, in the malicious majority setting, a lot less is known. The work of Nielsen and Ranellucci is the first to provide constant-overhead in the communication complexity when a majority of parties are malicious; their result demonstrates feasibility, but is quite complex and impractical.

In this work, we construct a new MPC protocol in the pre-processing model. We introduce a new middle-ground: our protocol has low communication and provides robustness when a majority of parties are honest, and gives security with abort (possibly with higher communication cost) when a majority of players are malicious. Robustness is impossible when a majority of parties are malicious; viewing the increased communication complexity as a form of denial of service, similar to an abort, we view our result as providing the “best of both worlds”.

1 Introduction

After a decade of improvements in the computational cost of secure multiparty computation, we have reached a point where the primary performance bottleneck is the communication complexity, even when computing with only a moderate number of parties. Most constructions require that n participants communicate a total of $O(Cn^2)$ field elements to compute a circuit of size C . The n^2 term stems from point-to-point communication at every multiplication gate, which, at first glance, seems hard to avoid. Amazingly, when a majority of parties are honest, there are several constructions that require communicating only $O(C)$ field elements.¹ Very broadly, these constructions make use of two ideas to

¹ Additionally, they have an additive term that is polynomial in n .

lower communication cost. First, by using a randomly chosen dealer, they can reduce the communication channels from $O(n^2)$ to $O(n)$. This requires care, to ensure that a malicious dealer cannot corrupt the computation. Second, by using “packed secret sharing”, the participants can communicate just one field element to compute $O(n)$ multiplication gates. In a bit more detail, multiple wire values are simultaneously encoded using a single threshold secret sharing scheme: to encode ℓ wire values, w_1, \dots, w_ℓ , a random polynomial p is chosen such that $p(-j) = w_j$. As usual, $p(1), \dots, p(n)$ define the secret shares of the n parties, and, for a degree $t + \ell$ polynomial, all ℓ secrets remain perfectly hidden against t colluding parties. Since $t + \ell < n$, this provides a tradeoff between security and efficiency; as more values are packed into the secret sharing, the number of corruptions that can be tolerated decreases. With a small blowup in the circuit description, these polynomials can be used to compute ℓ multiplication gates at a time, cutting the communication cost by a factor of $\ell = O(n)$ [9].

In the malicious majority setting, a lot less is known about reducing the communication complexity. The recent work of Nielsen and Ranellucci is the first and only protocol with constant communication cost per circuit gate [16]. The result of their work is exciting, as it demonstrates feasibility for the first time. However, as the authors state, their protocol “is solely of theoretical interest”; it has constants that are large and difficult to compute, and, conceptually, it requires parsing a complex composition of player emulations and subprotocols.

In this work, we propose an optimistic approach to communication complexity. Our protocol has constant *expected* communication complexity if a majority of players are *honest*. However, unlike prior work in the honest majority setting, we stress that our protocol also remains secure when a majority of players are malicious, albeit with higher communication complexity. At a high level, the variation in communication complexity stems from the following feature of our approach. We choose a random dealer and hope that they are honest. If the dealer happens to be malicious, he can force a re-start of the protocol, and if $O(n)$ consecutive dealers are malicious, then they can force the communication complexity to blow up. Taking the view that this increased communication cost is simply a form of denial of service, we view our result as providing “the best of both worlds” with respect to denial of service; when a majority of parties are malicious, it is impossible to prevent a denial of service attack, as the adversary can always force an abort. While Nielsen and Ranellucci show that, technically, it is possible to achieve low communication when a majority of players are malicious, the benefit of our relaxation is that it allows us to construct a much simpler protocol, both in concept and in concrete complexity.

The phrase “best of both worlds” has been used before in the MPC literature, referring to the more common notion of denial of service: guaranteed output delivery [8, 13, 15]. With only a few exceptions, protocols for secure multiparty computation are usually designed with a particular corruption threshold in mind. They either provide security with guaranteed output delivery when a majority of parties are honest, but provide no security at all when a majority are malicious, or they provide security with abort when a majority of parties are

corrupt, but allow a denial of service even if only a single party is corrupt. Our protocol provides the best of both worlds in this sense as well, giving security with guaranteed output delivery when the adversary fails to corrupt a majority of parties, and security with abort when a majority are corrupt.

Our construction relies on offline (data independent) preprocessing that, currently, we do not know how to compute with constant overhead (short of using Nielsen and Ranellucci). While we hope this reliance can be removed in future work, we note that there are settings where it might be very reasonable to use such preprocessing. The obvious case is where the parties can afford to send a lot of data prior to the arrival of their inputs, but another setting in which preprocessing is available is where the parties have access to some trusted setup.

Formal Description of Our Result. For privacy threshold t_p , and packing parameter ℓ , our protocol enables n players to compute any arithmetic circuit C , guaranteeing security with abort when fewer than $t < t_p$ players are corrupt. It achieves guaranteed output delivery (aka: robustness, full security) when $t < t_r$, where $t_r = (n - t_p - 2 \cdot \ell) / 2$. In addition, if $t < t_r$ and $\ell \in \Omega(n)$, then for a circuit C of size $|C|$ and depth d , our protocol has expected communication complexity of $O(|C| \log |C| + \text{poly}(n, d))$.

Related Work. Our work follows from two lines of work. The first line focuses on achieving low overhead computation in the majority setting, this includes the work of [2, 14]. The paper of [3] achieved a sublinear overhead in the number of players, but only in the computational setting and with overhead in the security parameter that is not sublinear. The paper of [2] showed how by selecting $\ell \in \Omega(n)$, it was possible to construct a protocol for n parties with communication overhead of $O(|C| \log |C| + \text{poly}(n, d))$ for a circuit C of size $|C|$ and depth d .

The second line of work, [8, 13, 15] focuses on finding MPC protocols with tradeoffs between how many corruptions can be tolerated before privacy is compromised (t_p), and how many corruptions can be tolerated before the robustness guarantee is lost (t_r). Ishai et al. demonstrated that this is possible when there is some slack in the parameters: there exist n -party protocols where, for $t_p + t_r < n$, the protocol maintains security with guaranteed output delivery against t_r malicious players and security with abort against t_p malicious parties [13]. In the same work, they demonstrated that this slackness is inherent, by giving an example of a function that cannot be securely computed with these same guarantees if $t + s = n$.

In parallel to our work, the work of [12] used the assumption that a certain number of parties are honest to improve the efficiency of semi-honest GMW and BMR-style MPC protocols. Other approaches that use preprocessing (such as [4, 6, 7]) require each player to communicate one field element per multiplication since they do not use packing.

1.1 Technical Overview

In this section we present a high level overview of our protocol. We begin by describing a semi-honest version of our protocol, in order to provide insight into how we achieve low communication complexity. (Note that we never give a formal description of this semi-honest version, and it is meant purely for intuition.) Borrowing techniques from [2,3,5], we use a t_p -private packed Shamir secret sharing scheme with packing parameter ℓ . These polynomials have degree $t_p + \ell$, and we will maintain this degree as we compute the circuit.

To compute multiplication gates, our protocol uses a special designated party (called the dealer), and Beaver triples $[\mathbf{a}], [\mathbf{b}], [\mathbf{c}]$, which are secret sharings of values $a, b, c \in \mathbb{F}^\ell$, where a, b are randomly sampled and $c = a \cdot b$ (introduced in [1]). These triples are shared using a t_p -private Shamir packed secret sharing scheme with packing parameter ℓ . The packing parameter ℓ allows players to compute pointwise multiplication on vectors of field elements by having each player compute and send a constant number of field elements to the dealer.

Our protocol evaluates an arithmetic circuit C in topological order from the input to the output gates. Since packed Shamir secret sharing is linear, the players can locally compute on their shares in order to evaluate the addition gates of C . To compute the product $[\mathbf{z}] = [\mathbf{x}] \cdot [\mathbf{y}]$, the players execute the following steps, using a Beaver triple $[\mathbf{a}], [\mathbf{b}], [\mathbf{c}]$. First, the players locally compute shares of $\mathbf{x} - \mathbf{a}$ and $\mathbf{y} - \mathbf{b}$ and send them to the dealer. The dealer reshares $\mathbf{x} - \mathbf{a}$, $\mathbf{y} - \mathbf{b}$ and $(\mathbf{x} - \mathbf{a}) \cdot (\mathbf{y} - \mathbf{b})$ using degree ℓ polynomials. By resharing and packing those values instead of sending them in the clear, we cut down the communication cost by a factor ℓ ; the secret sharing in this step has nothing to do with privacy. The players then compute shares of $\mathbf{w} \leftarrow \mathbf{y} \cdot (\mathbf{x} - \mathbf{a}) + \mathbf{x} \cdot (\mathbf{y} - \mathbf{b}) + r$, where the random mask r is sampled and secret shared during preprocessing, using a degree $t_p + \ell$ polynomial. Since \mathbf{x} and \mathbf{y} are of degree $t_p + \ell$, and $(\mathbf{y} - \mathbf{b})$ and $(\mathbf{x} - \mathbf{a})$ are of degree ℓ , it follows that \mathbf{w} is of degree $t_p + 2\ell$. The players send their shares of \mathbf{w} to the dealer. The dealer re-shares \mathbf{w} using a degree ℓ polynomial, and the players compute $\mathbf{z} = \mathbf{w} - r$. Since r is shared using a degree $t_p + \ell$ polynomial, this results in shares of a degree $t_p + \ell$ polynomial, maintaining the invariant.

The use of a dealer allows each user to send secret shares to one party, instead of n parties, cutting the cost per gate from $O(n^2)$ to $O(n)$. Packed secret sharing further reduces the complexity from $O(n)$ to $O(n/\ell)$. However, this also forces us to increase the degree of the polynomial to $t_p + \ell$, which creates a tradeoff between privacy and efficiency: the closer t_p is to n , the smaller ℓ must be.

Attacks by Malicious Adversaries. The protocol above is only secure against a semi-honest adversary. At a high level, an active adversary, which instructs the players or the dealer to deviate from the protocol specification, can mount two types of attacks.

Additive Attacks. The first class of attacks occur either when a corrupt dealer re-shares the wrong value, or when malicious players send invalid shares to an

honest dealer, thereby causing the dealer to reconstruct and re-share the wrong value. As we describe in our proof sketch in Sect. 4, these attacks are actually instances of additive attacks, in which an adversary can tamper with the evaluation of circuits by adding or subtracting values on individual wires, but cannot impact the computation in any other way. See the full version of this paper for more detail. By running the protocol on an additively secure circuit, obtained from the compiler of Genkin et al. [10], we are able to construct a protocol for MPC that renders such an attack ineffective. At a high level, the compiler of Genkin et al. takes any circuit and transforms it into a new circuit that will output \perp if the adversary applies an additive attack (i.e. tampers with the value of any wire). By showing that any attack on our protocol is equivalent to an additive attack, we can apply the protocol of [10] to make it secure. We note that [10] has a constant overhead.

Divide and Conquer Attacks. The second class of attacks can only be performed by a malicious dealer. At a high level, during the evaluation of multiplication gates, instead of re-sharing values using a degree- ℓ polynomial, the dealer can create two sets of shares, each consistent with a different degree- ℓ polynomial.

More formally, consider the following situation: let n be the number of parties, let M be a set of corrupted parties, and let S_1, S_2 be distinct sets of honest parties (not necessarily disjoint). The adversarial dealer sends shares to S_1 such that the secret recovered from those shares is $x - a$. He sends shares to S_2 such that the secret recovered from those shares is $x - a + 1$. Then, when the players try to compute shares of $(x - a) \cdot y + r$, where r is a random mask, note that both $S_1 \cup M$ and $S_2 \cup M$ give the dealer enough shares to reconstruct the blinded secret: from $S_1 \cup M$, the dealer can recover $(x - a) \cdot y + r$ and from the shares of $S_2 \cup M$, the dealer can recover $(x - a + 1) \cdot y + r$. By subtracting $(x - a + 1) \cdot y + r$ from $(x - a) \cdot y + r$ the malicious dealer can recover y , even though the value of $(x - a) \cdot y$ is supposedly hidden by a random mask.

The Degree-Test Protocol. Dealing with this second type of attack is one of our main technical contributions. In Sect. 3 we present a novel *degree-test* protocol that takes secret shares from the dealer and transmits them to the players if only if the shares of the honest players are consistent with a polynomial of degree- d . This degree test is also efficient, requiring each player to exchange only a constant number of field elements with the dealer. The main idea behind this protocol is as follows. During preprocessing, all parties learn a portion of a secret that is encoded in a degree $n - 1$ polynomial, w . Additionally, they receive shares of a degree $n - d - 1$ polynomial, v , such that $v(0) = 0$. To prove that he shared a degree d polynomial, the dealer collects n shares of z , defined as $z \leftarrow p \cdot v + w$. If p is of appropriate degree, this suffices to learn $w(0)$, revealing the secret value, while if the degree of p is too high, $w(0)$ remains hidden and the dealer fails to prove that he acted honestly.

2 Best of Both Worlds Security

We prove our protocols secure under the ideal-world, real-world paradigm. We define f_C as the ideal functionality that takes an input x from the players and outputs $C(x)$. The functionality f_C^A takes an input x from the players, and a vector A from the adversary. It also evaluates C on x , but it allows an adversary to tamper with the evaluation by adding values on individual wires; the variable A specifies the values that are added to each wire.

Definition 1. Let $t_p \leq n$ be positive integers, let SD denote the statistical distance, and let $0 \leq \epsilon \leq 1$. We say that an n -party protocol π (t_p, ϵ)-securely realizes a functionality \mathcal{F} if for every PPT real-world adversary A which corrupts at most t_p players, there exists a simulator S such that

$$\text{SD}(\text{Real}_{\pi,A}, \text{Ideal}_{\mathcal{F},S}) \leq \epsilon.$$

We naturally extend this definition to protocol in the g -hybrid model by replacing $\text{Real}_{\pi,A}$ above with $\text{Real}_{\pi,A,g}$. In this case we say that π (t_p, ϵ)-securely realizes \mathcal{F} in the g -hybrid model.

Definition 2. Let $n \geq t_p \geq t_r$ be positive integers and let $0 \leq \epsilon \leq 1$. We say that an n -party protocol π (t_r, t_p, ϵ)-robustly realizes (f_C, f_C^A) if it meets the following two conditions.

1. **Security.** If $t_p > t \geq t_r$ then π (t_p, ϵ)-securely realizes f_C^A as per Definition 1. This property does not guarantee that players receive outputs, because the adversary can cause the protocol to abort in the real world.
2. **Robustness.** If $t_r > t$ then π (t_p, ϵ)-securely realizes f_C , and it is guaranteed that the protocol will successfully terminate, with each honest player receiving output. More formally, if less than t_r players are corrupt, the output generated in the real world is the same that is produced by the functionality f_C in the ideal world where (i) each honest player P_i provides input x_i to the functionality, and (ii) the ideal functionality selects a default input for each corrupted player that does not provide an input x_i .

3 Degree Test

Our degree test protocol is an interactive proof between a single prover (dealer) and multiple verifiers (players). The dealer sends a field element to each player, and proves that these elements are consistent with a polynomial p of degree at most d . We construct a proof where the prover can only convince a given verifier with probability $2^{(-4s/n-t_p-\ell)}$. The aim is that at least $n - t_p - \ell$ verifiers will not be convinced by a cheating prover. The protocol proceeds as follows. The preprocessing functionality randomly samples a binary string of size $\frac{4s}{n-t_p-\ell}$, encodes it as $\mathbf{secret} \in \mathbb{F}$, and sends a portion of \mathbf{secret} to each player. After sharing the polynomial p , the players will interact with the dealer in a manner

that allows the dealer to learn this `secret` if and only if p is of degree d or less. The dealer then proves that he learned `secret` by sending to each player the portion of the binary string that they received during the preprocessing. If some player does not receive the correct part of the secret that was given to him during preprocessing, the player complains about the dealer.

In more detail, the preprocessing phase will generate a random degree- $(n - d - 1)$ polynomial v such that $v(0) = 0$. Additionally, the preprocessing phase generates a random string, encodes it in \mathbb{F} , and shares `secret` $\in \mathbb{F}$ using a random degree- $(n - 1)$ polynomial w (that is, $w(0) = \text{secret}$). Finally, the individual bits of the binary string are distributed among the n participating players (we assume a large enough field \mathbb{F} to facilitate this). Upon receiving $p(i)$ from the dealer in the online phase, the player P_i computes $z(i) \leftarrow p(i) \cdot v(i) + w(i)$ and sends it to the dealer. In case $(p(1), \dots, p(n))$ are not consistent with any degree- d polynomial, the dealer cannot reconstruct the value `secret` since the degree of z is larger than $n - 1$. As a result, the dealer would only be able to break soundness with a small number of players. The remaining players will complain, and conclude that the dealer is a cheater. On the other hand, if the dealer shared a low degree polynomial, the dealer can reconstruct `secret` by interpolating $z(1), \dots, z(n)$, and can then use `secret` as a proof that indeed $(p(1), \dots, p(n))$ define a degree- d polynomial. This can be done by sending each P_i its portion of the binary string encoded as `secret`.

Attack on Shares by Corrupt Players. Even if the dealer gives shares $p(1), \dots, p(n)$ consistent with a low degree polynomial, it may be that corrupt players would send back bad shares to prevent the dealer from reconstructing the correct secret, or by refusing to send shares altogether. To solve this problem, we allow the dealer to verify shares and eliminate players that send bad shares.

We allow the dealer to verify that P_i sent a share that equals $p(i) \cdot v(i) + w(i)$ by (1) having the preprocessing phase authenticate the shares $v(i), w(i)$ that it sends to each player P_i , (2) using the linearly homomorphic MAC from SPDZ, and (3) by giving the verification keys to the dealer. When a dealer complains about a player, the player will be eliminated and will no longer take part in any future degree tests with that dealer.² We use E to denote the set of eliminated players.

Properties About the Set of Eliminated Players. We need certain guarantees about the set of eliminated players. First, if the dealer is honestly sharing a low degree polynomial, then no honest player will complain about the dealer. Second, if the dealer is malicious and does not share a low degree polynomial, then a large number of honest players will eliminate themselves. Third, we must ensure that every player has a consistent view of the set of eliminated players. We satisfy this last property using a secure broadcast anytime a player is eliminated. If a large number of players are eliminated, then the dealer is replaced

² If the protocol re-starts because the dealer is thrown out, the party will rejoin the computation.

and protocol restarts with a new dealer. We can safely remove the dealer in this case, because, either the dealer is corrupt, or there are enough corrupt players that we can give up on robustness.

Recovering from Eliminated Players. The fact that the dealer can eliminate players creates a new problem: how does the dealer reconstruct **secret** when a few players have been eliminated? Recall that **secret** is shared using a degree- $n-1$ polynomial, and that eliminated players no longer provide shares to the dealer. In order to replace the eliminated players, we have the non-eliminated players send additional information that will allow the dealer to recover the missing shares of z . A natural approach for this is to have each remaining player send the dealer a share (generated during the preprocessing phase) of the eliminated player's share. While this solution works, it is too costly, as it introduces a quadratic overhead in the number of players. This overhead stems from two facts: first, a linear number of players could be eliminated, and second, for each execution of the degree test, for each eliminated player, each non-eliminated player would have to send one share to the dealer.

Reducing Recovery Overhead. We employ a couple of strategies to reduce the communication required of the honest players when they help the dealer to reconstruct the shares of eliminated players. First, we will reuse the same v for each execution of the degree test. Now, when a player P_i is eliminated by the dealer, each player will only need to send a share of v_i to the dealer once, ensuring that the dealer learns v_i for all further executions of the degree test protocol. Next, we notice that the dealer recovers **secret** from the shares of z by performing Lagrange interpolation, which is a linear operation. That is, the dealer computes $\mathbf{secret} = \sum_{i=1}^n \alpha_i z_i = \sum_{i=1}^n \alpha_i (p(i) \cdot v_i + w_i)$ where $\alpha_1, \dots, \alpha_n$ are the Lagrange interpolation coefficients. Rewriting the above equation,

$$\begin{aligned} \mathbf{secret} &= \sum_{i=1}^n \alpha_i z_i = \sum_{P_i \notin E} \alpha_i z_i + \sum_{P_i \in E} \alpha_i z_i = \sum_{P_i \notin E} \alpha_i z_i + \sum_{P_i \in E} \alpha_i (p(i) \cdot v_i + w_i) \\ &= \sum_{P_i \notin E} \alpha_i z_i + \sum_{P_i \in E} \alpha_i p(i) \cdot v_i + \sum_{P_i \in E} \alpha_i w_i. \end{aligned}$$

Since the dealer knows $p(i)$ for all players, knows v_i for all eliminated players, and has the shares z_i for all non-eliminated players, he only needs to learn $\bar{c} = \sum_{i \in E} \alpha_i w_i$. Thus, each non-eliminated player can locally compute a single share of \bar{c} , using a share of w_i for each $P_i \in E$. Sending just this single share to the dealer, instead of one share for every eliminated player, allows us to avoid the linear overhead that arose in the naive approach previously suggested.

3.1 Formal Description of the Degree Test Protocol

In this section we formally present and analyze our degree test protocol. Let H be the set of honest players and let E denote a global, shared variable, indicating

the set of eliminated players. We denote the inputs to the degree test protocol by $p(1), \dots, p(n)$. For some honest player, P_i , we let η denote the probability that a malicious dealer wrongly convinces P_i that p is of degree less than or equal to d . Finally, we denote the set of parties complaining about the dealer by C . The ideal functionality, \mathcal{F}_{dt} , is formally described in Fig. 1, and the degree test protocol realizing this functionality is described in Fig. 2. Consider the following theorem.




Fig. 1. Degree test functionality \mathcal{F}_{dt}

Theorem 1. π_{dt} securely realizes \mathcal{F}_{dt} in the preprocessing-hybrid model.

In order to prove Theorem 1, we provide two simulators, one simulator for the case when the dealer is honest, and a second simulator for the case when the dealer is corrupt. In each case the simulator simply follows the description of the protocol, determines if players or the dealer needs to complain, and adds players to the set of eliminated players E that would be eliminated. We recall that H denotes the set of honest players. E denotes the set of eliminated parties. The point $v_i \in \mathbb{F}$ is a share of a degree $n - t_p - \ell$ polynomial that evaluates to zero at zero. The point $w_i \in \mathbb{F}$ is a share of a polynomial that evaluates to a random value **secret**. The share $v_{i,j}$ is a resharing of v_i that will help the dealer to reconstruct v_i if P_i is eliminated.

Degree Test Simulation Honest Case. The simulator queries the ideal functionality and receives $p(i)$ for each $P_i \in \bar{H}$.

1. The simulator simulates the preprocessing by following its description.
2. The simulator sends $p(i)$ to each non-eliminated corrupt player $P_i \in \bar{E} \cup \bar{H}$

The degree test protocol is parameterized by security parameter s , n players, threshold t , degree d , and field \mathbb{F} where $|\mathbb{F}| \geq \frac{2ns}{n-t}$. Each player gets a point $p(i) \in \mathbb{F}$ from the dealer. The global variable E stores the set of eliminated players and may be updated. If E is updated, all parties will halt and discard their shares. Players in E do not send shares to dealer. Otherwise, they know that $\{p(i)\}_{P_i \in H}$ lies on a polynomial of degree at most d . We define $\alpha_1, \dots, \alpha_n$ as the lagrange coefficients. We use a homomorphic mac scheme (`setup`, `keygen`, `mac`, `verify`)

One-time preprocessing

1. $(v_1, \dots, v_n) \leftarrow \text{share}_{n-d}(0)$
2. Send $\gamma \leftarrow \text{setup}(1^s, \mathbb{F})$ to the dealer. (global key for mac scheme).
3. Send $k_i \leftarrow \text{keygen}(1^s, \mathbb{F})$ to the dealer. (key used to authenticate $p(i) \cdot v_i$)
4. Send $v_i, t_i \leftarrow \text{mac}(\gamma, k_i, v_i)$ to P_i
5. Generate shares to reconstruct values in case players are eliminated.
 $v_{i,j} \leftarrow \text{share}_t(v_i)$ and send $v_{1,j}, \dots, v_{n,j}$ to P_j

Preprocessing (per instance)

1. Set $m = \frac{2ns}{n-t-t}$ and sample a random `secret` $\in \{0,1\}^m$.
2. $(w_1, \dots, w_n) \leftarrow \text{share}_n(\text{secret})$ (Hides the secret.)
3. Send $\delta_i \leftarrow \text{keygen}(1^s, \mathbb{F})$ to the dealer.
4. $\tau_i \leftarrow \text{mac}(\gamma, \delta_i, w_i)$ ($\tau_i + p(i) \cdot t_i$ produces a mac tag for $p(i) \cdot v_i + w_i$)
5. Send (w_i, τ_i) and `secret` $_{m/n \cdot (i-1)+1, \dots, m/n \cdot i}$ to P_i .
6. Generate shares to recover in case players are eliminated.
 $w_{i,1}, \dots, w_{i,n} \leftarrow \text{share}_{t_p}(w_i)$ and send $w_{1,j}, \dots, w_{n,j}$ to P_j .

Protocol online phase

1. **dealer**
 Send $p(i)$ to P_i for each $P_i \notin E$.
2. **players**
 - i. $z_i \leftarrow p(i) \cdot v_i + w_i$ (z_i is a share of z such that $z(0) = \text{secret}$)
 - ii. Send $z_i, \sigma_i \leftarrow p(i) \cdot t_i + \tau_i$ to the dealer.
 - iii. Generate shares when players are eliminated.
 Send $a_i \leftarrow \sum_{P_j \in E} \alpha_j \cdot w_{j,i}$ to the dealer. (local shares of $\sum_{P_j \in E} \alpha_j w_j$)
3. **dealer**
 - i. Dealer complains about each non-eliminated player P_i who either: did not send a z_i with a tag σ_i such that $\text{verify}(\gamma, p(i) \cdot k_i + \delta_i, z_i, \sigma_i) = \text{accept}$; or did not send an a_i when players have been eliminated. If the dealer complained, E is updated to include all the players that he complained about, and the dealer skips to 5.i.
 - ii. Generate shares when players are eliminated.
 $\bar{c} \leftarrow \text{recover}([\mathbf{a}_i \mid P_i \notin E])$ ($\bar{c} = \sum_{P_j \in E} \alpha_j \cdot w_j$)
 $\bar{c} \leftarrow \bar{c} + \sum_{P_i \in E} \alpha_i (v_i \cdot p(i))$
 - iii. `secret` $\leftarrow (\sum_{P_i \notin E} \alpha_i \cdot z_i) + \bar{c}$ and send `secret` $_{m \cdot (i-1)+1, \dots, m \cdot i}$ to P_i
4. **players**
 - i. If dealer did not complain, and P_i received a different value for `secret` from the dealer then what he got in the preprocessing, then he complains. P_i is added to E . Each player P_i who received a complaint from the dealer (resp. P_j) about P_j (resp. Dealer) sends $(i, j, v_{j,i})$ to the dealer.
5. **dealer**
 - i. The dealer reconstructs v_j from the $v_{j,i}$ he receives using Reed-Solomon decoding. If the dealer cannot reconstruct v_j , the dealer broadcasts failure, and then E is set to be the set of all players.

Fig. 2. Degree test π_{dt}

3. The simulator await that corrupted non-eliminated player $P_i \in \bar{E} \cup \bar{H}$ sends $(z_i, \mathbf{m}(z_i))$ and a_i to the dealer.
4. The simulator computes $\mathbf{k}(z_i) \leftarrow p(i) \cdot \mathbf{k}(v_i) + \mathbf{k}(w_i)$, assigns to S the subset of corrupted non-eliminated players P_i who either did not send a z_i with a valid mac tag or who did not send an a_i . All players in S are added to E . If any players were added to E , he runs the player elimination simulation (below).
5. Send $\mathbf{secret}_{m \cdot (i-1)+1, \dots, m \cdot i}$ to each non-eliminated corrupt player P_i .
6. For each player P_i who complains about the dealer, the simulator sends $(\text{bad_proof_complaint}, i)$ to the functionality and then add P_i to E . The simulator then runs the Player elimination simulation (below).

Honest Dealer Elimination Simulation. Whenever a player is eliminated, we require that the simulator do the following. After a set S of players are added to E , the simulator awaits $(i, j, v'_{j,i})$ from each non-eliminated corrupt player for each $P_j \in S$. The, for each $P_j \in S$, the simulator tries to reconstruct v_j from the $v'_{j,i}$ that come from non-eliminated corrupt players, and the $v_{j,i}$ that were generated in the preprocessing for the honest players. If the simulator does not reconstruct a valid share v_j , the dealer broadcasts failure, and the full set of players is added E . The simulation then halts.

Description of the Simulator When Dealer Is Corrupt

1. The simulator simulates the preprocessing by following its description.
2. The simulator await that the dealer send $p(i)$ to each non-eliminated honest player P_i .
3. The simulator computes $z_i \leftarrow p(i) \cdot v_i + w_i$, $a_i \leftarrow \sum_{P_j \in E} \alpha_j \cdot w_{j,i}$ (local shares of $\sum_{P_j \in E} \alpha_j w_j$), $\mathbf{m}(z_i) \leftarrow p(i) \cdot \mathbf{m}(v_i) + \mathbf{m}(w_i)$ and sends $(z_i, \mathbf{m}(z_i))$ and a_i to the dealer.
4. For each non-eliminated player P_i that the dealer complains about, the simulator send $(\text{bad_proof_complaint}, i)$ to the functionality. The simulator executes the player elimination simulation.
5. The simulator await $\mathbf{secret}_{m \cdot (i-1)+1, \dots, m \cdot i}$ from the dealer for each honest non-eliminated player P_i .
6. For each non-eliminated honest player $P_i \in \bar{E} \cup H$, if the dealer did not send the same value of \mathbf{secret} that would have been to P_i , the simulator sends $(\text{bad_proof_complaint}, i)$ to the functionality. The simulator executes the player elimination simulation.

Corrupt Dealer Elimination Simulation. Whenever a player is eliminated players we require that the simulator do the following. After a set S of players are added to E , the simulator sends $(i, j, v'_{j,i})$ for each eliminated player $P_j \in S$ and non-eliminated honest player P_i . If the corrupt dealer broadcasts failure, then E is set to be the set of all players. The simulation then halts.

3.2 Properties of the Degree-Test Protocol

We already know that the degree test protocol securely realizes the degree test functionality. Within the context of our main protocol, we want to show that our degree test protocol has more features than what is directly provided by the functionality. The first is that the online cost of the degree test is low, namely that if we use the protocol many times, the overhead of the degree test per player will be small. The second condition that we want is that if the dealer is honest, and less than some threshold of players are dishonest, then in each execution of the degree test, either it succeeds, or some malicious party is eliminated.

The third condition that we are interested in is that if a corrupt dealer cheats by sharing a high degree polynomial, and does not complain about the shares and tags given to him by honest players, then less than half of the honest player's will accept the secret. (Recall, if he does complain about some of the shares and tags that he was given, then all parties are eliminated and the protocol re-starts with a new dealer.)

Lemma 1. *The total communication cost of running m executions of the degree test with the same dealer is $O(s \cdot n \cdot m + \text{poly}(n))$ bits.*

Proof. We enumerate over each item that is communicated and compute its associated communication overhead. A player will broadcast a complaint about the dealer at most once ($O(n^2)$). The dealer will broadcast a complaint about a player at most once. ($O(n^2)$). Each player will send a constant number of field elements to the dealer per execution of the degree test ($O(s \cdot n \cdot m)$). The dealer will send a constant number of field elements to each player per execution of the degree test ($O(s \cdot n \cdot m)$).

The communication complexity of all these items is $O(s \cdot n \cdot m + \text{poly}(n))$. This completes the proof of this lemma.

Lemma 2. *If the dealer is honest, and less than $\frac{(n-t_p-2\ell)}{2}$ players are corrupt, then both of the following conditions will be met (except with negligible probability).*

1. *No honest player will be eliminated.*
2. *The degree test will succeed, or at least one corrupt player will be eliminated.*

Proof. First, we show that an honest player will not be eliminated by an honest dealer except with negligible probability. Since honest players always send correct shares, the dealer would only eliminate an honest player if he reconstructs an incorrect value for the secret in step 3.iii. This can only occur if the adversary is able to successfully forge a mac tag in step 2.ii. Otherwise, the dealer would complain about a corrupt player and the degree test would terminate. Since forging a mac tag only succeeds with negligible probability, this completes the first part of the proof.

Next we proceed to show that either the degree test will succeed, or at least one corrupt player will be eliminated. If the dealer reconstructs the correct secret, the dealer will send the correct part of the secret to each honest player in step

3.iii, and each honest player will accept the secret in step 4.i. This leaves only two strategies for the adversary to prevent the degree test from succeeding: he can either send bad shares, or not send shares at all. In either case, the dealer will complain about corrupt players in step 3.i and the dealer will eliminate corrupt players. This completes the second part of the proof.

Lemma 3. *If less than t_p players are corrupt and the following conditions all hold, then more than $\frac{n-t_p-2\ell}{2}$ honest players will be eliminated.*

1. *The dealer is malicious and does not complain about a player in step 3.i.*
2. *The degree of the polynomial p is greater than ℓ .*

Proof. First, we show that if all the above conditions hold then the dealer cannot learn any information about the secret. Since by condition 2, the dealer shares a polynomial p of degree higher than d , then the degree $p \cdot v$ is greater than n . Since w was selected at random, and less than t_p players are corrupt, then the dealer cannot recover the secret from $(p \cdot v + w)(0)$.

By the first condition of the lemma, the dealer did not complain in step 3.i. This means that the dealer, to convince an honest player that he is honest must correctly guess the part of the secret given to that player. Since the probability of correctly guessing the secret for a given player is $2^{-\frac{4s}{n-t_p-\ell}}$, we can finally show that more than half the honest players will abort.

By combining the following two statements with the lemma below, we have what we want: (1) the probability of correctly guessing a player’s secret is $p = 1 - 2^{-4s/(n-t_p-\ell)}$ and (2) the random variables associated to the dealer successfully guessing players’ part of the secret are independent.

Lemma 4. *Given $s, m \in \mathbb{N}$, let X_1, \dots, X_m be independant Bernoulli variables with success probability $p = 1 - 2^{-4s/m}$ and let $X = \sum_{i=1}^m X_i$ then*

$$\Pr [X < \frac{m}{2}] \leq 2^{-\theta(s)}$$

Proof. If $m \geq s$, we can directly apply Chernoff’s bound to get this result. We have that $\mu = m \cdot (1 - 2^{-4s/m})$ and let $\delta = \frac{1}{2(1-2^{-4s/m})}$. We note that $\delta \geq \frac{1}{2}$ and that $\mu \cdot \delta = \frac{m}{2}$ and thus we have that

$$\Pr [X \leq \frac{m}{2}] = \Pr [X < (1 - \delta) \cdot \mu] \leq e^{-\frac{\delta^2 \mu}{3}} \leq e^{-\frac{m}{12}} \in 2^{-\theta(s)}$$

This leaves only the case where $m < s$ and we show that this also holds by using the following combinatorial argument.

$$\begin{aligned} \Pr [X < \frac{m}{2}] &= \sum_{i=0}^{m/2} \binom{m}{i} (1 - 2^{-4s/m})^i (2^{-4s/m})^{m-i} \leq \sum_{i=0}^{m/2} \binom{m}{i} (2^{-4s/m})^{m-i} \\ &\leq \sum_{i=0}^{m/2} 2^m (2^{-4s/m})^{m/2} \leq 2^{-2s+m+\log m} \in 2^{-\theta(s)} \end{aligned}$$

4 Additively-Secure Protocol

We now construct a protocol which is secure in the preprocessing-hybrid model, aside from allowing additive attacks. (Recall, these are then handled using the compiler of Genkin et al. [10].) The players will randomly elect (without repetition) a dealer that will be used to run the computation. If at some point too many players claim the dealer is cheating, the protocol will be restarted with a new dealer. During the evaluation phases (routing, multiplication and addition), the players will add, multiply and subtract shares locally, and they will also send and receive shares to and from the dealer. The dealer will be responsible for receiving, reconstructing values and resharing them. In particular, the dealer will be responsible for reconstructing values when less than t_r shares are corrupted. This will be done by having the dealer apply Reed-Solomon decoding to the shares he receives. If at any point, the dealer fails to reconstruct the secret, the dealer will eliminate himself and the protocol will be restarted with a new dealer. The protocol employs the degree testing protocol to ensure that when the dealer reshares values, he cannot use a polynomial of degree greater than ℓ . Since degree testing involves eliminating players, the protocol will need to keep track of who has been eliminated.

At the beginning, the players will randomly elect a dealer. While that dealer is active, each party will keep track of a set of eliminated players denoted by the variable E . Players can eliminate themselves if they detect malicious behavior, or they can be eliminated, either by an honest dealer for acting maliciously, or by a malicious dealer, arbitrarily. If the set of eliminated players grows too big, all parties kick out the dealer and rejoin the protocol with a new dealer (chosen without replacement). To simplify exposition, we assume that the set E is a global variable, and that all honest parties agree on its members. In practice, this can be achieved using a broadcast channel, without impacting the claimed communication cost. Our main protocol consists of four phases: the preprocessing phase, the input phase, the evaluation phase, and the output phase. The input, evaluation, and output phase will all rely on values generated by the preprocessing. We will not describe the preprocessing phase in its entirety but rather describe which values each of the other phases need from the preprocessing.

Throughout the computation, parties hold shares of wire values, encoded using polynomials of degree $t = t_p + \ell$. This ensures that t_p parties cannot learn anything about these values. Because we need to multiply these polynomials by degree ℓ polynomials that encode masked wire values, the degree of the polynomials becomes $t_p + 2\ell$ during the evaluation. This allows us to error-correct in the presence of less than $\frac{n-t_p-2\ell}{2}$ corruptions, maintaining robustness as claimed in our theorem. When we do not specify the degree of a sharing, we mean that the polynomial has degree $t_p + \ell$ (Figs. 3 and 4).

Input Phase. In the input phase, a sender provides his input \mathbf{x} , and the other parties receive shares of that input, which can then be used in the evaluation phase. The preprocessing functionality randomly samples $\mathbf{r} \in \mathbb{F}^\ell$, gives the value

Below, we let $t = t_p + \ell$.

Input

1. Randomly sample $r \in \mathbb{F}^\ell$ and send it to the sender.
2. Send $[\mathbf{y}] \leftarrow \text{share}_t(r)$ to players

ρ -gate

1. Randomly sample $\mathbf{u}, \mathbf{v} \in \mathbb{F}^\ell$ such that $\rho(\mathbf{u}) = \mathbf{v}$.
2. Send $[\mathbf{c}] \leftarrow \text{share}_t(\mathbf{u}), [\mathbf{z}] \leftarrow \text{share}_t(\mathbf{v})$ to players.

Multiplication

1. Randomly sample $\alpha, \beta, \gamma, v \in \mathbb{F}^\ell$ such that $\alpha \cdot \beta = \gamma$.
2. $[\mathbf{a}] \leftarrow \text{share}_t(\alpha), [\mathbf{b}] \leftarrow \text{share}_t(\beta), [\mathbf{c}] \leftarrow \text{share}_t(\gamma),$
 $[\mathbf{u}] \leftarrow \text{share}_t(v)$
3. Send $[\mathbf{a}], [\mathbf{b}], [\mathbf{c}], [\mathbf{u}]$ to players.

Output

1. $r \in_R \mathbb{F}^\ell. [\mathbf{r}] \leftarrow \text{share}_t(r)$
2. Send $[\mathbf{r}], [[\mathbf{r}]]$ to players. ($[[\mathbf{r}]]$ is a VSS sharing of r .)

Fig. 3. Preprocessing

to the sender, and provides $[\mathbf{r}]$ to the other players. The sender broadcasts $\mathbf{y} = \mathbf{x} - \mathbf{r}$. The players then compute $[\mathbf{x}] = [\mathbf{r}] + \mathbf{y}$. Due to a lack of space, we provide the full description in the full version of the paper.

Output Phase. In the output phase, parties take shares $[\mathbf{x}]$ of the output and reconstruct \mathbf{x} . We have to limit the adversary to an additive attack on the revealed output, ensuring that the adversary cannot arbitrarily choose the output. The preprocessing functionality creates two shares of a value $\mathbf{r} \in \mathbb{F}^\ell$, once using packed Shamir secret sharing resulting, in $[\mathbf{r}]$, and once using a VSS, resulting in $[[\mathbf{r}]]$. The players will use $[\mathbf{r}]$ and $[[\mathbf{r}]]$ to mask and then unmask \mathbf{x} . That is, they locally, homomorphically add r to the output by adding their shares of $[\mathbf{x}]$ and $[\mathbf{r}]$, and then reveal r by opening the VSS sharing. Because VSS is binding, the adversary can only modify the value before it is unmasked. As such, the only attack that can be done is an additive attack. Due to a lack of space, we provide the full description in the full version of our paper.

Multiplication. The multiplication is the most complex operation, the goal is to take shares $[\mathbf{x}], [\mathbf{y}]$ and produce shares of $[\mathbf{x} \cdot \mathbf{y}]$. To do so, we will use beaver triple $[\mathbf{a}], [\mathbf{b}], [\mathbf{a} \cdot \mathbf{b}]$, and a sharing of a random $\mathbf{r} \in \mathbb{F}^\ell, [\mathbf{r}]$. First the players will send $[\mathbf{x} - \mathbf{a}], [\mathbf{y} - \mathbf{b}]$. The dealer will reconstructs values $x - a, y - b, (x - a)(y - b)$, and re-shares them using degree ℓ polynomials. The players verify that the shares given to them by the dealer are of degree ℓ , using the degree test protocol. The players will then compute $[\mathbf{u}] = [\mathbf{x} - \mathbf{a}]_\ell \cdot [\mathbf{y}] + [\mathbf{y} - \mathbf{b}]_\ell \cdot [\mathbf{x}] + [\mathbf{r}]$ and send the shares to the dealer. The dealer re-shares \mathbf{u} using a degree ℓ polynomial, and the players again test that the degree is no more than ℓ . Finally, the players will compute $[\mathbf{x} \cdot \mathbf{y}] = [\mathbf{u}] - [\mathbf{a} \cdot \mathbf{b}] - [\mathbf{r}]$.

The players want to evaluate a SIMD circuit C with m input gates g_1, \dots, g_m with output gates $g_{|C|-O+1}$ to $g_{|C|}$. We denote $g_{m+1}, \dots, g_{|C|-O}$ as evaluation gates. For each input gate g , we denote p_g as the player who provides input x_g and we will use the variable $[\mathbf{i}_g]$ to store the shares generated by \mathcal{F}_{pre} for the gate g .

For each multiplication gate g , we will use the variable $[\mathbf{a}_g], [\mathbf{b}_g], [\mathbf{c}_g]$ and $[\mathbf{u}_g]$ to store the shares generated by \mathcal{F}_{pre} for the gate g . For each ρ -gate, we will use the variable $[\mathbf{v}_g], [\mathbf{w}_g]$ to store the shares generated by \mathcal{F}_{pre} for gate g . We use t_g to denote the shares generated by evaluating the gate g . For each output gate, We will use the variable $[\mathbf{r}_g]$ to store the shares generated by packed secret sharing for the output phase and define $[\mathbf{\Phi}_g]$ to denote the shares generated using the VSS.

We use E to denote the set of eliminated players; it is a global variable which is sometimes updated in the sub-protocols. Each round will have a fresh dealer. We denote \mathcal{D} as the set of eliminated dealers.

Preprocessing

1. For each input gate g
 The players execute $\mathcal{F}_{\text{pre}}(\text{input})$ where p_g taking the role of the sender. The sender stores the value he receives as r_g and the shares received by players are assigned to $[\mathbf{i}_g]$.
2. For each ρ -gate g , $[\mathbf{v}_g], [\mathbf{w}_g] \leftarrow \mathcal{F}_{\text{pre}}(\rho)$
3. For each multiplication gate g , $[\mathbf{a}_g], [\mathbf{b}_g], [\mathbf{c}_g], [\mathbf{u}_g] \leftarrow \mathcal{F}_{\text{pre}}(\text{multiplication})$
4. For each output gate g ,
 - (a) $[\mathbf{y}_g], [[\mathbf{r}_g]] \leftarrow \mathcal{F}_{\text{pre}}(\text{output})$
 - (b) $[\mathbf{r}_g] \leftarrow [\mathbf{y}_g] \quad [\mathbf{\Phi}_g] \leftarrow [[\mathbf{r}_g]]$

Dealer replacement rule.

If at any point $|E_i| \geq \frac{n-t_p}{2} - \ell$, the player sets $\mathcal{D} \leftarrow \mathcal{D} \cup \{D\}$ and return to step 3 of the evaluation.

Evaluation

1. For each input gate g , $[\mathbf{t}_g] \leftarrow \text{Input}(x_g, r_g, [\mathbf{i}_g])$ where x_g and r_g are the inputs of sender p_g .
2. $\mathcal{D} \leftarrow \emptyset, E \leftarrow \emptyset$.
3. Sample a random dealer not in \mathcal{D} .
4. For each evaluation layer L and for each gate $g \in L$.
 - (a) If L is an addition layer then $[\mathbf{t}_g] \leftarrow [\mathbf{x}] + [\mathbf{y}]$
 - (b) Otherwise
 - If L is a routing layer then

$$[\mathbf{x}] \leftarrow [\mathbf{t}_{\text{input}(g)}]$$

$$[\mathbf{t}_g] \leftarrow \text{routing}([\mathbf{x}], [\mathbf{v}_g], [\mathbf{w}_g])$$
 - If L is a multiplication layer then

$$[\mathbf{x}] \leftarrow t_{\text{left}(g)}, [\mathbf{y}] \leftarrow t_{\text{right}(g)}$$

$$[\mathbf{t}_g] \leftarrow \text{multiplication}([\mathbf{x}], [\mathbf{y}], [\mathbf{a}_g], [\mathbf{b}_g], [\mathbf{c}_g], [\mathbf{u}_g])$$
 - If $|E| \geq \frac{n-t_p}{2} - \ell$, $\mathcal{D} \leftarrow \mathcal{D} \cup \{D\}$ and return to step 3 of the evaluation.
 - If $[\mathbf{t}_g] = \perp$, return to step 4.b (This occurs when a player is eliminated during a degree test.)
5. For each output gate g
 - (a) $[\mathbf{x}] \leftarrow [\mathbf{t}_{\text{input}(g)}]$
 - (b) $\mathbf{v} \leftarrow \text{output}([\mathbf{x}])$
 - (c) Output \mathbf{v}

Fig. 4. Main protocol

Routing. The input is $[\mathbf{x}]$ and the output should be $[\rho(\mathbf{x})]$. The preprocessing functionality generates shares $[\mathbf{r}], [\mathbf{r}']$ such that $\rho(\mathbf{r}) = \mathbf{r}'$. Then when provided $[\mathbf{x}]$, the players will send $[\mathbf{x} + \mathbf{r}]$ to the dealer who will reshare $[\rho(\mathbf{x} + \mathbf{r})]$. The players will then verify that the dealer reshared $\mathbf{x} + \mathbf{r}$ using a low degree polynomial via the degree test. The players will then compute $[\mathbf{x}] = [\rho(\mathbf{x} + \mathbf{r})] - [\mathbf{r}']$. Due to a lack of space, we provide the full description in the full version of our paper.

Formally, we prove that our protocol realizes two things. First, we show that the protocol securely realizes f_C with low expected communication overhead if less than t_r players are corrupt. Second, we prove that our protocol securely realizes $f_C^{\mathbf{A}}$ (the functionality that allows the adversary to tamper with each wire individually) if less than t_p players are corrupt. Then, by running our protocol on a circuit secure against tampering on individual wires, our protocol securely realizes f_C . The compilers of [10, 11] allow us to compile any circuit into an equivalent circuit that is secure against individual tampering with only a constant blowup in circuit size. As a result, it is easy to see that by employing our protocol with the results of [10, 11], we achieve the desired security properties as well as the desired level of efficiency (Fig. 5).

Theorem 2. *For any number of players n , privacy threshold $n \geq t_p \geq \frac{n}{2}$, packing parameter $\ell < \frac{n-t_p}{2}$, $\pi_{\text{mpc}}(t_p, O(1/|\mathbb{F}|))$ -securely realizes $f_C^{\mathbf{A}}$ with abort in the \mathcal{F}_{pre} -hybrid model.*

Theorem 3. *For any number of players n , privacy threshold $n \geq t_p \geq \frac{n}{2}$, robustness threshold $t_r \leq \frac{n-t_p-2\ell}{2}$, $\ell < \frac{n-t_p}{2}$. $\pi_{\text{mpc}}(t_p, t_r, O(1/|\mathbb{F}|))$ -securely realizes $(f_C, f_C^{\mathbf{A}})$ for arithmetic circuit C with depth d in the \mathcal{F}_{pre} -hybrid model, with full security, and expected communication overhead*

$$O(|C| \log(|C|) \cdot \frac{n}{\ell} + d^2 \cdot n + \text{poly}(n, s)).$$

Due to space constraint, we only provide a short summary of how we prove our main protocol secure. A more complete argument appears in the full version of the paper.

Security Under Honest Dealer. Since our protocol is in a hybrid-model, the simulator can simulate a run of the preprocessing functionality and store the generated values. This allows the simulator to extract the adversary's inputs. The simulator runs the honest parties with dummy inputs and determines whether the adversary causes the honest dealer to abort, or causes an additive attack. We show that if the adversary sends bad shares to the dealer, then the simulator can determine, based solely on these shares, which of these three things happen: (1) the dealer aborts because he failed to reconstruct a secret, (2) the bad shares can be ignored (which is the case if enough players are honest), or (3) the attack by the adversary can be mapped to an additive attack. We prove the previous statement by using the fact that Shamir secret sharing is a linear error correcting code, and from the following fact about such codes.

The players input sharings $[\mathbf{x}], [\mathbf{y}]$, a triple of shares $[\mathbf{a}], [\mathbf{b}]$ and $[\mathbf{c}]$ such that $\mathbf{c} = \mathbf{a} \cdot \mathbf{b}$, and sharing $[\mathbf{r}]$ of a random value $\mathbf{r} \in \mathbb{F}^\ell$. All of these values are shared using polynomials of degree $t_p + \ell$. The shares $[\mathbf{r}]$ are used to hide values from the dealer. The global variable E , which stores the set of eliminated players, may be updated. If it is updated, all players terminate, discard their shares, and output \perp . Otherwise, the players produce shares $[\mathbf{z}]$ such that $\mathbf{z} = \mathbf{x} \cdot \mathbf{y}$.

Evaluation

1. **players**
 - i. $[\mathbf{d}] \leftarrow [\mathbf{y}] - [\mathbf{b}]$
 - ii. $[\mathbf{e}] \leftarrow [\mathbf{x}] - [\mathbf{a}]$
 - iii. Send $[\mathbf{d}], [\mathbf{e}]$ to dealer.
2. **dealer**
 - i. $\mathbf{v} \leftarrow \text{recover}([\mathbf{d}]) \cdot \text{recover}([\mathbf{e}])$
 - ii. $[\boldsymbol{\alpha}]_\ell \leftarrow \text{share}_\ell(\mathbf{v})$
 - iii. $[\boldsymbol{\epsilon}]_\ell \leftarrow \text{reshare}_\ell([\mathbf{e}])$
 - iv. $[\boldsymbol{\delta}]_\ell \leftarrow \text{reshare}_\ell([\mathbf{d}])$
3. **Degree Test**
 - i. Execute degree-testing with dealer input $[\boldsymbol{\alpha}]_\ell$.
 - ii. Execute degree-testing with dealer input $[\boldsymbol{\delta}]_\ell$.
 - iii. Execute degree-testing with dealer input $[\boldsymbol{\epsilon}]_\ell$.
 - iv. If the degree test outputs \perp in any of these executions, then all players terminate and output \perp .
4. **players**
 - i. $[\mathbf{g}_1] \leftarrow [\boldsymbol{\delta}]_\ell \cdot [\mathbf{x}] + [\boldsymbol{\epsilon}]_\ell \cdot [\mathbf{y}]$
 - ii. $[\mathbf{g}] \leftarrow [\mathbf{g}_1] - [\boldsymbol{\alpha}]_\ell + [\mathbf{r}]$
 - iii. Send $[\mathbf{g}]$ to dealer.
5. **dealer**
 - i. $[\mathbf{h}]_\ell \leftarrow \text{reshare}_\ell([\mathbf{g}])$
6. **Degree Test**
 - i. Execute degree-testing with dealer input $[\mathbf{h}]_\ell$.
 - ii. If the degree test outputs \perp , then all players terminate and output \perp .
7. **players**
 - i. Set $\mathbf{O} \leftarrow [\mathbf{h}]_\ell + [\mathbf{c}] - [\mathbf{r}]$
 - ii. **Output \mathbf{O}**

Recovery failure subroutine

1. If $\text{recover}(\cdot)$ ever fails to output a polynomial of degree $t_p + \ell$ that agrees with $\frac{n+t_p}{2} + \ell$ shares, then the dealer broadcasts “failure”, sets $E \leftarrow P$, and all players terminate with output \perp . The parameters in this subroutine are chosen so that it is guaranteed that either the Berklaml-welch algorithm reconstruct a correct value for the secret or that the recover operation will output \perp .

Fig. 5. Multiplication

Let $(\text{encode}, \text{decode})$ be a linear error-correcting, let $c = \text{encode}(m)$ be an encoding of m , and let be μ be an error vector. By linearity of the error-correcting code, we have that $\text{decode}(\text{encode}(m) + \mu) = \text{decode}(\text{encode}(m)) + \text{decode}(\mu)$. In particular, this implies that $\text{decode}(\text{encode}(m) + \mu) = \perp$ if and only if

$\text{decode}(\mu) = \perp$. The error vector μ in this case represents the difference between the shares the adversary should have sent, versus the shares it actually sent.

Security Under Malicious Dealer. At a high level, the simulation of a malicious dealer is similar to that of an honest dealer. The main difference is that this simulator must ensure that the dealer does not share a polynomial of too high a degree. This is easily detected by inspecting the shares sent to the degree-test functionality, and the dealer can then be replaced. If the dealer's polynomial is of the appropriate degree, the simulator can compute the value of an additive attack by reconstructing the shared secret and comparing it with the secret that the dealer should have sent.

Acknowledgments. This material is based upon work supported by the National Science Foundation under Grants No. #1564088, #1111599, #1514261, #1652259 and #1563722. Daniel Genkin was also supported by financial assistance award 70NANB15H328 from the U.S. Department of Commerce, NIST, the 2017–2018 Rothschild Postdoctoral Fellowship, and DARPA Contract #FA8650-16-C-7622.

References

1. Beaver, D.: Efficient multiparty protocols using circuit randomization. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 420–432. Springer, Heidelberg (1992). https://doi.org/10.1007/3-540-46766-1_34
2. Damgård, I., Ishai, Y., Krøigaard, M.: Perfectly secure multiparty computation and the computational overhead of cryptography. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 445–465. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13190-5_23
3. Damgård, I., Ishai, Y., Krøigaard, M., Nielsen, J.B., Smith, A.: Scalable multiparty computation with nearly optimal work and resilience. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 241–261. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-85174-5_14
4. Damgård, I., Keller, M., Larraia, E., Pastro, V., Scholl, P., Smart, N.P.: Practical covertly secure MPC for dishonest majority – or: breaking the SPDZ limits. In: Crampton, J., Jajodia, S., Mayes, K. (eds.) ESORICS 2013. LNCS, vol. 8134, pp. 1–18. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40203-6_1
5. Damgård, I., Nielsen, J.B.: Scalable and unconditionally secure multiparty computation. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 572–590. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-74143-5_32
6. Damgård, I., Pastro, V., Smart, N., Zakarias, S.: Multiparty computation from somewhat homomorphic encryption. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 643–662. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32009-5_38
7. Damgård, I., Zakarias, S.: Constant-overhead secure computation of Boolean circuits using preprocessing. In: Sahai, A. (ed.) TCC 2013. LNCS, vol. 7785, pp. 621–641. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-36594-2_35

8. Dowsley, R., Müller-Quade, J., Otsuka, A., Hanaoka, G., Imai, H., Nascimento, A.C.A.: Universally composable and statistically secure verifiable secret sharing scheme based on pre-distributed data. *IEICE Trans.* **94–A**(2), 725–734 (2011)
9. Franklin, M.K., Yung, M.: Communication complexity of secure computation (extended abstract). In: *Proceedings of the 24th Annual ACM Symposium on Theory of Computing*, pp. 699–710 (1992)
10. Genkin, D., Ishai, Y., Polychroniadou, A.: Efficient multi-party computation: from passive to active security via secure SIMD circuits. In: Gennaro, R., Robshaw, M. (eds.) *CRYPTO 2015*. LNCS, vol. 9216, pp. 721–741. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48000-7_35
11. Genkin, D., Ishai, Y., Prabhakaran, M., Sahai, A., Tromer, E.: Circuits resilient to additive attacks with applications to secure computation. In: *Symposium on Theory of Computing, STOC 2014*, pp. 495–504 (2014)
12. Hazay, C., Orsini, E., Scholl, P., Soria-Vazquez, E.: Efficient MPC from syndrome decoding (or: Honey, I shrunk the keys). *IACR Cryptology ePrint Archive 2018:208* (2018)
13. Ishai, Y., Katz, J., Kushilevitz, E., Lindell, Y., Petrank, E.: On achieving the “best of both worlds” in secure multiparty computation. *SIAM J. Comput.* **40**(1), 122–141 (2011)
14. Ishai, Y., Kushilevitz, E., Prabhakaran, M., Sahai, A., Yu, C.-H.: Secure protocol transformations. In: Robshaw, M., Katz, J. (eds.) *CRYPTO 2016*. LNCS, vol. 9815, pp. 430–458. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53008-5_15
15. Katz, J.: On achieving the “best of both worlds” in secure multiparty computation. In: *Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, pp. 11–20 (2007)
16. Nielsen, J.B., Ranellucci, S.: On the computational overhead of MPC with dishonest majority. In: Fehr, S. (ed.) *PKC 2017*. LNCS, vol. 10175, pp. 369–395. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54388-7_13