

Constant-Round MPC with Fairness and Guarantee of Output Delivery

S. Dov Gordon*, Feng-Hao Liu**, and Elaine Shi***

Abstract. We study the round complexity of multiparty computation with fairness and guaranteed output delivery, assuming existence of an honest majority. We demonstrate a new lower bound and a matching upper bound. Our lower bound rules out any two-round fair protocols in the standalone model, even when the parties are given access to a common reference string (CRS). The lower bound follows by a reduction to the impossibility result of virtual black box obfuscation of arbitrary circuits.

Then we demonstrate a three-round protocol with guarantee of output delivery, which in general is harder than achieving fairness (since the latter allows the adversary to force a fair abort). We develop a new construction of a threshold fully homomorphic encryption scheme, with a new property that we call “flexible” ciphertexts. Roughly, our threshold encryption scheme allows parties to adapt flexible ciphertexts to the public keys of the non-aborting parties, which provides a way of handling aborts without adding any communication.

1 Introduction

Secure multi-party computation (MPC) allows mutually distrusting parties to securely compute a function on their inputs with several desired properties, including: **correctness** (honest parties should not receive a wrong output), and **privacy** (corrupted parties cannot learn anything beyond the prescribed output). In addition to these two basic properties, one might further require **fairness** (corrupted parties receive their output only if all honest parties receive output), or the stronger **guarantee of output delivery** (corrupted parties cannot prevent honest parties from receiving their output). Alternatively, a relaxed security notion is often used, called **security with abort** – it is possible that the attacker can prevent the honest parties from receiving output. All of these requirements can be formalized in an Ideal/Real paradigm [5, 13], which provides a nice way to analyze security.

* Dept. of Computer Science, George Mason University. This work was done when the author was a research scientist at Applied Communication Sciences. Email: crypto@dovgordon.com

** Dept. of Computer & Electrical Engineering and Computer Science, Florida Atlantic University. The work was done when the author was a postdoc at the University of Maryland. Email: fenghao.liu@fau.edu.

*** Dept. of Computer Science, Cornell University. Email: runting@gmail.com.

In this work, we explore the *round complexity* required for achieving these various properties. For the setting of *security with abort*, we already understand the round complexity fairly well – Asharov et al. [1] constructed a 3-round protocol (in the common reference string model) under the learning with error (LWE) assumption; Garg et al. [10] constructed a 2-round protocol for general computation (in the CRS model) using indistinguishable obfuscation; it is well-known that one-round protocols are in general not possible.

However, for protocols with *fairness* and *guarantee of output delivery*, our understanding of round complexity is still incomplete. Regarding *feasibility*, everything is well understood: if there is no honest majority, Cleve proved [6] that fair MPC for general computation is not possible. In the setting of an honest majority, we know that we can always achieve fairness [4], and, assuming a broadcast channel, we can always guarantee output delivery [7]. However, the optimal round complexity for this setting (of an honest majority) is still unknown¹. Asharov et al. [1] show that their basic protocol can be extended to achieve security with guarantee of output delivery (and thus fairness) in 5 rounds, assuming there exists an honest majority. By slightly modifying the multi-key FHE protocol of Lopez-Alt, Tromer, and Vaikuntanathan [18], we can obtain a 5-round protocol with guarantee of output delivery, assuming there exists an honest majority. This is the best known round complexity for achieving fairness for any $t < N/2$. For the lower-bounds, Gennaro et al. [11] showed that there are functionalities that cannot be computed (fairly) by 2-round protocols, even in the CRS model. Recently, Garg et al. [10] claimed that their 2-round protocol (in the CRS model) also achieves fairness, but the claim contradicts the lower-bound of [11], as well as the (stronger) lower bound we present here.

1.1 Our Results

Our main two results are matching upper and lower bounds for three-round multiparty computation with guaranteed output delivery with security against a malicious minority of parties. More specifically:

- We show that 2-round, fair MPC for general functions is impossible, even if there is an honest majority. We strengthen the impossibility result of Gennaro et al. [11], demonstrating impossibility even when a fail-stop adversary corrupts only a single party. Both our result and the result of Gennaro et al. extend to the CRS model. (Section 3.)
- There exists a 3-round MPC with guaranteed output delivery for general functions in the CRS model, secure against a minority of semi-honest fail-stop adversaries. The security relies on the learning with errors (LWE) assumption. (Section 4.)

¹ For more restricted corruption settings, we do know how to construct 2-round protocols. See related work for more discussions.

- If parties have access to an authenticated broadcast channel², then the above 3-round protocol can be upgraded to one that is secure against malicious adversaries, without any additional rounds. (Section 5.)
- Additionally, we show that security of the two-round protocol by Garg et al. [10] can be based on witness encryptions for general NP statements, which is weaker than indistinguishable obfuscation for general circuits³ as presented in their work. Together with an idea in Section 1.2, we can construct a three-round *fair* protocol (but not guarantee of output delivery) based on witness encryptions for general NP statements. Due to space limit, we present the results in the full version of this paper [14].

In summary, 2-round general fair MPC is not possible, and 3-round general MPC with guarantee of output delivery can be constructed under a falsifiable assumption. Guarantee of output delivery implies fairness (by definition), and thus 3-round fair MPC can also be constructed under the same falsifiable assumption.

All of our positive results are UC-secure [5]. Our protocols, along with those appearing in the prior work of Garg et al. [10] and Asharov et al. [1], require a CRS.

1.2 Overview of our Techniques

Impossibility of Fairness in Two Rounds. We show that a two round, fair, polynomial-time protocol for general functions yields a construction of virtual black box (VBB) secure program obfuscation for P/Poly, in contradiction of the well-known impossibility result of Barak et al. [2].

Consider a symmetric 3-ary functionality $f(x_1, x_2, x_3)$ that interpret x_1 as a circuit C , ignores x_3 and outputs $C(x_2)$. Suppose there exists a two-round fair protocol π that computes f with fairness, then we make the following observations. We assume the three parties are Alice, Bob, and Charlie.

- If the adversary (only) corrupts Alice and instructs her to abort in the second round, then after Bob and Charlie send their messages in the second round, the adversary can learn the output $C(x_2)$.
- By the property of fairness, Bob and Charlie must be able to learn the output $C(x_2)$, since the adversary in the above case has learned the output.
- It follows that Alice’s second message is redundant. Whether she sends her second message or not, the other parties can compute the outcome.

Using the above observations, we can construct a program obfuscator for general circuits: we view Alice’s first message as the obfuscation of C . To evaluate $C(x)$, we just simulate Bob and Charlie with Bob’s input x . Since Alice’s first

² An authenticated broadcast channel enables a party to send a message to all other parties, ensuring that each party knows both the identity of the sender and that all other parties have received the same message.

³ Indistinguishable obfuscation for general circuits is a stronger assumption. We know that indistinguishable obfuscation for general circuits implies witness encryption for general NP statements, but the other way is unclear.

message is independent of the other parties' inputs, we can rewind Bob and Charlie and compute $C(x)$ repeatedly on arbitrary values of x .

We note that Garg et al. state (without proof) that their two-round protocol achieves fairness [10], and one can see why this mistake might have been made. Their protocol works by collapsing some protocol with greater round complexity into a two-round protocol through the use of obfuscation, and they state that if the underlying protocol is fair, then the resulting two round protocol will also be fair. Speaking very roughly, in their construction, each party sends a commitment to their input and their randomness in round one, and in round two, they each send obfuscations of the next message functions from the underlying fair protocol. They then each finish the protocol locally, using the obfuscated programs to generate the correct protocol messages. At first glance, it would seem that this preserves fairness, because if a party aborts in round two, the other parties can simply generate next-messages as though he aborted, and, by the fairness of the underlying protocol, fairness should be preserved. In fact, this misses the following subtlety. If a party aborts in round two, he still receives all of the obfuscated programs, and can still compute the output of the function: this is equivalent to aborting in the very last round of the underlying fair protocol. On the other hand, because he never sent his obfuscated next-message programs, the other parties will be forced to treat him as though he aborted in round one of the underlying protocol, perhaps replacing his input with some default value. In particular, then, it could be that the malicious party learns $f(x_1, \dots, x_N)$ while the other parties learn $f(\perp, x_2, \dots, x_N)$.

Fairness in Three Rounds. The construction of Garg et al. can be modified slightly to get a three-round fair protocol, as we now outline. However, we note that there is no clear way to guarantee output delivery without increasing the round complexity; our main technical result is a new protocol for achieving guaranteed output delivery in three rounds.

To achieve fairness in three rounds, we can start with the protocol of Garg et al., but instead of sending obfuscations of the next message functions that compute the underlying secure computation, the parties will send obfuscations that compute an $N/2$ -out-of- N secret sharing of the output. They then add one additional round to reconstruct the output. Now, if the adversary aborts in round two, even though he learns all of the next message functions, he still cannot recover the output (since there is an honest majority). If he aborts in round three, the honest parties already have enough shares to reconstruct the output on their own.

In general, we can compile any fair protocol into one that guarantees output delivery, assuming a broadcast channel [7], but we cannot necessarily preserve the round complexity. In the particular protocol just described, note that the obfuscated programs sent in round two have commitments to the parties' inputs embedded inside of them. If a party aborts in round two, the other parties would need to replace their obfuscations, embedding a commitment to some default input value in place of the aborting party's true input. But this will incur additional communication rounds.

Guarantee of Output Delivery. Before we describe our protocol, we first give an overview the approach by Asharov et al. [1]. Asharov et al. proposed a new primitive called *Threshold Fully Homomorphic Encryption* (TFHE), which is essentially a distributed version of fully homomorphic encryption (FHE). For their TFHE, there is a joint public key \mathbf{pk}^* whose secret key is shared among all parties, i.e. $\mathbf{sk}^* = \mathbf{sk}_1 + \mathbf{sk}_2 + \dots + \mathbf{sk}_N$. (There is also an evaluation key, but we omit it for simplicity of exposition). The keys $(\mathbf{pk}^*, \mathbf{sk}^*)$ constitute an FHE key pair, so the encryption and evaluation algorithms can remain the same as those used in the original FHE scheme. To decrypt, parties need to run a *threshold decryption protocol*, since the secret key is shared among all parties.

Using the TFHE scheme, their basic three round protocol has the following structure: (1) in the first round parties establish a joint public key \mathbf{pk}^* ; (2) in the second round parties output an encryption of their inputs, i.e. $\text{Enc}_{\mathbf{pk}^*}(x_i)$; (3) in the third round, parties perform the homomorphic operations (for computing f) to obtain an evaluated ciphertext C^* , and then run the threshold decryption protocol to decrypt C^* . Asharov et al. [1] presented a simple idea to make the basic protocol fair (and to guarantee output delivery) in the first round, the parties also secret share their inputs and all random coins. If any party aborts in the second or third round, the honest majority would reconstruct his states and resume the protocol. (Note that if a party aborts in the first round, he is simply ignored). This approach will add two additional rounds for the worst case.

We note that their construction uses an N -out-of- N sharing of the secret key \mathbf{sk}^* , so they require all parties in order to decrypt. This means, if any party aborts, the other parties need to reconstruct his view to resume. Thus, these two additional rounds seem inherent if we follow this approach. To get a 3-round protocol, we need a new approach. In particular we propose and construct a new variant of TFHE with more fine-grained features. Using it as a building block, we are able to get around the barriers mentioned above. We highlight our new ideas below.

Instead of establishing a “fixed” joint public key, our new TFHE uses $\mathbf{pk}_{[N]} = \{\mathbf{pk}_i\}_{i \in [N]}$ as the public keys, where \mathbf{pk}_i is contributed by party P_i . Then with $\{\mathbf{pk}_1, \dots, \mathbf{pk}_N\}$, P_i can encrypt the input x_i and produce a **flexible** ciphertext C_i . We introduce a new algorithm $\text{TransCT}(C; S)$ that transforms a flexible ciphertext C into C' , where C' is with respect to the public keys $\mathbf{pk}_S = \{\mathbf{pk}_j : j \in S\}$. Intuitively, a *flexible* ciphertext is one that is not yet committed to a set of public keys, and a *transformed* one commits to some \mathbf{pk}_S and can be homomorphically evaluated. Finally, our threshold decryption protocol works when there is an honest majority of parties (as opposed to the previous one which requires all parties).

Using the new TFHE, our protocol has the following structure: (1) in the first round, parties generate $\{\mathbf{pk}_1, \dots, \mathbf{pk}_N\}$; (2) in the second round, each party output a flexible ciphertext $C_i = \text{Enc}(x_i)$; (3) let S be the parties that did not abort in the second round. Now each party transforms the ciphertexts to C'_i with respect to \mathbf{pk}_S and performs the homomorphic evaluation for computing f . Then they perform the threshold decryption to obtain the output.

Intuitively, if a party aborts in the first round, then he is simply ignored. If he aborts in the second round, he is also ignored: since the other parties output flexible ciphertexts, these can be transformed to a public key representing the set of non-aborting parties. Those remaining parties can then proceed to perform the homomorphic computation. Finally, if a party aborts at the end, then it is too late – our threshold decryption algorithm only requires an honest majority of parties. We describe our three round protocol in Section 4.3.

Constructing TFHE. Our construction of TFHE is a distributed variant of the FHE scheme by Gentry, Sahai, and Waters [12]. We inherit from their scheme that our TFHE does not need the evaluation keys that Asharov et al. required, which allows for a cleaner presentation. We outline some of the technical aspects of our construction here, after we recall the GSW construction. The public key in their construction is a matrix \mathbf{B} and a vector $\mathbf{b} = \mathbf{B}\mathbf{s} + \mathbf{e}$ of the LWE form; the secret key is the LWE secret \mathbf{s} . To encrypt a bit m , the algorithm generates a random 0-1 matrix \mathbf{R} , and outputs $C = \text{Flatten}(m \cdot I_D + \text{BitDecomp}(\mathbf{R} \cdot \mathbf{b} \parallel \mathbf{R} \cdot \mathbf{B}))$, where I_D is the identity matrix. (We will define BitDecomp and BitDecomp^{-1} in Section 2, but, essentially, these functions act as their names suggest, decomposing a field element into a binary representation, and building a field element from a binary string.) To decrypt, the algorithm takes row β (where roughly $2^\beta >$ some noise bound) and parses the row into $(C_{\beta,1}, C_{\beta,2}) \in \mathbb{Z}_q^\ell \times \mathbb{Z}_q^{n \cdot \ell}$. (The parameters ℓ, n, q will be set in the scheme. Here for exposition, we can omit them.) Then it outputs

$$\left\lfloor \frac{\text{BitDecomp}^{-1}(C_{\beta,1}) - \langle \text{BitDecomp}^{-1}(C_{\beta,2}), \mathbf{s} \rangle}{2^\beta} \right\rfloor.$$

The homomorphic evaluation of the GSW scheme is surprisingly simple and beautiful! For addition it is $C + C'$ and for multiplication $C \cdot C'$.

As we discussed, our TFHE does not immediately determine a public key with respect to all parties, as done by Asharov et al. [1]. Instead, we set the public parameter (CRS) to be \mathbf{B} , and let each party P_i output $\text{pk}_i = \mathbf{b}_i = \mathbf{B}\mathbf{s}_i + \mathbf{e}_i$. Note that each $(\mathbf{B}, \mathbf{b}_i)$ is a GSW public key. The next challenge is how to generate flexible ciphertexts. A first natural idea would be: for P_i to generate a flexible ciphertext on some message m , P_i encrypts m under all GSW-type public keys $\{(\mathbf{B}, \mathbf{b}_i)\}_{[N]}$ to get $C = (C_1, \dots, C_N)$. To transform C with respect to a set S , we simply output $\{C_i\}_{i \in S}$. However, this is not secure since it allows every party, independently, to decrypt P_i 's ciphertext. A next idea would be P_i encrypts m under the key $(\mathbf{B}, \mathbf{b}_i)$ corresponding to his public key, and encrypts 0 for other keys $\{(\mathbf{B}, \mathbf{b}_j)\}_{j \neq i}$. The transform algorithm works the same. Intuitively, semantic security holds since P_i does not encrypt m under other people's keys. However, it is not clear how to jointly evaluate two ciphertexts from two parties, since the essential messages are encrypted under two different GSW public keys.

Our new idea to solve such challenge modifies C_j 's for $j \neq i$: instead of generating $\text{Enc}(0)$'s under $\{(\mathbf{B}, \mathbf{b}_i)\}_{[N] \setminus \{i\}}$, P_i outputs some *hints* for the transformation algorithm, but such hints will not hurt security. More specifically, we have the following design: P_i generates $C_i = \text{Flatten}(m \cdot I_D + \text{BitDecomp}(\mathbf{R} \cdot \mathbf{b}_i \parallel \mathbf{R} \cdot \mathbf{B}))$

and $C_j = \text{BitDecomp}(\mathbf{R} \cdot \mathbf{b}_j \parallel \mathbf{0})$, for $j \neq i$, where the same \mathbf{R} is used for all $\{C_j\}_{j \in [N]}$. Since each C_j only decreases the entropy of \mathbf{R} by $|\mathbf{R} \cdot \mathbf{b}_j|$, we can still use a leftover-hash-lemma style approach to argue that m is hidden.

Then given a set S (including i), we can compute $C_S = \sum_{j \in S} C_j$. By unfolding the equation, we can see:

$$C_S = \left(m \cdot I_D + \text{BitDecomp} \left(\mathbf{R} \cdot \left(\sum_{j \in S} \mathbf{b}_j \right) \parallel \mathbf{R} \cdot \mathbf{B} \right) \right),$$

which is of the form $\text{Enc}(m)$ under the GSW public key $(\mathbf{B}, \sum_{j \in S} \mathbf{b}_j)$! This means any flexible ciphertext, after being transformed, results in an encryption under the GSW public key. Therefore, ciphertexts from different parties can be jointly computed after transformed to ones with respect to the same set S .

Our threshold decryption protocol needs to work for any set S of participants such that $|S| > [N/2]$. So the parties should distribute the secret \mathbf{s}_i 's to all the other parties using a threshold secret sharing scheme. The challenging part is to design a *one-round* protocol. We use the fact that the decryption algorithm of the GSW scheme is essentially computing inner product (of a publicly known vector and the secret key), and Shamir's secret sharing scheme is highly compatible with inner product computation. In particular, each party P_i shares \mathbf{s}_i into $(\mathbf{p}_i(1), \mathbf{p}_i(2), \dots, \mathbf{p}_i(N))$ and sends $\mathbf{p}_i(j)$ to P_j , where \mathbf{p} is a vector of polynomials for Shamir's shares. To compute $w = \langle \mathbf{u}, \sum_{j \in S} \mathbf{s}_j \rangle$ for some publicly known vector \mathbf{u} (think of it as part of a ciphertext), each party can output $w_i = \langle \mathbf{u}, \sum_{j \in S} \mathbf{p}_j(i) \rangle$. Then it is not hard to see that these w_i 's form shares of w , so after receiving a majority of shares each party can run the reconstruction without interaction!

Finally, we need to handle an additional technicality to deal with noise of evaluated ciphertexts, as pointed out by Asharov et al. [1]. Intuitively, an evaluated ciphertext $\text{Enc}(f(x))$ might contain noise that is related to the original input x , so we need to add additional smudging noise to eliminate any such link. In the decryption protocol of the work [1], each party adds independent small noise to the output. However, this method will not work for our case because in our reconstruction procedure, these noise values are multiplied by the Lagrange coefficient, which can be too large. To solve this issue, we let each party P_i secret share some small noise η_i into $(r_i(1), \dots, r_i(N))$ and send the shares to the other parties (where r_i is a random polynomial for the shares). Then each party P_i adds $\sum_{j \in S} r_j(i)$ to their output. By the linearity of the Shamir's sharing scheme, this is equivalent to adding $\sum_{j \in S} \eta_j$ to the original reconstructed output value. In Section 4.2 we go through this construction in detail. The new TFHE may be of independent interests.

1.3 Related Work

There is a long line of work studying the round complexity of secure computation, both in the semi-honest and malicious models, the two-party and multi-party

settings, the honest majority and honest minority settings, and even in a variety of other models. We will not aim to survey all of this work, but mention what we know to be the best round complexity in the most relevant settings.

Constant round protocols have been known since Yao’s original two-round construction for the two-party, semi-honest setting [21], and Beaver et al.’s constant round protocol for the setting of a malicious minority [3]. In the two-party, malicious setting, Katz and Ostrovsky give a five-round protocol and demonstrate that this is tight [16]. There are several works demonstrating constant round protocols in the multiparty, malicious majority setting ([17, 19] Of course, with a malicious majority (including the two-party case), fairness is unachievable, so these results are in the security-with-abort model, and are not directly relevant to our own work.

In the multiparty setting with a malicious minority, the best known round complexity is achieved by the two-round protocol of Garg et al. [10], but, as we outlined above, their result does not ensure fairness. For $t < N/5$ corruptions, Damgård and Ishai give a three-round protocol with a guarantee of output delivery [8], though they require private point-to-point channels, and establishing these would add at least one additional round. For $t < N/2$, the exact round complexity of their protocol is a bit hard to discern, but it is greater than four (and we believe more); in this domain, the five-round protocol of Asharov et al. [1], which also guarantees output delivery, is the best known. For $t = 1$ corruption, Ishai et al. [15] showed that $N \geq 5$ parties are sufficient to securely compute general functionalities with guarantee of output delivery. The work [15] also showed 2-round protocols (guarantee of output delivery) for general functionalities in the server-client model, with a more restricted corruption pattern (e.g. one corrupted client and coalitions of $t < N/3$ servers). In the semi-honest, two-party setting, Yao’s original construction already achieves two-rounds.

Very recently and independent of this paper, Mukherjee and Wichs [] constructed 2-round protocols (in the CRS model) that achieve security with abort against any number of corruptions. In the setting of an honest majority, their protocol can be easily modified to achieve guarantee of output delivery in 3 rounds, assuming private communication channels, and in 4 rounds without private communication channels.

Gennaro et al. [11] provide a lower bound on the round complexity of fair protocols whenever $1 < t < N/2$. Our lower-bound strengthens theirs, ruling out even a fail-stop adversary that corrupts a single party.

2 Preliminaries

In this section, we present basic vector operations. Due to space limit, we describe the security definitions for MPC and the LWE assumptions in the full version of this paper [14].

2.1 Elementary Vector Operations

We define a number of vector/matrix operations that we describe below. Let \mathbf{a}, \mathbf{b} be vectors of dimension k . Let $\ell = \lfloor \log q \rfloor + 1$ for some modulus q . Note that the operations we describe are also defined over matrices, operating *row by row* on the matrix, and that all arithmetic is over \mathbb{Z}_q .

$\text{BitDecomp}(\mathbf{a})$ = the $k \cdot \ell$ dimensional vector $(a_{1,0}, \dots, a_{1,\ell-1}, \dots, a_{k,0}, \dots, a_{k,\ell-1})$ where $a_{i,j}$ is the j^{th} bit in the binary representation of a_i , with bits ordered from least significant to most significant.

$\text{BitDecomp}^{-1}(\mathbf{a}')$ For $\mathbf{a}' = (a_{1,0}, \dots, a_{1,\ell-1}, \dots, a_{k,0}, \dots, a_{k,\ell-1})$, let

$\text{BitDecomp}^{-1}(\mathbf{a}') = \left(\sum_{j=0}^{\ell-1} 2^j a_{1,j}, \dots, \sum_{j=0}^{\ell-1} 2^j a_{k,j} \right)$, but defined even when \mathbf{a}' isn't binary.

$\text{Flatten}(\mathbf{a}') = \text{BitDecomp}(\text{BitDecomp}^{-1}(\mathbf{a}'))$

$\text{Powersof2}(\mathbf{b}) = (b_1, 2b_1, 4b_1, \dots, 2^{\ell-1}b_1, \dots, b_k, \dots, 2^{\ell-1}b_k)$.

3 Impossibility Result

In this section, we are going to show that it is impossible to construct a two-round secure protocol for general multi-party computation with fairness, even with an honest majority of players. Our impossibility results holds in the *standalone* model, even with non-rushing fail-stop adversaries with access to a CRS. Our result strengthens that of Gennaro, as it holds even for adversaries corrupting only a single party, while their result cannot rule out the case where $t = 1$.

We assume that the players have both point-to-point channels and a public broadcast channel, but they do not have private point-to-point channels – an eavesdropper can listen to all channels.⁴ We note that our three round protocol from Section 4 can be collapsed into a three round protocol if we give the users access to a PKI of the appropriate form, so the assumption of non-private channels in our lower-bound is natural.⁵ A more formal proof follows.

Theorem 1. *Let \mathcal{C} be a family of circuits, and let Π be a polynomial-time, 2-round, 3-party secure protocol for computing $U(C, x, 0) = C(x)$ for any $C \in \mathcal{C}$, with fairness in the standalone model. Then there exists a virtual black-box obfuscator for general circuits⁶.*

Proof. We describe a VBB obfuscator \mathcal{O} for all circuits in \mathcal{C} . Before doing that, we define some notation and our next message functions. We let \mathcal{M} denote the set of valid messages in the secure computation. We let $\perp \in \mathcal{M}$ denote a special

⁴ Our lower bound holds even when the eavesdropper only listens to some channels.

⁵ Although we allow the eavesdropping adversary to corrupt two private channels at once, we do not allow it to corrupt the parties themselves, so we do still maintain an honest majority. However, there is still room to consider a weaker model where the eavesdropper can only listen to a single channel.

⁶ We describe the definition of VBB obfuscation in the full version of this paper [14].

abort symbol, and we let \emptyset denote the empty transcript (before any messages have been sent). A *partial incoming transcript*, is either \emptyset (if no messages have been sent yet), or of the form $(\mathcal{M}, \mathcal{M})$, where each message is received from one of the two other parties in the first round of the protocol. A *partial outgoing transcript* is of the same form, but represents the two messages sent by a single party in the first round, each going to one of the other parties. A *full incoming transcript* is of the form $((\mathcal{M}, \mathcal{M}), (\mathcal{M}, \mathcal{M}))$, where the first pair of messages are those received in the first round, and the second pair are those received in the second round. We define the following set of circuits.

$\pi_{i,j}(x, \tau, r)$: for parties $i, j \in \{1, 2, 3\}$, on input value x , partial incoming transcript τ and randomness r , the circuit outputs i 's next message to j .

$\pi_{\text{out}}(x, \tau, r_2)$: the circuit computes P_2 's output in the secure computation, given input x , full incoming transcript τ and randomness r_2 .

The VBB obfuscation of circuit C is as follows. $\mathcal{O}(C)$ chooses randomness r_1 and computes $\alpha_2 = \pi_{12}(C, \emptyset, r_1)$, $\alpha_3 = \pi_{13}(C, \emptyset, r_1)$. Note that they are the first-round messages from P_1 to P_2 and P_3 . Then the obfuscator outputs the following circuit $\Gamma_{\alpha_2, \alpha_3}(x; r_2, r_3)$, which, on input x and randomness r_2, r_3 , performs the following computations:

- $\gamma^{(1)} = \pi_{32}(\emptyset, \emptyset, r_3)$; $\beta = \pi_{23}(x, \emptyset, r_2)$. (The relevant first round messages.)
- $\gamma^{(2)} = \pi_{32}(\emptyset, (\alpha_3, \beta), r_3)$. (The relevant second round message.)
- Output $\pi_{\text{out}}(x, ((\alpha_2, \gamma^{(1)}), (\perp, \gamma^{(2)})), r_2)$. (P_2 's output, given his full incoming transcript.)

Basically, the circuit simulates P_2, P_3 's messages when P_1 sends out α_2, α_3 and then aborts in the second round.




Fig. 1. A depiction of the messages used in the circuit $\Gamma_{\alpha_2, \alpha_3}$. α messages are sent by party P_1 , β messages by P_2 , and γ messages by P_3 . A subscript i indicates that the recipient is P_i , and a superscript indicates a round number. Since we do not need all protocol messages, we drop subscripts and superscripts where we can.

We claim that for any $C \in \mathcal{C}$, $\Gamma_{\alpha_2, \alpha_3}$ is a secure VBB obfuscation of C . Efficiency of $\Gamma_{\alpha_2, \alpha_3}$ follows from the fact that the secure computation is polynomial-time. Correctness follows from the fairness of the underlying secure computation protocol – we note that by the correctness of the protocol, P_1 can learn the output after he sees all the incoming messages, regardless of whether or not he aborts in the second round. Thus by fairness of the protocol, P_2 should also receive the output, regardless of whether or not P_1 aborts in the second round. Thus, given the transcript $(x, ((\alpha_2, \gamma^{(1)}), (\perp, \gamma^{(2)})))$, P_2 can compute the output.

To prove the VBB property, recall that we need to prove that for any adversary $\mathcal{A}_{\mathcal{O}}$, for any circuit C , there exists a simulator $\mathcal{S}_{\mathcal{O}}$ and a negligible function ϵ such that

$$|\Pr[\mathcal{A}_{\mathcal{O}}(\Gamma_{\alpha_2, \alpha_3}) = 1] - \Pr[\mathcal{S}_{\mathcal{O}}^C(1^{|C|}) = 1]| < \epsilon(|C|)$$

By the security of the underlying secure computation, we know there exists an ideal-world simulator, which we will denote by $\mathcal{S}_{\mathbb{E}}$, that simulates the view of an eavesdropper who listens to the channels between P_1 and P_2 , and between P_1 and P_3 . We will denote by $\mathcal{S}_{\mathbb{E}}^{(1)}$ the result of running $\mathcal{S}_{\mathbb{E}}$ and restricting the output to the partial outgoing transcript, i.e. the first round messages sent from P_1 . Then, $\mathcal{S}_{\mathcal{O}}$ gets $(\tilde{\alpha}_2, \tilde{\alpha}_3) \leftarrow \mathcal{S}_{\mathbb{E}}^{(1)}$, and constructs the circuit $\Gamma_{\tilde{\alpha}_2, \tilde{\alpha}_3}$, as described above; we note that neither C nor r_1 are needed, once $\tilde{\alpha}_2$ and $\tilde{\alpha}_3$ are computed. Finally, $\mathcal{S}_{\mathcal{O}}$ outputs $\mathcal{A}_{\mathcal{O}}(\Gamma_{\tilde{\alpha}_2, \tilde{\alpha}_3})$.

Suppose that this does not meet the above security requirement. It follows that there exists a distinguisher \mathcal{D} that distinguishes between a real world execution of the protocol and the ideal simulation of $\mathcal{S}_{\mathbb{E}}$. This follows immediately, because the only difference between the true obfuscation and the simulated obfuscation is the way in which α_2 and α_3 are generated. Therefore, on input transcript τ , \mathcal{D} simply takes the messages α_2, α_3 that constitute the first round messages sent from P_1 to P_2 and P_3 respectively, and he completes the construction of $\Gamma_{\alpha_2, \alpha_3}$ himself. \mathcal{D} then runs $\mathcal{A}_{\mathcal{O}}$ on the resulting circuit and determines from the output whether τ was simulated.

Remark. The lower bound proof can be extended to rule out protocols in the CRS model, using the same idea. The obfuscated circuit will now embed crs as a common reference string, and in the security proof, the simulator will simulate the string, i.e. $(\tilde{\text{crs}}, \tilde{\alpha}_2, \tilde{\alpha}_3) \leftarrow \mathcal{S}_{\mathbb{E}}^{(1)}$.

Two Round Feasibility with a PKI. Note that the proof breaks down if the parties have access to private channels, including in the scenario where they have access to a PKI. This is because we need *both* first-round messages sent from P_1 in order to simulate round two of the protocol. In particular, without access to α_3 , we could not correctly simulate $\gamma^{(2)}$ (as sent from P_3 to P_2 in round two), and therefore we could not guarantee the correct output of the obfuscation. The only way to gain access to both α_2 and α_3 is either by eavesdropping on multiple channels, or by corrupting two parties, but this latter approach would violate our assumption of an honest majority. Indeed, as we mentioned previously, and

as we will see later, our construction in Section 4 can be collapsed to two rounds if we have access to a PKI, with public keys of a particular form. It is still an open question whether a two-round protocol with guaranteed output delivery is possible, given access only to private channels.

4 Towards Fairness and Guarantee of Output Delivery

The previous section shows that two-round fair protocols are in general impossible. As discussed in the introduction, we can construct a three-round fair protocol by adding one more round to the protocol by Garg et al. [10]; yet it is unclear how to construct three-round protocols with guarantee of output delivery. In this section, we present our main contribution – we construct a new threshold FHE scheme, which extends the notion of threshold FHE by Asharov et al. [1] with enriched features. We elaborate on these below.

4.1 New Threshold Fully Homomorphic Encryption Scheme

As discussed in the introduction, our TFHE introduces a new idea of *flexible* and *transformed* ciphertexts that play an important role in our 3-round MPC construction. Here we first present the syntax: a threshold fully homomorphic encryption scheme (TFHE) is basically a homomorphic encryption scheme, with the difference that the key generation and decryption are N -party protocols instead of algorithms. We will consider protocols defined in terms of some common parameter pp .

- $\text{TFHE.Gen}(\text{pp})$ (**Key Generation Protocol.**) Initially each party holds some parameter pp . At the conclusion of the protocol, each party P_i for $i \in [N]$ publishes a public key pk_i , and keeps a private key sk_i .
- $\text{TFHE.Dec}_S(C; \mathbf{v})$ (**Threshold Decryption Protocol.**) Let S be a set in $[N]$, and $\mathbf{v} = \{v_i : i \in S\}$ be some secret values, each held by one party. The protocol is run among parties $\{P_i : i \in S\}$. Initially each party holds a secret input $\mathbf{v}[i]$, a secret key sk_i , and receives a ciphertext C as the public input. At the end, the parties in the set can compute the decrypted message m . Intuitively, the secret input \mathbf{v} is used for smudging the noise. Note: in the setting with honest majority, we assume that $|S| \geq [N/2] + 1$. For simplicity, we assume the input ciphertext C has already been transformed to one that corresponds to the set of public keys $\text{pk}_S = \{\text{pk}_i : i \in S\}$. See the syntax below for further exposition.
- $\text{TFHE.Enc}_i(\text{pp}, \text{pk}_1, \dots, \text{pk}_N; m)$ (**Encryption Algorithm.**) Let parties $\{P_i\}_{i \in [N]}$ participate in the protocol, and $\{\text{pk}_i\}_{i \in [N]}$ be the set of their public keys. The encryption algorithm is non-interactive and run by party P_i . The algorithm takes inputs the public parameter, the public keys $\{\text{pk}_i\}_{i \in [N]}$, a message m , and computes a ciphertext C .

We implicitly require that the ciphertexts here are **flexible** in the sense that they do not *commit* to a particular public key/secret key yet; in particular, we can use the algorithm below to transform a flexible ciphertext into one that corresponds to a set of public keys.

- **TFHE.TransCT($C; S$) (Ciphertext Transform Algorithm.)** The algorithm takes inputs a *flexible* ciphertext C (from the above encryption algorithm), and a set $S \subseteq [N]$ and outputs a *transformed* ciphertext C_S . The ciphertext can be thought as one under the set of joint keys: $\text{pk}_S = \{\text{pk}_i : i \in S\}$.
- **TFHE.Eval($f, C_1, \dots, C_t; S$) (Evaluation Algorithm.)** The evaluation algorithm is non-interactive. A party P_i (can be any party) receives inputs a function $f : \{0, 1\}^t \rightarrow \{0, 1\}$, flexible ciphertexts C_1, \dots, C_t , and a set $S \subseteq [N]$. He computes an evaluated ciphertext C'_S with respect to the set S , which can be thought as an evaluated ciphertext under the joint public key pk_S defined as above.

We summarize the main differences between our TFHE and that of the prior work [1].

1. Our key generation does not output a *joint* public key. Instead each party will only output their own public key pk_i . Then parties can run the Eval algorithm to homomorphically compute on the ciphertexts under pk_S for some set S decided later. As pointed out in the introduction, this is an important feature.
2. The construction of the prior work requires all parties to participate in the decryption protocol (in the non-interactive case). Here we allow a subset of parties to run the protocol; moreover, we allow a “threshold” type of decryption where a majority of parties can decrypt the ciphertext.

These new features play an important role: intuitively, when a party generates a ciphertext, he does not know who else might abort. The flexibility of ciphertexts handles this problem – the parties can generate ciphertexts first, and later on decide a set of public key (namely $\text{pk}_S = \{\text{pk}_i : i \in S\}$), so that the flexible ciphertexts can be transformed with respect to pk_S . Then the parties can perform homomorphic computation with respect to pk_S and run the threshold decryption algorithm.

Similar to the work [1], we do not define the security of TFHE on its own. The reason is similar: requiring that the above protocols securely realize some ideal key-generation and decryption functionalities is unnecessarily restrictive. Instead, we will show that our TFHE scheme is secure directly in the context of our implementation of general MPC in Section 4.3.

4.2 Construction of Our New TFHE

Following the intuition in the introduction, we describe our construction.

Common Parameter. All parties receive the common parameter pp of the form: let N be the number of parties, $L = \text{poly}(\kappa)$ be the maximum depth of the circuits supported by the TFHE evaluation algorithm. Then we choose a modulus q of $\text{poly}(L, N)$ bits, lattice dimension parameter $n = n(L, N)$, and error distribution $\chi = \chi(\kappa, L, N)$ appropriately for LWE security against 2^κ known attacks. Also, choose parameter $m = m(\kappa, L) = O((n + N) \log q)$. Let

the distribution χ be B_χ -bounded (i.e. with overwhelming probability, a sample from χ has the absolute value less than B_χ). Let $\ell = \lceil \log q \rceil + 1$, $D = (n + 1) \cdot \ell$, and $B_{\text{smug}} \in \mathbb{Z}$ be an integer bound, satisfying the following relations:

$$\frac{(D + 1)^L \cdot N \cdot B_\chi}{B_{\text{smug}}} = \text{negl}(\kappa), \quad B_{\text{smug}} < q/8.$$

Then $\text{pp} = (n, q, \chi, m, B_\chi, B_{\text{smug}}, \mathbf{B})$ where \mathbf{B} is sampled uniformly from $\mathbb{Z}_q^{m \times n}$.

TFHE.Gen(pp): This is a two-round protocol among N parties.

- **(Round 1)**: Each party P_i samples a random vector $\mathbf{s}_i \in \mathbb{Z}_q^n$, and computes $\mathbf{b}_i = \mathbf{B} \cdot \mathbf{s}_i + \mathbf{e}_i$ where $\mathbf{e}_i \leftarrow \chi^m$. Then P_i broadcasts $\text{pk}_i = \mathbf{b}_i$, and keep \mathbf{s}_i secretly.
- **(Round 2)**: Each party P_i secret shares \mathbf{s}_i using the Shamir Secret Sharing Scheme with threshold $\lfloor N/2 \rfloor + 1$. Let \mathbf{p}_i denote the random polynomial vector (of degree $\lfloor N/2 \rfloor + 1$) generated by P_i where $\mathbf{p}_i(0) = \mathbf{s}_i$ (This is how the Shamir Secret Sharing works). P_i sends $\mathbf{p}_i(j)$ to P_j for $j \in [N]$. At the end, P_i sets $\text{sk}_i = (\mathbf{p}_1(i), \mathbf{p}_2(i), \dots, \mathbf{p}_N(i))$.

Note that although we do not assume secure point-to-point channels, sending private message in the second round is achievable – everyone can send a public key in the first round, and later on every party encrypts the outgoing messages. For simplicity, we just assume there are secure point-to-point channels available in the second round.

TFHE.Dec_S(C; v): Let \mathbf{v} be a vector of $|S|$ numbers (error terms), where $\mathbf{v}[i]$ (the element indexed by i) is held by party P_i for $i \in S$; let C be a ciphertext. For $S \subseteq [N]$ such that $|S| \geq N/2$, this is a one-round protocol among parties $\{P_i : i \in S\}$. For simplicity, we assume that C is a transformed ciphertext that corresponds to pk_S .

- Each party P_i parses C as a matrix in $\mathbb{Z}_q^{D \times D}$. Then he picks the β -th row, C_β , where $\beta = \lfloor \log_2(q/2) \rfloor$. Note that $2^\beta \in (q/4, q/2]$. Then parse $C_\beta = (C_{\beta,1}, C_{\beta,2})$ where $C_{\beta,1} \in \mathbb{Z}_q^\ell$, $C_{\beta,2} \in \mathbb{Z}_q^{n \cdot \ell}$. Then he computes $\mathbf{z}_i = \sum_{j \in S} \mathbf{p}_j(i)$ and broadcasts $w_i = \langle \text{BitDecomp}^{-1}(C_{\beta,2}), \mathbf{z}_i \rangle + \mathbf{v}[i]$.
- At the end, each party picks an arbitrary subset $T \subseteq S$ such that $|T| = \lfloor N/2 \rfloor + 1$. Then they compute $w = \sum_{k \in T} \mu_k(0) w_k$, where μ_k is the Lagrange polynomial. Finally they output $\left\lfloor \frac{\text{BitDecomp}^{-1}(C_{\beta,1}) - w}{2^\beta} \right\rfloor$.

TFHE.Enc_i(pp, pk₁, ..., pk_N; m): This is the i -th party's encryption algorithm. Let $m \in \{0, 1\}$ be the input message.

- The algorithm parses pp as a matrix $\mathbf{B} \in \mathbb{Z}_q^{m \times n}$, $\text{pk}_j = \mathbf{b}_j \in \mathbb{Z}_q^m$ for $j \in [N]$. Then it samples a random matrix $\mathbf{R} \in \{0, 1\}^{D \times m}$, and computes $\mathbf{W}_j = \text{BitDecomp}(\mathbf{R} \cdot \mathbf{b}_j \parallel 0^{D \times n})$ for $j \neq i$. It computes $\mathbf{W}_i = \text{Flatten}(m \cdot I_D + \text{BitDecomp}(\mathbf{R} \cdot \mathbf{b}_i \parallel \mathbf{R} \cdot \mathbf{B}))$, where I_D is the identity matrix of dimension $D \times D$. It outputs $C = (\mathbf{W}_1, \dots, \mathbf{W}_N)$.

TFHE.TransCT($C; S$):

- The algorithm parses C as a N matrices $(\mathbf{W}_1, \dots, \mathbf{W}_N)$. It outputs $C_S = \sum_{j \in S} \mathbf{W}_j$.

TFHE.Eval($f, C_1, \dots, C_t; S$):

- For simplicity, we assume that all the ciphertexts C_1, \dots, C_t 's are transformed to ones that correspond to \mathbf{pk}_S (otherwise we can apply the above TFHE.TransCT first). We then observe that actually a transformed ciphertext is of the same form of the GSW scheme [12] where the public key is $(\sum_{k \in S} \mathbf{b}_k \parallel \mathbf{B})$. Thus, we can run exactly the same evaluation as the GSW scheme! More specifically, we represent f as a circuit (with all NAND gates). Then we can homomorphically compute $NAND(C, C')$ by outputting $\text{Flatten}(I_D - C \cdot C')$. See the work [12] for detailed explanation.

With out setting of parameters, we can argue that flexible ciphertexts do not leak the underlying messages to the other parties (and so do the transformed ciphertexts, since they can be obtained deterministically from flexible ciphertexts). This can be shown formally using the lemma below in a strait-forward way as done by the work [12]. See their work [12] for further exposition⁷.

Lemma 1 (Implicit in [20]). *Let n, m, χ, q be parameters such that the LWE $_{n,q,\chi}$ holds, and N be some polynomial. Then for $m = O((n + N) \log q)$, for any vectors $\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_{N-1} \in \mathbb{Z}_q^m$, then the distribution described as above $(\mathbf{B}, \mathbf{b}, \mathbf{R} \cdot (\mathbf{B} \parallel \mathbf{b}), \mathbf{R} \cdot (\mathbf{b}_1 \parallel \dots \parallel \mathbf{b}_{N-1}))$ is computationally indistinguishable from $(\mathbf{B}, \mathbf{u}, \mathbf{U}, \mathbf{R} \cdot (\mathbf{b}_1 \parallel \dots \parallel \mathbf{b}_{N-1}))$, where \mathbf{B} is uniform over $\mathbb{Z}_q^{m \times n}$, \mathbf{u} is uniform over \mathbb{Z}_q^m , \mathbf{U} is uniform over $\mathbb{Z}_q^{D \times (n+1)}$, and \mathbf{R} is uniform over $\{0, 1\}^{D \times m}$, $D = (n + 1) \cdot \ell$, $\ell = \lceil \log q \rceil + 1$.*

4.3 Three-round MPC with Guarantee of Output Delivery

Now we are ready to present our new three-round MPC for general functions using the new TFHE we have developed in the previous section. We first present a simpler case that considers MPC for polynomial-time deterministic boolean function f (where all parties receive the same bit). Moreover, the security holds against static *semi-malicious* fail-stop attackers⁸ corrupting less than half of the parties. In Section 5, we discuss how to handle general cases using standard techniques.

⁷ Our setting of parameters is slightly different from that of the work [12], so our parameters in the lemma are slightly different. The analysis is essentially identical.

⁸ Basically, a semi-malicious attacker is one whose behavior follows the protocol with *some* input and randomness he must know. Protocols that achieve such security can be upgraded to malicious security without adding a coin flipping round (c.f. Section 5). See the full version of this paper [14] for further details about the notion and its advantage.

Remark 1. Our protocol only needs a public broadcast channel⁹. For simplicity of presentation, we make the following two assumptions. First, there are secure point-to-point channels available. Second, when a party distributes shares to the other parties, he must either send messages to all parties or send messages to no one. These assumptions are not necessary, and we sketch how to achieve them in our protocol using the broadcast channel. We observe that our protocol will only use the secure channels to distribute shares in the second round. So in the first round everyone can publish a public key, and then in the second round, everyone broadcasts encryptions of the shares (under different parties' public keys). This can implement the secure channels, and ensure that parties will either abort (not broadcast at all) or distribute messages to all the other parties.

Our Construction. Let $f : \{0, 1\}^{(\ell_{\text{in}})^N} \rightarrow \{0, 1\}$ be a function computed by a depth L circuit, where ℓ_{in} is the input length of each party.

Input: Each party P_i holds some input $x_i \in \{0, 1\}^{\ell_{\text{in}}}$. The parties share the public parameter \mathbf{pp} as described in the TFHE scheme. (\mathbf{pp} can be viewed as the common reference string. The generation of \mathbf{pp} depends on L , since we need the TFHE to support circuits up to depth L).

The Protocol:

- **Round 1:** The parties execute the first round of the $\text{TFHE.Gen}(\mathbf{pp})$. If anyone aborts in this round, then he is simply ignored. Let $S_1 \subseteq [N]$ be the set of non-aborting parties at this round. At the end of this round, each party holds all $\{\mathbf{pk}_i\}_{i \in S_1}$.
- **Round 2:** The parties execute the following procedures at the same time:
 - The (currently non-aborting) parties execute the second round of the $\text{TFHE.Gen}(\mathbf{pp})$.
 - For $i \in S_1$, P_i broadcasts an encryption of his input using the algorithm $\text{TFHE.Enc}_i(x_i)$ (encrypt it bit-by-bit). Note that these are a flexible ciphertxts.
 - Each P_i samples a uniformly random error term from $\eta_i \leftarrow [-B_{\text{smug}}, B_{\text{smug}}]$, and compute random Shamir secret shares (with the same threshold $[T/2] + 1$). Denote the polynomial as r_i (note that $r_i(0) = \eta_i$). Then each P_i sends $r_i(j)$ to party P_j for $j \neq i$.

Let $C_i = (C_{i,1}, C_{i,2}, \dots, C_{i,\ell_{\text{in}}})$ be the broadcasted ciphertxts from P_i , and $(r_i(1), \dots, r_i(N))$ be the shares from P_i to the other parties.

If anyone aborts at this round, either not sending the second round of $\text{TFHE.Gen}(\mathbf{pp})$, the ciphertxts, or the shares of error terms, then he (and his input) are again ignored. Let $S_2 \subseteq S_1$ be the set of non-aborting parties.

- **Round 3:** Now each non-aborting party in S_2 first transforms the ciphertxts he received to ones that correspond to \mathbf{pk}_{S_2} . Let $\{C_{j,k}\}_{j \in S_2, k \in [\ell_{\text{in}}]}$ be the

⁹ In the semi-malicious setting, this can be easily implemented by reliable public point-to-point channels, where an eavesdropper can listen to the all channels but cannot modify the messages

broadcasted ciphertexts. For $i \in S_2$, P_i first computes $C_{j,k}^{S_2} = \text{TFHE.TransCT}(C_{j,k}; S_2)$ for $j \in S_2, k \in [\ell_{\text{in}}]$.

Let f^{S_2} be the residual function where the inputs of $[N] \setminus S_2$ are replaced with the default values. P_i homomorphically computes the residual function, i.e. $C^* = \text{TFHE.Eval}(f^{S_2}, \{C_{j,k}^{S_2}\}_{j \in S_2, k \in [\ell_{\text{in}}]})$.

Then each P_i computes $v_i = \sum_{k \in S_2} r_k(i)$. Finally, they run the threshold decryption $\text{TFHE.Dec}_{S_2}(\{C^*; \mathbf{v}_{S_2}\})$, where \mathbf{v}_{S_2} denotes the vector of the following set $\{v_j : j \in S_2\}$.

Recall that the protocol TFHE.Dec handles situations when parties abort. In this round, parties broadcast some messages, and a majority of them is sufficient to recover the output.

Theorem 2. *Let f be any deterministic functionality with N inputs and one output. Let \mathbf{pp} be parameters sampled according to the choice as the TFHE above, and the corresponding LWE assumption holds. Then the above protocol π UC-realizes the ideal functionality \mathcal{F}_f with guarantee of output delivery, in the presence of any static (semi-malicious) fail-stop adversary who corrupts less than $\lfloor N/2 \rfloor$ parties.*

As explained in the introduction, the transformed ciphertexts $\{C_{j,k}^{S_2}\}_{j \in S_2, k \in [\ell_{\text{in}}]}$ are GSW ciphertexts under the public key $(\mathbf{B}, \sum_{i \in S_2} \mathbf{b}_i)$. Therefore, by applying the evaluation algorithm, C^* is a ciphertext of the output y . Each party in our threshold decryption protocol, as explained, outputs a share of y by computing some inner product with the shares (and subtraction). Thus, the correctness holds.

To prove security, we need to construct a simulator \mathcal{S} that generates the views of the honest parties. We sketch the construction: the simulator simulates the public parameter faithfully, and generates the messages in each round as follows. Let \mathcal{I} be the set of corrupted set.

- **(First round).** \mathcal{S} simulates the public keys of honest parties' by random vectors \mathbf{u}_i for $i \notin \mathcal{I}$.
- **(Second round).** \mathcal{S} simulates the encrypted ciphertexts by $\text{TFHE.Enc}(0)$, and simulates the error terms and shares of secret keys by sending random values (or vectors).
- **(Third round).** \mathcal{S} then reads the witness tapes of the adversary to get secret keys and inputs from the corrupted parties. He sets the aborting parties' inputs to be the default value, and then queries the ideal functionality to receive the output y . From the output y and the secret keys of the corrupted parties, \mathcal{S} then figures out consistent outputs of the honest parties.

Intuitively, the LWE assumption guarantees that the simulation in round 1 is indistinguishable, and Lemma 1 guarantees that $\text{TFHE.Enc}(0)$ is indistinguishable from the encryptions in the real world. The last step is the most challenging, and we will further explain the ideas in the appendix.

In the full version of this paper [14], we present the detailed analyses of correctness and security with further exposition.

5 Variants and Generalizations

In this section, we discuss variants of our basic protocol in the following aspects: (1) how to handle functionalities with longer inputs, (2) how to handle randomized functionalities, (3) how to compile a protocol that is secure against semi-malicious adversaries into one that is secure against malicious adversaries, and (4) how to reduce one round by using a PKI setup. These issues can be handled using standard techniques as presented in the work of Asharov et al. [1]. We highlight the ideas and refer curious readers to their work for further details.

Functions with longer outputs. Let $f : \{0, 1\}^{(\ell_{\text{in}})^N} \rightarrow \{0, 1\}^{\ell_{\text{out}}}$ be an N -ary functionality. We consider ℓ_{out} boolean functionalities $\left\{ f_i : \{0, 1\}^{(\ell_{\text{in}})^N} \rightarrow \{0, 1\} \right\}_{i \in [\ell_{\text{out}}]}$

where each f_i outputs the i -th bit of f . Let π_i be the protocol computing f_i as we described in Section 4. To compute f , we simply run $\pi_1, \dots, \pi_{\ell_{\text{out}}}$ in parallel, and we treat an abort in any one of the execution as an abort in all executions. To argue that the resulting protocol is secure against an arbitrary semi-malicious adversary, we also require the adversary to include proofs, in the form of witnesses written to their witness tape, of input-consistency across the parallel executions. This is to enforce that the adversary is using the same inputs for all the subprotocols. Below we will describe a compiler that upgrades the protocol to one against malicious adversaries.

Randomized functionalities. Our basic MPC protocol only considers deterministic functionalities where all the parties receive the same output. It can be generalized to handle with randomized functionalities and individual outputs via a standard transformation. Basically in this transformation, instead of computing some randomized function $f(x_1, \dots, x_N; r)$, the parties compute the deterministic function $f'((x_1, r_1), \dots, (x_N, r_N)) = f(x_1, \dots, x_N; \oplus_{i \in [N]} r_i)$. This transformation does not add additional rounds.

Semi-malicious security to malicious security. Our basic MPC protocol is only secure in the semi-malicious setting. Asharov et al. [1] presents a simple and general round-preserving compiler from semi-malicious to fully malicious security using UC NIZKs [9] in the CRS model. In particular, in each round, the attacker must prove (in zero-knowledge) that it is following the protocol consistently with *some* setting of the random coins. In particular, we present the theorem of Asharov et al. [1]:

Theorem 3 ([1]). *There is a generic round-preserving compiler such the following holds. Let \mathcal{F} be an N -ary functionality and π be an N -party protocol. Suppose π t -securely computes \mathcal{F} against semi-malicious fail-stop adversaries with guarantee of output delivery (or fairness), then the compiled protocol π' t -securely computes \mathcal{F} against malicious adversaries with guarantee of output delivery (or fairness, respectively) in the CRS, \mathbb{F}_{ZK} , and authenticated broadcast-hybrid model. Moreover, π' has the same round complexity as π .*

Together with Theorem 2, we are able to achieve the following corollary:

Corollary 1. *Assume that the LWE assumption holds and UC-NIZK exists. Then there exists a three-round MPC in the CRS and authenticated broadcast hybrid model, with a guarantee of output delivery, and providing security against a malicious adversary that corrupts less than half of the parties.*

Two rounds with PKI. We recall that in the first round of our protocol, each party just publishes some public key $\mathbf{b}_i = \mathbf{B} \cdot \mathbf{s}_i + \mathbf{e}_i$, which is independent of the input. If there is an additional setup *public-key infrastructure* (PKI), then we can move the first round to the PKI. Thus the entire MPC execution would consist only of the remaining two rounds. The resulting PKI is very simple and does not require a trusted party for setup; we just need a trusted party to choose a CRS, and then each party can choose its own public key individually (possibly maliciously). Moreover, the PKI can be reused for many MPC executions of arbitrary functions f with arbitrary inputs.

The security analysis is exactly the same as that of our original three-round protocol in the CRS model, just by noting that the first round there consists of broadcast message, which does not depend on the inputs of the parties (and hence we can think of it as a public key). In the malicious case, the parties need to provide a zero-knowledge proof of knowing some randomness of their public keys registered in the PKI. This is similar to our original protocol (without PKI) where the parties need to provide a zero-knowledge proof of knowing some randomness of their first round messages.

Acknowledgments

This research is partially supported by an NSF grant CNS-1314857, a Sloan Fellowship, and Google Research Awards.

References

1. G. Asharov, A. Jain, A. López-Alt, E. Tromer, V. Vaikuntanathan, and D. Wichs. Multiparty computation with low communication, computation and interaction via threshold FHE. In D. Pointcheval and T. Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 483–501. Springer, Apr. 2012.
2. B. Barak, O. Goldreich, R. Impagliazzo, S. Rudich, A. Sahai, S. P. Vadhan, and K. Yang. On the (im)possibility of obfuscating programs. In J. Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 1–18. Springer, Aug. 2001.
3. D. Beaver, S. Micali, and P. Rogaway. The round complexity of secure protocols (extended abstract). In *22nd ACM STOC*, pages 503–513. ACM Press, May 1990.
4. M. Ben-Or, S. Goldwasser, and A. Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In *20th ACM STOC*, pages 1–10. ACM Press, May 1988.
5. R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145. IEEE Computer Society Press, Oct. 2001.

6. R. Cleve. Limits on the security of coin flips when half the processors are faulty (extended abstract). In *Proceedings of the 18th Annual ACM Symposium on Theory of Computing, May 28-30, 1986, Berkeley, California, USA*, pages 364–369, 1986.
7. R. Cohen and Y. Lindell. Fairness versus guaranteed output delivery in secure multiparty computation. In P. Sarkar and T. Iwata, editors, *ASIACRYPT 2014, Part II*, volume 8874 of *LNCS*, pages 466–485. Springer, Dec. 2014.
8. I. Damgård and Y. Ishai. Constant-round multiparty computation using a black-box pseudorandom generator. In V. Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 378–394. Springer, Aug. 2005.
9. A. De Santis, G. Di Crescenzo, R. Ostrovsky, G. Persiano, and A. Sahai. Robust non-interactive zero knowledge. In J. Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 566–598. Springer, Aug. 2001.
10. S. Garg, C. Gentry, S. Halevi, and M. Raykova. Two-round secure MPC from indistinguishability obfuscation. In Y. Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 74–94. Springer, Feb. 2014.
11. R. Gennaro, Y. Ishai, E. Kushilevitz, and T. Rabin. On 2-round secure multiparty computation. In M. Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 178–193. Springer, Aug. 2002.
12. C. Gentry, A. Sahai, and B. Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In R. Canetti and J. A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 75–92. Springer, Aug. 2013.
13. O. Goldreich. *Foundations of Cryptography: Basic Applications*, volume 2. Cambridge University Press, Cambridge, UK, 2004.
14. S. D. Gordon, F.-H. Liu, and E. Shi. Constant-round mpc with fairness and guarantee of output delivery. Cryptology ePrint Archive, Report 2015/371, 2015.
15. Y. Ishai, E. Kushilevitz, and A. Paskin. Secure multiparty computation with minimal interaction. In T. Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 577–594. Springer, Aug. 2010.
16. J. Katz and R. Ostrovsky. Round-optimal secure two-party computation. In M. Franklin, editor, *CRYPTO 2004*, volume 3152 of *LNCS*, pages 335–354. Springer, Aug. 2004.
17. J. Katz, R. Ostrovsky, and A. Smith. Round efficiency of multi-party computation with a dishonest majority. In E. Biham, editor, *EUROCRYPT 2003*, volume 2656 of *LNCS*, pages 578–595. Springer, May 2003.
18. A. López-Alt, E. Tromer, and V. Vaikuntanathan. On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In H. J. Karloff and T. Pitassi, editors, *44th ACM STOC*, pages 1219–1234. ACM Press, May 2012.
19. R. Pass. Bounded-concurrent secure multi-party computation with a dishonest majority. In L. Babai, editor, *36th ACM STOC*, pages 232–241. ACM Press, June 2004.
20. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In H. N. Gabow and R. Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005.
21. A. C.-C. Yao. How to generate and exchange secrets (extended abstract). In *27th FOCS*, pages 162–167. IEEE Computer Society Press, Oct. 1986.