A Framework for Outsourcing of Secure Computation
(Draft Version)

Thomas P. Jakobsen, Jesper Nielsen, and Claudio Orlandi

Aarhus University, Denmark, {tpj,jbn,orlandi}@cs.au.dk

Abstract We study the problem of how to efficiently outsource a sensitive computation on secret inputs to a number of untrusted workers, under the assumption that at least one worker is honest.
In our setting there is a number of clients $C_1, \ldots, C_n$ with inputs $x_1, \ldots, x_n$. The clients want to delegate a secure computation of $f(x_1, \ldots, x_n)$ to a set of untrusted workers $W_1, \ldots, W_m$. We want do so in such a way that as long at there is at least one honest worker (and everyone else might be actively corrupted) the following holds 1) the privacy of the inputs is preserved 2) the output of the computation is correct (in particular workers cannot change the inputs of honest clients). We propose a solution where the clients' work is minimal and the interaction pattern simple (one message to upload inputs, one to receive results), while at the same time reducing the overhead for the workers to a minimum. Our solution is generic and can be instantiated with any underlying reactive MPC protocol where linear operations are "for free". In contrast previous solutions were less generic and could only be instantiated for specific numbers of clients/workers.

1 Introduction

In this paper we will use the term Secure Multiparty Computation (MPC) to refer to any problem where a number of parties wants to compute a function f on inputs $x_1, \ldots, x_n$ while guaranteeing interesting security properties such as the privacy of the inputs and the correctness of the outputs. In particular we will consider the setting where n parties (the clients) provide inputs and receive outputs, in the presence of m additional parties (the workers) who act as helpers to reduce the computational burden on the clients. Clients do not trust each other, and they wish to trust the workers as little as they have to.

A notable example is the case of *verifiable delegation of computation* [GGP10,PHGR13,BSCG+13,BFR13, BFR13] where one (or more) computationally bounded clients want to perform a computation on an untrusted cloud provider, and therefore wish to perform this computation in a way that the work required to verify the correctness of the result is much less than the work needed to compute the function itself, while also protecting the privacy of the inputs. Traditionally, the problem of verifiable delegation of computation is studied in the presence of a single untrusted worker. However in this case the only known way of protecting the privacy of the inputs is by using fully-homomorphic encryption schemes. This introduces a huge computational overhead for the worker. (If one is interested *only* in verifying the correctness of the output, recent studies show that using SNARKs (succinct non-interactive arguments of knowledge) can be made much more practical than expected a few years ago [PHGR13, BSCG+13]).

Another important application is *large-scale* secure computation, where one wants to run a secure computation on thousands or millions of secure inputs. In this setting a (significant) number of clients $C_1, \ldots, C_n$ with inputs $x_1, \ldots, x_n$, want to securely evaluate $f(x_1, \ldots, x_n)$. However running any existing MPC protocols for general functionalities between all the clients would require that all parties are online at the same time [HLP11], and the communication overhead of every practical protocol for dishonest majority scales quadratically with the number of parties. Instead, the clients can delegate their computation to a (small) set of untrusted workers $W_1, \ldots, W_m$. This is a relevant scenario in practice, and many real-world uses of secure computation follow this paradigm e.g., the Danish sugar beet auction [BCD+09], Sharemind [BLW08], MEVAL [CMF+14], etc. A limitation of these solutions is that they require a majority of the workers to be honest and only guarantee security against passive corruptions (in particular, a dishonest worker can arbitrarily change the input of an honest client).

Here instead we want to tolerate that all but one of the workers might be corrupted: this allows to use less workers to achieve the same security, which might be important in practice as the main cost of the system (probably) will be the price to rent computing time on the workers. Since we want to tolerate that all but one of the workers can be corrupted, we cannot use a protocol that guarantees termination. In fact, if we want to tolerate that all but one worker might be corrupted, it must provably be the case that a single worker can deadlock the system. This, however, can be detected and then other workers can be rented next time. However, our protocol guarantees termination whenever all workers are honest, independently of how many clients are corrupted.

There is a lot of prior work looking at this and related problems, both in terms of concrete [KMR11, KMR12, PTK13, CLT14, KMRS14] and asymptotic efficiency [Gen09, BV11, LTV12, GHRW14]. We will compare to related work of the first kind after presenting our protocol. The latter kind of work heavily relies on advanced cryptographic tools such as fully-homomorphic encryption: while this "Swiss Army knife" of cryptography allows for wonderful and surprising results in terms of feasibility and asymptotic complexity, it introduces a huge computational overhead for the workers and therefore it is worth studying alternative solutions that can be used in practice.

2 Technical Overview

We describe here the main idea of our framework. It will be instructive to think of a simple 3-party setting *a la* Salus [KMR12], where a client, running on a computationally limited device (e.g., a mobile phone) wants to engage in a computation with a server, and outsource most of the work to a worker (e.g., a cloud provider) which he does not fully trust. The client has input x and the server has input y. At the end of the protocol the client is supposed to learn $z = f(x, y)$ for some function f agreed upon by the parties (and nothing else about y). The server and the worker should not learn anything.

Ideally, we would like a protocol that satisfies the security requirements even if all but one party are (actively) corrupted. At the same time, as the client is computationally limited, we want to make sure that the work performed by the client is minimal – in particular *independent of the size of the function to be computed*. As already discussed, this is possible (and with optimal asymptotic efficiency) using fully-homomorphic encryption. However this will incur a huge computational overhead for the worker and the server. So, following the approach of [KMR11, KMR12, PTK13, CLT14] we seek for a protocol where the client has to trust that at least one among the worker and the server is honest.

Instead of designing a specific protocol to solve the problem, we propose a more generic approach to this problem, which can be instantiated using different building blocks depending on the particular application. This gives more flexibility and allows for a greater range of applications (for instance, solutions based on garbled circuits are typically limited to two parties).

Our main building block will be a protocol for reactive secure computation (that is, a protocol where it is possible to open intermediate values) and where linear operations are for free. Many protocols of this kind are known (e.g., [DO10, BDOZ11, NNOB12, DPSZ12, DKL+13])[1].

It is clear that the overall efficiency will be highly impacted by the efficiency of the underlying protocol, and in this paper we do not try to improve on this (but there is plenty of ongoing research on the subject). Instead, we consider *only* the (somewhat orthogonal) problem of how to let clients provide inputs to the workers in such a way that the clients' work is minimal and the overhead induced on the underlying MPC protocol is as limited as possible. We believe this modular approach is useful, both from a conceptual point of view, and also from a practical point of view e.g., one can imagine that improvements on the underlying MPC protocols for the workers would not require one to update the software on the client side.

2.1 A Simple but Inefficient Solution

We first note that there exists a simple solution to the problem of outsourcing computation, namely we can let the client additively secret share his input between the worker and the server. We will refer to the

[1] Also Yao's protocol can be made to fit this framework using standard techniques [HL10].

worker and the server as W_1 and W_2, as the role of the worker and the server is almost symmetrical (the only difference is that, in this application, W_2 also has an input y). Then the client C picks random x_1, x_2 such that $x_1 + x_2 = x$ and sends x_i to W_i. In addition, to make sure neither W_1, W_2 learn the output of the function, we let C send one-time pads r_1, r_2 to W_1, W_2.

Now W_1, W_2 run their favorite secure computation protocol (which guarantees security against actively corrupted parties) to securely evaluate the function

$$g((x_1, r_1), (x_2, r_2, y)) = f(x_1 + x_2, y) + r_1 + r_2$$

and send the output to C who can reconstruct the output by removing the masks. Note that under the assumption that linear operations are "for free" securely evaluating g is as efficient as securely evaluating f.

This solution only works if at most one of W_1, W_2 is passively corrupted, as a malicious adversary can input a share $x_i^* = x_i + \epsilon$ to the secure computation, thus being able to add an error ϵ to the client's input. This can be easily fixed by having the client send a (shared) MAC together with his shares. That is, now the client picks a key k, computes a MAC $t = \mathrm{Tag}_k(x)$ and secret shares the MAC between the server and the worker, who now securely compute the function:

$$g'(\ (x_1, t_1, k_1, r_1), (x_2, t_2, k_2, r_2, y)\ ) \\ := \begin{cases} f(x, y) + r, & \text{if } \mathrm{Ver}_k(t, x) = 1 \\ \texttt{abort}, & \text{else} \end{cases}$$

Unfortunately this requires that the MAC is verified by the secure computation protocol, and this will increase the circuit size significantly. Even using simple, information theoretic MACs of the form $t = a \cdot x + b$ (where $k = (a, b)$) might add a significant number of multiplications (e.g., if one uses garbled circuits as the underlying protocol, then we need to add a number of garbled gates quadratic in the size of the inputs and linear in the number of clients.).

2.2 Our Solution

Our solution relies on the following observation: in the previous protocol MACs and keys are only required to check that the worker and the server do not lie about the client's inputs. After those values are given as input to the secure computation protocol, they are essentially "committed" and cannot be changed anymore. Therefore, at this point the key can simply be revealed, which together with a careful choice of MAC scheme will turn the MAC verification into a linear computation which does not have any significant impact on the efficiency of the overall protocol. Here is a high level description of our protocol:

Client Input Phase: Let K be a sufficiently large finite field, which is *efficient to compute in* securely using the underlying MPC protocol.[2] Assume that all inputs are from K. (If not, simply parse the input as several elements from K and continue as follows for each element.) The client C, on input $x \in K$, picks a random key $k \in K$ and computes a simple algebraic tag[3] $t = \mathrm{Tag}_k(x) = k \cdot x$; The client also picks a uniformly random mask $r \in K$ and sends additive secret shares of the input, the key, the tag and the mask to the workers. The addition is over K.

Workers Computation Phase: The workers W_1, W_2 input the shares they receive to the MPC protocol, then they "open" the key k and check that the tag is correct i.e., they compute (and open) the output of $\mathrm{Ver}_k(x, t)$. If the output is `false`, output `abort`; Else, they compute and open the "encrypted output"

$$c = f(x, y) + r\ ;$$

[2] By K being *efficient to compute in* we mean that taking a multiplication between a secret value and a public value should be very efficient. As an example, if the underlying protocol is Yao garbled circuits, we can take $K = GF(2^k)$. Then multiplication with a public value will just be taking XOR of some of the bits of the secret value, which is essentially for free using the *free XOR* technique [KS08]: no communication and no additional data storage. If the underlying protocol is based on secret sharing or encrypted data over some field, then K can be taken to be that field.

[3] Or an AMD code as introduced in [CDF+08].

Client Output Phase: Finally each worker sends the output to the client. If the client receives the same output from *all* workers, he outputs the unmasked value $z = c - r$, else output `abort`.

The protocol is secure for the client as long as at most one of W_1, W_2 is (actively) corrupted: in a nutshell, a corrupted W_i cannot change its share without breaking the security of the MAC scheme. Of course, a corrupted worker can make the protocol abort and prevent termination, but this is unavoidable in the dishonest majority setting.[4] Note that revealing the key of the MAC has no impact on the security, as by that time a corrupted worker has already committed to his share of the input and the MAC. Moreover, no (selective failure) attacks can be mounted: using a wrong key as input to the protocol always makes the protocol abort. Finally, as the output is masked, the workers do not learn any information about it. Note that a corrupted worker might try to modify the output value by sending $c_i \neq c$ to the client during the output phase – this could be solved by adding a MAC to the output, but in fact a simpler way exist, namely having the client simply check that the two outputs he receives are the same (remember, at least one of W_1, W_2 is honest). Of course a malicious worker can then prevent the client from getting the output, but this is possible anyway as we assume that all but one worker might be corrupted, in which case it is impossible to guarantee termination of the secure multiparty computation protocol run between the workers. Hence a malicious worker might just make the secure computation deadlock, which would have the same effect of the client not getting its output.

In terms of efficiency, the Workers Computation Phase requires only (on top of the complexity of securely computing f) 5 secure additions, 2 secure openings and the cost of computing $\mathrm{Ver}_k(x, t)$. However, as described above, given that the key k is public this can be done by generating a random shared value s and computing $\beta = s \cdot (t - k \cdot x)$. Now $\beta = 0$ iff[5] the MAC is correct and a uniformly random value otherwise. As k is a public value at this stage, this requires only *one* additional multiplication of secret values.

The framework can be used in several settings, choosing appropriate number of clients and workers. As discussed earlier, the single-client/many-workers setting can be used for private and verifiable delegation of computation. This is to be compared with single-server verifiable delegation of computation protocols [PHGR13, BSCG+13], which is getting extremely close to practice if one is only interested in correctness of the result, but requires the use of FHE to achieve input privacy. In addition, current solutions do not extend to the case of multiple clients, while our solutions naturally generalizes. The multiple-client/few-workers setting can be used for large-scale secure computation. As we will show in Section 3.4, using our framework we can even guarantee termination in the relevant setting where a malicious client is trying to make the computation abort by using an invalid input. This is particularly relevant when the numbers of clients is much bigger than the number of workers, and it therefore is undesirable that a single, corrupted client can make the whole computation abort. Think e.g., of an electronic election: a single invalid vote should not prevent all honest parties from reaching a consensus, instead it should be counted as a void vote. This of course introduces new challenges, as a malicious worker should not be able to claim that a client is corrupted and therefore replace the input of a honest party.

The main idea behind our solution is the following: If multiple clients are present, we check that *all* of the MACs are valid using a single multiplication. If $\beta \neq 0$ we recursively split the MACs in half and we search for the incorrect MACs. This takes $\log n$ multiplications times the number of incorrect MACs. Once we identify the set of incorrect MACs, we need to decide whether this is due to a corrupted worker or a corrupted client. Note that this is not trivial: one might think that it is enough to let clients sign their messages, but then a malicious worker could claim he did not receive the message. See Section 3.4 for our solution.

Finally, we note that while it is clear that it is not possible to achieve any privacy if all workers are corrupted, it is possible to achieve (a flavour of) correctness by combining the techniques of this paper with those in [BDO14].

[4] In a later section, we will discuss how to distinguish between a cheating client and a cheating worker. This will be useful in the multi-client setting.

[5] There is a negligible probability of error of $1/|K| \leq 2^{-k}$.

2.3 Eventual Output Consensus

One problem with the above protocol is that it might happen that some honest clients learn their outputs and some other honest clients do not learn their outputs. This can be provoked by a single malicious worker. If the outputs are used as input in a later protocol (for instance, a later execution of the outsourcing protocol itself), this might cause some of the honest clients to start the later protocol while some other honest clients will never start the later protocol. This might lead to problems with functionality, deadlock and also security problems, as it might wash out the fraction of honest clients participating in the ensuring computation.

We will describe a generic and efficient way to achieve a property which we call *eventual output consensus*, meaning that *either* no honest clients receive an output *or* all honest clients will eventually receive their output. This will, e.g., ensure that either no honest clients will continue with a later protocol, or all honest clients will eventually start executing the later protocol.

We now describe our solution. As a first modification, all workers will send $c = (c^1, \ldots, c^n)$ to all clients, the workers will sign the value c, and the clients will accept only if they receive the same c from all workers along with valid signatures from all workers. If so, the client C_j will retrieve c^j from c and output $z^j = c^j - r^j$ as usual. Whenever a client C^j outputs a value c^j, it will forward c to all other clients, along with the signatures of all workers. If a client C^j has not yet given output and receives such a c signed by all workers, it will in turn retrieve c^j from c, output $z^j = c^j - r^j$, and forward c to all clients along with the signatures. It is easy to see that this is secure under the assumptions that at least one worker is honest, as all values used to determine outputs are signed by all workers and therefore also the honest worker.

One might wonder if we can do better than eventual agreement on the output. It turns out that we can not. In our model all but one worker can be corrupted and any number of clients can be corrupted. This means that all-in-all more than half of the participants might be corrupted. It is well-known that in a setting without honest majority, generic MPC cannot guarantee termination or fairness. I.e., we cannot ensure that all honest clients will learn their outputs and we cannot even ensure that it does not happen that the corrupted clients learn their outputs and no honest clients learn their outputs. The best we can hope for is therefore that at least the honest clients have consensus on whether outputs were gotten or not. Since we consider asynchronous communication, where we assume that all messages between honest clients are eventually delivered but has no upper bound on the communication delay, we can only hope for this consensus to eventually arise, which is exactly what our protocol achieves.

3 Our Framework

3.1 Notation and Preliminaries

We write $x \leftarrow K$ to say that x is sampled uniformly from a finite field K. When we write $x + y$, the addition refers to the finite field K (therefore, $k \leftarrow K$, $c = x + k$ is a "one-time-pad" of x). We divide our parties into clients $C_1, \ldots, C_n$ with inputs $x^1, \ldots, x^n \in K$ respectively, and workers $W_1, \ldots, W_m$ with no input. Note that clients and workers need not be disjoint sets, as in the example in the previous section. The clients send a single message (to each worker), then the workers securely evaluate the function $f : K^n \to K^n$ on the clients' inputs, and at the end they send one single message to each client. Therefore the communication pattern for the clients is optimal: one message to provide input, one to receive output.

We require in addition that the workload of the clients should not depend on the size of the function f to be evaluated. We only assume that at least one of the workers is honest. In particular, clients are *not* assumed to be honest. We only consider static corruptions. We prove security in the UC framework [Can01].

3.2 The Underlying MPC Protocol

As discussed before, our framework can be instantiated with any secure computation protocol that allows for reactive computation and where linear operations (additions) are "for free" that is, their efficiency can be essentially ignored when considering the overall complexity of the protocol. In order to keep generality, we will describe our protocol assuming that the workers have access to an ideal functionality for reactive

computation as in Figure 1. Thanks to the UC composition theorem, one can replace the functionality with any protocol that UC-implements it, and the overall protocol will still be secure. This allows for a modular presentation and to separate the issues of the clients interacting with the workers (giving inputs and receiving outputs) without worrying about which specific protocol is used by the workers.

Functionality $\mathcal{F}_{\text{MPC}}$
The functionality is for the parties $W_1, \ldots, W_m$.

Rand: On input $(rand, P_i, vid)$ from all parties P_i, with vid a fresh identifier, the functionality outputs $(rand, vid)$ to the adversary. On a later input $(complete, vid)$ from the adversary, the functionality samples $r \leftarrow K$, stores (vid, r) and outputs $(complete, vid)$ to all parties.

Input: On input $(input, W_i, vid, x \in K)$ from W_i and $(input, W_i, vid, ?)$ from all other parties, with vid a fresh identifier, the functionality outputs $(input, W_i, vid, ?)$ to the adversary. On a later input $(complete, vid)$ from the adversary, the functionality stores (vid, x) and outputs $(complete, vid)$ to all parties.

Secret Eval: On command

$$(eval, vid_1, \ldots, vid_p, f, vid_{p+1}, \ldots, vid_{p+q})$$

from all parties (if $vid_{p+1}, \ldots, vid_{p+q}$ are present in memory and $vid_1, \ldots, vid_p$ are not), the functionality outputs $(eval, vid_1, \ldots, vid_p, f, vid_{p+1}, \ldots, vid_{p+q})$ to the adversary. On a later input

$$(complete, vid_1, \ldots, vid_p, f, vid_{p+1}, \ldots, vid_{p+q})$$

from the adversary, the functionality retrieves $(vid_{p+1}, x_1), \ldots, (vid_{p+q}, x_q)$, computes $(z_1, \ldots, z_p) \leftarrow f(x_1, \ldots, x_q)$ and stores $(vid_1, z_1), \ldots, (vid_p, x_p)$. Then it outputs $(complete, vid_1, \ldots, vid_p, f, vid_{p+1}, \ldots, vid_{p+q})$ to all parties.

Output: On input $(output, W_i, vid)$ from all parties (if vid is present in memory), the functionality outputs $(output, W_i, vid)$ to the adversary. On a later input $(complete, vid)$ from the adversary, the functionality retrieves (vid, x) and outputs it to W_i.

Further adversarial behavior: All the completion messages are delivered at the will of the adversary only, i.e., a message like $(complete, \ldots)$ is not delivered to an honest party until the adversary signals that it should be delivered, and the adversary might deliver to different parties at different times.

Figure 1. The ideal functionality for reactive MPC.

We use some short-hand notation: $[x]$ is a secret representation of $[x]$, i.e., a value uploaded to the ideal functionality using the Input command or computed via the Secret Eval command. The representation is assumed to be *cheap to compute on* using linear operations on elements from K, so we will write $[ax + by] = a[x] + b[y]$ for publicly known $a, b \in K$ and secrets $x, y \in K$, and we will not count these operations towards the complexity of the protocol. We assume that K has size at least 2^k, where k is the security parameter. We use this notation:

Input: $[x] \leftarrow \mathsf{Input}(P_i, x)$ allows party P_i to input the value $x \in K$ to the computation; We also define a command $[r] \leftarrow \mathsf{Rand}()$ which can be simply implemented by $[r_i] \leftarrow \mathsf{Input}(P_i, r_i)$; for all i and random $r_i \in K$ and $[r] = \sum_i [r_i]$;

Eval: $([z_1], \ldots, [z_p]) \leftarrow f([x_1], \ldots, [x_q])$ allows to compute an agreed upon function f of q inputs and p outputs on secret representations, producing again secret representations. This is done via the Secret Eval command.

Linear: For public $a, b \in K$ and secret $x, y \in K$ the command $[z] \leftarrow a[x] + b[y]$ allows parties to compute a linear combination in K. This is a special case of Secret Eval, but we single it out for notational convenience and because the command is assumed to be essentially for free for our framework to make sense.

Multiplication: $[z] \leftarrow \mathsf{Mul}([x],[y])$ allows parties to compute a representation of $z = x \cdot y$. This is a special case of Secret Eval, but we single it out for notational convenience and because the command is assumed to be not too expensive for our framework to make sense.

Open: $x \leftarrow \mathsf{Open}([x])$ publicly reveals the value inside $[x]$; We also define $x \leftarrow \mathsf{OpenTo}(P_i,[x])$ which allows to reveal a value only to party P_i, and can always be implemented doing $[r] \leftarrow \mathsf{Input}(P_i, r)$ for uniform random r chosen by P_i and $c \leftarrow \mathsf{Open}([k+r])$, and then party P_i outputs $x = c - r$;

All algebraic notation denotes operations in K.

Clients Input Phase: Each client C^j with input x^j:

1. Pick random $\{x_i^j\}_{i=1}^m$ from K s.t., $\sum_{i=1}^m x_i^j = x^j$;
2. Pick random $\{k_i^j\}_{i=1}^m$ from K; let $k^j = \sum_{i=1}^m k_i^j$;
3. Pick random $\{r_i^j\}_{i=1}^m$ from K; let $r^j = \sum_{i=1}^m r_i^j$;
4. Compute $t^j = \mathrm{Tag}_{k^j}(x^j) = k^j \cdot x^j$;
5. Pick random $\{t_i^j\}_{i=1}^m$ from K s.t., $\sum_{i=1}^m t_i^j = t^j$;
6. Send the values $v_i^j = (x_i^j, t_i^j, k_i^j, r_i^j)$ to W_i for $i \in 1,\ldots,m$.

Workers Computation Phase: The workers $W_1,\ldots,W_m$ do:

1. Each worker W_i, for $i \in 1,\ldots,m$, waits until receiving input (*eval*) and then proceeds as below.
2. Each worker W_i, for $i \in 1,\ldots,m$: $[x_i^j] \leftarrow \mathsf{Input}(W_i, x_i^j)$; $[t_i^j] \leftarrow \mathsf{Input}(W_i, t_i^j)$; $[k_i^j] \leftarrow \mathsf{Input}(W_i, k_i^j)$; $[r_i^j] \leftarrow \mathsf{Input}(W_i, r_i^j)$;
3. Compute: $[x^j] = \sum_{i=1}^m [x_i^j]$ for all $j \in 1,\ldots,n$;
4. Compute: $[t^j] = \sum_{i=1}^m [t_i^j]$ for all $j \in 1,\ldots,n$;
5. Compute: $[k^j] = \sum_{i=1}^m [k_i^j]$ for all $j \in 1,\ldots,n$;
6. Compute: $[r^j] = \sum_{i=1}^m [r_i^j]$ for all $j \in 1,\ldots,n$;
7. $k^j \leftarrow \mathsf{Open}([k^j])$ for all $j \in 1,\ldots,n$;
8. $[\alpha^j] = [t^j] - k^j \cdot [x^j]$ for all $j \in 1,\ldots,n$;
9. $[s] \leftarrow \mathsf{Rand}()$;
10. $[\beta] = [s] \cdot [\sum_{j \in 1,\ldots,n} \alpha^j]$;
11. $\beta \leftarrow \mathsf{Open}([\beta])$; If $\beta \neq 0$ output `abort`, else continue;
12. Compute: $([z^1],\ldots,[z^n]) = f([x^1],\ldots,[x^n])$;
13. Compute: $[c^j] = [z^j] + [r^j]$ for all $j \in 1,\ldots,n$;
14. $c^j \leftarrow \mathsf{Open}([c^j])$ for all $j \in 1,\ldots,n$;

Client Output Phase:

1. (Each worker W_j) Send c^j to C_j;
2. Let c_i^j be the output that C_j receives from W_i;
3. If the vector $(c_1^j,\ldots,c_n^j)$ is not constant, C_j outputs `abort`, else let c^j be the constant value;
4. C^j outputs $z^j = c^j - r^j$.

Figure 2. The protocol.

3.3 Protocol Analysis

The protocol is given in Figure 2 (see Section 2 for a high-level description of the protocol).

Theorem 1. *Let π be the protocol in Figure 2. We prove that π securely implements the ideal functionality $\mathcal{F}_{\mathrm{FE}}^f$ against any static adversary corrupting any number of clients and at most $m-1$ workers.*

Functionality $\mathcal{F}_{\text{FE}}^{f}$
The functionality is for the parties $C_1, \ldots, C_n, W_1, \ldots, W_m$.

Input: On input $(input, C_i, x^i)$ from C_i and $(input, C_i, ?)$ from all W_j, the functionality outputs $(input, C_i, ?)$ to the adversary. On a later input $(complete, C_i)$ from the adversary, the functionality stores $(input, C_i, x^i)$ and outputs $(complete, input, C_i)$ to all W_j.

Eval: On command $(eval)$ from all W_j at a point where $(input, C_i, x^i)$ is stored for all $i \in 1, \ldots, n$, compute $(z^1, \ldots, z^n) = f(x^1, \ldots, x^n)$ and store $(output, C_i, z^i)$ for all $i \in 1, \ldots, n$. Then output $(eval, \{z^j\}_{j \in J})$ to the adversary, where J is the set of indices of corrupted parties.

Output: On command $(deliver, C_i)$ from the adversary, at a point where $(output, C_i, z^i)$ is stored, output z^i to C_i and delete $(output, C_i, z^i)$.

Figure 3. The ideal functionality for function evaluation of f.

We do the proof in the UC framework [Can01]. We prove static security against an adversary corrupting any number of clients and up to all but one of the workers. Recall that we have the following proof burden.

There is a real world where we run the protocol. The parties of the protocol has access to secure point-to-point channels (which can in turn be implemented using cryptography) plus a copy of $\mathcal{F}_{\text{MPC}}$. In the real world there is an adversary A attacking the protocol. It is the adversary A which controls the corrupted parties, i.e., it sends messages on behalf of the corrupted parties and sees all messages sent to the corrupted parties, including the messages to and from $\mathcal{F}_{\text{MPC}}$. In addition the adversary has access to the adversarial behavior allowed by $\mathcal{F}_{\text{MPC}}$, like deciding when messages are delivered. There is also an environment Z. It is the environment which provides the inputs to the honest parties of the protocol and which sees the outputs of the honest parties of the protocol. The environment Z can also interact with A, in any way that they desire and at any time. The interaction is via exchanging messages. At the end of the interaction, the environment outputs a bit. We denote the distribution of this bit by $\text{EXEC}_{\pi,A,Z}(k)$, where k is the security parameter. Both A and Z are restricted to poly-time computations.

In the ideal process there are three entities, the ideal functionality $\mathcal{F}_{\text{FE}}^{f}$, the adversary S and the environment Z. It is the environment which provides inputs to $\mathcal{F}_{\text{FE}}^{f}$ via the (dummy) honest parties, and it sees their outputs from $\mathcal{F}_{\text{FE}}^{f}$. It is the adversary S which provides inputs to $\mathcal{F}_{\text{FE}}^{f}$ on behalf of the corrupted parties, and it sees the outputs to the corrupted parties from $\mathcal{F}_{\text{FE}}^{f}$. In addition it has access to the adversarial behavior allowed by $\mathcal{F}_{\text{FE}}^{f}$. Besides this, Z and S can interact by exchanging messages. At the end of the execution, Z will output a bit. We use $\text{EXEC}_{\mathcal{F}_{\text{FE}}^{f},S,Z}(k)$ to denote the distribution of this bit. Both S and Z are restricted to poly-time computations.

To prove security of the protocol we have to construct for all adversaries A for the real world an adversary S for the ideal process such that no Z can guess whether it interacts with π and A or $\mathcal{F}_{\text{FE}}^{f}$ and S. We also call this adversary S a simulator. Technically we require that for all A there exists S such that for all Z the value $|\Pr[\text{EXEC}_{\pi,A,Z}(k) = 1] - \Pr[\text{EXEC}_{\mathcal{F}_{\text{FE}}^{f},S,Z}(k) = 1]|$ goes to 0 faster than any inverse polynomial in k.

We proceed to the proof. Assume we are given any A. The simulator S works as follows:

Simulated Protocol S runs internally a copy of π, i.e., a copy of each party $C_1, \ldots, C_n, W_1, \ldots, W_m$ along with a copy of $\mathcal{F}_{\text{MPC}}$. We call this the *simulated protocol*. To distinguish the simulated parties from the corresponding parties in the real execution we write $\tilde{C}_1, \ldots, \tilde{C}_n, \tilde{W}_1, \ldots, \tilde{W}_m$ and $\tilde{\mathcal{F}}_{\text{MPC}}$ for the parties and the ideal functionality and $\tilde{\pi}$ for the simulated protocol as a whole.

Simulated Adversary S also runs internally a copy of A, we call this the *simulated adversary* and denote it by $\tilde{A}$.

Monitor Corrupted Parties S lets the simulated adversary $\tilde{A}$ and the simulated parties interact exactly as in the real execution, i.e., whenever $\tilde{A}$ instructs a corrupted party to send a given message, the simulator performs this command in the simulated protocol, and whenever a corrupted simulated party receives a message in the execution of the simulated protocol, S gives this message to A. Notice that as

a consequence of this simulation strategy, S knows all messages sent and received by corrupted parties, including the messages to and from $\tilde{\mathcal{F}}_{\text{MPC}}$. The simulator also knows the internal state of $\tilde{\mathcal{F}}_{\text{MPC}}$, as it is S which runs the copy $\tilde{\mathcal{F}}_{\text{MPC}}$. We use these facts later.

Relay Between Adversary and Environment S lets the simulated adversary A and Z interact exactly as in the real execution, i.e., whenever $\tilde{A}$ sends a message to its environment S passes it on to Z, and whenever Z sends a message to S, the simulator just passes it on to A.

Dummy Honest Inputs Whenever Z gives an input x^i to an honest C_i, the simulator S is given $(input, C_i, ?)$. It then picks a dummy input $\tilde{x}^i$ for $\tilde{C}_i$, e.g., $\tilde{x}^i = 0$ or some other legal input. Then it simply runs $\tilde{\pi}$ according to the protocol, but with this dummy input $\tilde{x}^i$ to $\tilde{C}_i$ instead of the correct input x^i (which is unknown to S by the rules of the game).

Eval Whenever Z gives an input $(eval)$ to an honest W_i, the simulator S is given $(eval, W_i)$. It then simply inputs $(eval)$ to $\tilde{W}_i$ in the simulated protocol and then runs $\tilde{W}_i$ according to the protocol.

Extracting Corrupted Inputs Recall that it is S which must provide input to $\mathcal{F}_{\text{FE}}^f$ on behalf of the corrupted parties. It must try to make these inputs consist with the "input" of the corrupted $\tilde{C}_j$. Note that $\tilde{C}_j$ has no explicit input, it is defined via its behavior in the protocol. The job of the simulator is hence to extract this implicit input. The simulator extracts the input x^j of a corrupted $\tilde{C}_j$ as follows: It waits until all workers has input the values $\tilde{x}_i^j, \tilde{t}_i^j, \tilde{k}_i^j, \tilde{r}_i^j$ to $\tilde{\mathcal{F}}_{\text{MPC}}$. Then it computes $\tilde{x}^j = \sum_{i=1}^m \tilde{x}_i^j$ and $\tilde{r}^j = \sum_{i=1}^m \tilde{r}_i^j$. Then it inputs $x^j = \tilde{x}^j$ to $\mathcal{F}_{\text{FE}}^f$ on behalf of the corrupted C^j. It saves $\tilde{r}^j$ for later use.

Patching Corrupted Outputs Recall that A sees the outputs of corrupted parties in the simulated protocol and can talk to Z which sees the inputs and outputs of of $\mathcal{F}_{\text{FE}}^f$. Hence it is important that the outputs in the simulated protocol are consistent with the inputs and outputs of $\mathcal{F}_{\text{FE}}^f$. This is not necessarily the case, as we ran the simulated protocol on dummy inputs for all honest parties. The simulator deals with this as follows. Assume that the simulated protocol reaches Step 14 in the Workers Computation Phase. Before executing this step, the simulator S will modify the value of c^j inside $\tilde{\mathcal{F}}_{\text{MPC}}$ for each corrupted $\tilde{C}_j$ before executing Step 14. Observe that if the simulated protocol reaches Step 14, then all parties in the simulated protocol must have given inputs, or the simulated protocol would not have passed Step 2 in the Workers Computation Phase. Hence S will have given an input x^j to $\mathcal{F}_{\text{FE}}^f$ on behalf of each corrupted C_j at this point. Furthermore, since it is S which gives (dummy) inputs to the honest $\tilde{C}_j$ in the simulated protocol and since S only does so when Z gives input to the corresponding C_j on $\mathcal{F}_{\text{FE}}^f$, we can conclude that when $\tilde{\pi}$ reaches Step 14, the environment gave input to $\mathcal{F}_{\text{FE}}^f$ on behalf of all honest C_j. So, all in all, when the simulated protocol reaches Step 14, all C_j received an input x^j in $\mathcal{F}_{\text{FE}}^f$. Using a similar line of reasoning we see that if $\tilde{\pi}$ reaches Step 14, then Z must have input $(eval)$ to all honest W_i on $\mathcal{F}_{\text{FE}}^f$. So, now S can input $(eval)$ to all corrupted W_i on $\mathcal{F}_{\text{FE}}^f$. In response $\mathcal{F}_{\text{FE}}^f$ computes $(z^1, \ldots, z^n) = f(x^1, \ldots, x^n)$ and outputs $(eval, \{z^j\}_{j \in J})$ to S, where J is the set of corrupted parties. Now S computes $\tilde{c}^j = z^j + \tilde{r}^j$, where $\tilde{r}^j$ was computed and stored in Extracting Corrupted Inputs. Then S changes the internal state of $\tilde{\mathcal{F}}_{\text{MPC}}$ to hold the value $c^j = \tilde{c}^j$. Then it simply runs the simulated protocol according to the protocol.

Honest Output Delivery Whenever an honest client $\tilde{C}_i$ reaches Step 4 in Client Output Phase, the simulator inputs $(deliver, C_i)$ to $\mathcal{F}_{\text{MPC}}$, which makes $\mathcal{F}_{\text{MPC}}$ output z^i to Z on behalf of C_i.

That completes the description of the simulator. We now prove that

$$|\Pr[\text{EXEC}_{\pi,A,Z}(k) = 1] - \Pr[\text{EXEC}_{\mathcal{F}_{\text{FE}}^f,S,Z}(k) = 1]| \leq 2^{-k+1} \ .$$

Let E be the event that some client C_j has its input replaced, i.e., C_j runs with input x^j, but after Step 2 in Workers Computation phase it holds for the values x_i^j in $\mathcal{F}_{\text{MPC}}$ that $x^j \neq \sum_{i=1}^m x_i^j$. We can also define E in the simulation, but via the (dummy) input $\tilde{x}^j$ and the values $\tilde{x}_i^j$ in $\tilde{\mathcal{F}}_{\text{MPC}}$ and say that E occurs when $\tilde{x}^j \neq \sum_{i=1}^m \tilde{x}_i^j$. Let $\bar{E}$ denote the event that E did not occur. We clearly have that

$$\begin{aligned}\Pr[\text{EXEC}_{\pi,A,Z}(k) = 1] = \\ &\Pr[E] \Pr[\text{EXEC}_{\pi,A,Z}(k) = 1|E] + \\ &(1 - \Pr[E]) \Pr[\text{EXEC}_{\pi,A,Z}(k) = 1|\bar{E}]\end{aligned}$$

and

$$\begin{aligned}\Pr[\textsc{Exec}_{\mathcal{F}_{\text{FE}}^f,S,Z}(k)=1] = \\ \Pr[E]\Pr[\textsc{Exec}_{\mathcal{F}_{\text{FE}}^f,S,Z}(k)=1|E]+ \\ (1-\Pr[E])\Pr[\textsc{Exec}_{\mathcal{F}_{\text{FE}}^f,S,Z}(k)=1|\bar{E}]\ .\end{aligned}$$

We will first show *Claim 1*: the probability $\Pr[E]$ is the same in the real execution and in the ideal process. We will also prove *Claim 2*: $\Pr[E] \leq 2^{-k+1}$. Then we show *Claim 3*: It holds that

$$\Pr[\textsc{Exec}_{\mathcal{F}_{\text{FE}}^f,S,Z}(k)=1|\bar{E}] = \Pr[\textsc{Exec}_{\pi,A,Z}(k)=1|\bar{E}]\ .$$

From these three claims it follows that

$$\begin{aligned}&|\Pr[\textsc{Exec}_{\pi,A,Z}(k)=1] - \Pr[\textsc{Exec}_{\mathcal{F}_{\text{FE}}^f,S,Z}(k)=1]| \\ &\quad\leq |\Pr[E]\Pr[\textsc{Exec}_{\pi,A,Z}(k)=1|E]- \\ &\qquad \Pr[E]\Pr[\textsc{Exec}_{\mathcal{F}_{\text{FE}}^f,S,Z}(k)=1|E]| \\ &\quad\leq \Pr[E] \leq 2^{-k+1}\end{aligned}$$

as desired.

Proof of Claims 1 and 2 We prove that if the adversary uses values x_i^j such that $x^j \neq \sum_{i=1}^m x_i^j$, then $\beta \neq 0$ with probability $\leq 2^{-k+1}$, and the probability is independent of x^j. The same proof applies to the simulation, as the simulation is just a run of the real protocol but on different inputs. From this Claims 1 and 2 then follows.

Let $\alpha = \sum_j \alpha^j$. There are two (non-disjoint) ways it can happen that $\beta = 0$, namely $\alpha = 0$ and $s = 0$. Since $\Pr[s = 0] = 2^{-k}$, independently of x^j, it is sufficient to prove that $\Pr[\alpha = 0] \leq 2^{-k}$ and that the probability is independent of x^j.

If all clients are corrupted, there is nothing to prove. So, since corrupting more parties gives the adversary strictly more powers and since the role of all clients is symmetric, let us assume without loss of generality that all clients except C_1 are corrupted. Our model assumes that at least one worker is honest. Since the role of all workers are symmetric, let us assume without loss of generality that all workers except W_1 are corrupted. Finally, if $x^1 = \sum_{i=1}^m x_i^1$ there is nothing to prove, so assume that this is not the case. Finally, since Z and A can communicate, we cannot assume that A does not know x^1, and since getting x^1 clearly cannot make A worse off, let us assume without loss of generality that A knows x^1. To distinguish between the correct values of x_i^1 and the wrong ones chosen by corrupted workers, use x_i^1 to denote the values chosen by C_1 and use $\bar{x}_i^1$ to denote the values in $\mathcal{F}_{\text{MPC}}$. Note that $\bar{x}_1^1 = x_1^1$. We use similar notation for the values k_i^j and t_i^j.

Notice that for $i > 1$ the adversary knows both x_i^1 and $\bar{x}_i^1$, as it received x_i^1 and it chose $\bar{x}_i^1$. Hence it also knows the relative error $X_i^1 = \bar{x}_i^1 - x_i^1$. Notice that $\bar{x}_i^1 = x_i^1 + X_i^1$. I.e., the adversary inputs the correct input plus some known error. Similarly we can write $\bar{k}_i^1 = k_i^1 + K_i^1$ and $\bar{t}_i^1 = t_i^1 + T_i^1$ for values K_i^1 and T_i^1 known by the adversary. We use $X^1 = \sum_{i=2}^m X_i^1$ to denote the sum of relative errors. Note that X^1 is known by the adversary. Similarly, let $T^1 = \sum_{i=2}^m T_i^1$ and let $K^1 = \sum_{i=2}^m K_i^1$. Note that $\tilde{x}^1 = x^1 + X^1$, $\tilde{t}^1 = t^1 + T^1$, and $\tilde{k}^1 = k^1 + K^1$. We have that $t^1 - k^1x^1 = 0$, by design. Hence $\tilde{t}^1 - \tilde{k}^1\tilde{x}^1 = t^1 + T^1 - (k^1 + K^1)(x^1 + X^1) = t^1 - k^1x^1 - k^1X^1 - K^1x^1 + T^1 = -k^1X^1 - (K^1x^1 + T^1)$. Since we assume that A knows x^1, the value $K^1x^1 + T^1$ is known by the adversary. Finally, observe that $\tilde{\alpha} = \tilde{\alpha}^1 + \sum_{i>1} \tilde{\alpha}^i$, where $\tilde{\alpha}^i$ is known to A for $i > 1$. Hence, $\tilde{\alpha} = -k^1X^1 - (K^1x^1 + T^1) + \sum_{i>1} \tilde{\alpha}^i$, where $-(K^1x^1 + T^1) + \sum_{i>1} \tilde{\alpha}^i$ is known to A. Notice then that $\tilde{\alpha} = 0$ iff $k^1X^1 = -(K^1x^1 + T^1) + \sum_{i>1} \tilde{\alpha}^i$. Since k^1 is uniformly random and independent of the view of A and $-(K^1x^1 + T^1) + \sum_{i>1} \tilde{\alpha}^i$ is known to A and $X^1 \neq 0$ by assumption, it follows that $k^1X^1 = -(K^1x^1 + T^1) + \sum_{i>1} \tilde{\alpha}^i$ with probability exactly 2^{-k}, independently of the value of X^1.

Proof of Claim 3 First consider the following mind game. Consider an execution $\text{EXEC}_{\mathcal{F}_{\text{FE}}^f, S^1, Z}(k)$, which runs exactly as the execution $\text{EXEC}_{\mathcal{F}_{\text{FE}}^f, S, Z}(k)$, except that each time where S is about to use a dummy input $\tilde{x}^i = 0$ on behalf of honest C_i, it instead cheats and inspects $\mathcal{F}_{\text{FE}}^f$ to get the real value x_i and then it uses $\tilde{x}^i = x^i$. Besides this cheat, everything runs as in the simulation. Note that in a simulation S is of course not allowed to perform the above cheat. However, we are here only defining a random variable $\text{EXEC}_{\mathcal{F}_{\text{FE}}^f, S^1, Z}(k)$ for sake of the proof, and we are of course free to define it as we want. We claim that $|\Pr[\text{EXEC}_{\mathcal{F}_{\text{FE}}^f, S, Z}(k) = 1|\bar{E}] - \Pr[\text{EXEC}_{\mathcal{F}_{\text{FE}}^f, S^1, Z}(k) = 1|\bar{E}]| = 0$. The reason is that the views of A and Z do not depend on the dummy inputs at all. To see this notice that $\tilde{x}^i$ is input to $\tilde{\mathcal{F}}_{\text{MPC}}$ and is not used anywhere else. And, the only values leaked by $\tilde{\mathcal{F}}_{\text{MPC}}$ which might depend on $\tilde{x}^i$ are the values c^j. As for the value c^j for all corrupted C_j, note that it is patched to $\tilde{c}^j$ before it is output from $\tilde{\mathcal{F}}_{\text{MPC}}$, hence it has the same distribution no matter whether S uses $\tilde{x}^j = 0$ or $\tilde{x}^j = x^j$. As for the value c^j for all honest C_j, note that $c^j = \tilde{z}^j + r^j$ for a uniformly random value r^j. Hence, even though $\tilde{z}^j$ might depend on whether S uses $\tilde{x}^j = 0$ or $\tilde{x}^j = x^j$, the value c^j does not, as it is one-time pad encrypted with r^j which is known only to the honest C_j.

Consider then the mind game $\text{EXEC}_{\mathcal{F}_{\text{FE}}^f, S^2, Z}(k)$, which runs exactly as the previous mind game $\text{EXEC}_{\mathcal{F}_{\text{FE}}^f, S^1, Z}(k)$, except that we change the step Patching Corrupted Outputs such that S does not perform the patching $c^j = \tilde{c}^j$. Instead it just runs the simulated protocol with the value c^j already inside $\tilde{\mathcal{F}}_{\text{MPC}}$. We claim that $|\Pr[\text{EXEC}_{\mathcal{F}_{\text{FE}}^f, S^1, Z}(k) = 1|\bar{E}] - \Pr[\text{EXEC}_{\mathcal{F}_{\text{FE}}^f, S^2, Z}(k) = 1|\bar{E}] = 0$. Notice that in $\text{EXEC}_{\mathcal{F}_{\text{FE}}^f, S^2, Z}(k)$ each honest $\tilde{C}_i$ is run with input $\tilde{x}^i = x^i$, where x^i is the input to C_i on $\mathcal{F}_{\text{FE}}^f$. Then for each corrupted $\tilde{C}_i$ define $\tilde{x}^i$ as in Extracting Corrupted Inputs. Then it follows from the fact that E does not happen that the values $(z^1, \ldots, z^n)$ computed by $\tilde{\mathcal{F}}_{\text{MPC}}$ are equal to $f(\tilde{x}^1, \ldots, \tilde{x}^n)$, as no honest party has its input replaced. Since the input to $\mathcal{F}_{\text{FE}}^f$ is $x^i = \tilde{x}^i$ for the honest party (qua the cheat) and $x^i = \tilde{x}^i$ for the corrupted party (by design of Extracting Corrupted Inputs), it follows that the values $(z^1, \ldots, z^n)$ computed by $\tilde{\mathcal{F}}_{\text{MPC}}$ are equal to the values $(z^1, \ldots, z^n)$ computed by $\mathcal{F}_{\text{FE}}^f$. Furthermore, when the values $(z^1, \ldots, z^n)$ computed by $\tilde{\mathcal{F}}_{\text{MPC}}$ are equal to the values $(z^1, \ldots, z^n)$ computed by $\mathcal{F}_{\text{FE}}^f$, the patching $c^j = \tilde{c}^j$ clearly has no effect as $c^j = z^j + \tilde{r}^j$ and $\tilde{c}^j = z^j + \tilde{r}^j$. Ergo, it does not matter whether we do the patching or not.

We finally claim that the following holds

$$|\Pr[\text{EXEC}_{\mathcal{F}_{\text{FE}}^f, S^2, Z}(k) = 1|\bar{E}] - \Pr[\text{EXEC}_{\pi, A, Z}(k) = 1|\bar{E}]| = 0 \ .$$

Notice that after replacing dummy inputs by true inputs and dropping the patching of corrupted outputs, the "simulated" protocol run by S^2 is actually just a correct run of π on exactly the same inputs x^i given to $\mathcal{F}_{\text{FE}}^f$ by Z. I.e., S is internally running π *exactly* as it is being run in $\text{EXEC}_{\pi, A, Z}(k)$. There might, however, still be a difference in the view of Z: In $\text{EXEC}_{\mathcal{F}_{\text{FE}}^f, S^2, Z}(k)$ the output of an honest client C_i to Z is the value z^i output by $\mathcal{F}_{\text{FE}}^f$. In $\text{EXEC}_{\pi, A, Z}(k)$ the output of an honest client C_i to Z is the value z^i output by C_i. We hence need to argue that in $\text{EXEC}_{\mathcal{F}_{\text{FE}}^f, S^2, Z}(k)$ the output of an honest client $\tilde{C}_i$ is identical to the value z^i output by $\mathcal{F}_{\text{FE}}^f$. To see this, recall that we argued above that the values $(z^1, \ldots, z^n)$ computed by $\tilde{\mathcal{F}}_{\text{MPC}}$ are equal to the values $(z^1, \ldots, z^n)$ computed by $\mathcal{F}_{\text{FE}}^f$. Hence, all we have to argue is that the output of an honest client $\tilde{C}_i$ is identical to the value z^i computed by $\mathcal{F}_{\text{MPC}}$. This is so, as $c^i = z^i + r^i$ and the correct value of c^i is sent to C_i by at least one honest worker, so C_i only accepts c^i, and then it outputs $c^i - r^i = z^i$.

3.4 Coping with Malicious Clients

In the previous protocol a malicious client can force the entire protocol to abort by using an invalid input (e.g., providing a MAC that is not consistent). As a result it will happen that $\beta \neq 0$ and the computation will abort. In an application where there are thousands or millions of clients it is undesirable that a single malicious (or faulty) client can prevent the whole computation from terminating. We therefore want a protocol which guarantees termination whenever all workers are honest. The problem is that when $\beta \neq 0$

we cannot see if it happened because of a malicious worker or a malicious client. We therefore need to add extra mechanisms for detecting who was cheating and for recovering when it was a client.

The high-level idea of our solution is to identify clients which provided invalid inputs and replace their inputs with a default value. This means that all attacks possible by a client are equivalent to choosing an alternative input, which is of course an allowed option for even an honest party. The only difference is that in case of cheating, all workers will learn that the client cheated, and all workers will learn the (alternative) input of the client. This should deter clients from cheating at all. "Cheating" might, however, also happen because a client is faulty, so we should expect a few cheating parties occasionally. Since we want to consider cases where n is huge, we will therefore present a solution optimized for the case where there are very few cheating parties relative to the size of n.

The first step is to identify the troublesome clients. Note that $\sum_{j\in 1,\ldots,n} \alpha^j \neq 0$ just shows that some $\alpha^j \neq 0$. We can find it by securely computing $\sum_{j=1}^{n/2} \alpha^j$ and $\sum_{j=n/2+1}^{n} \alpha^j$ and revealing blinded values to check which of them is non-zero, possibly both. Then continue like this recursively, until all j for which $\alpha^j \neq 0$ are known. At the point where $O(n)$ sums have been computed, blinded and opened, switch to a mode where we simply open blinded versions of all α^j. This way we never use more than $O(n)$ multiplications, and if there is a constant number of non-zero α^j, then we use only $O(\log n)$ multiplications. Let I be the set of indices j such that $\alpha^j \neq 0$. The above strategy finds I except with negligible probability. Now we must for each $j \in I$ find out whether C_j is corrupted or whether some worker input a different value to the computation than the one provided by C_j. Assume for now that 1) a worker can reveal which values it received from C_j without the worker being able to reveal a value different from the ones actually received from C_j and 2) it is OK to reveal the input x^j of C_j. In that case the solution is trivial: for each $j \in I$ each worker W_i will reveal the values (x_i^j, t_i^j, k_i^j) received from C_i and the workers will open the values $[x^j]$, $[t^j]$ and $[k^j]$ to check that W_i uploaded the right values to $\mathcal{F}_{\text{MPC}}$. If not, W_i is corrupted, and the protocol is terminated. Otherwise, it must be the case that $t^j \neq \text{Tag}_{k^j}(x^j)$, as $\alpha^j \neq 0$. So, since all the uploaded values were the ones received from the client, it follows that C_j is corrupted. In that case the protocol continues, using, e.g., $x^j = 0$ as input on behalf of C_j.We now discuss how to get rid of the two assumptions.

As for the second assumption, we will simply let each C_j split the input x^j into two random shares $y^{j,1}$ and $y^{j,2}$ for which $x^j = y^{j,1} + y^{j,2}$. E.g., pick $y^{j,1}$ uniformly at random and let $y^{j,2} = x^j - y^{j,1}$. Then proceed as before, but let C_j give both of the inputs $y^{j,1}$ and $y^{j,1}$ as above, i.e., with separate keys and MACs. Now, if any one of them turns out to be troublesome, open it and find out who was cheating. If it was C_j, use a default value. If it was W_i, terminate. Then compute $(z^1, \ldots, z^n) = f(y^{1,1} + y^{1,2}, \ldots, y^{n,1} + y^{n,2})$. Note that there is never a need to reveal both $y^{j,1}$ and $y^{j,2}$: if they are both troublesome, simple reveal $y^{j,1}$ and use that value to make the decision. Furthermore, revealing a single of the values $y^{j,1}$ or $y^{j,2}$ is secure, as they are individually uniformly random and independent of x^j.

As for the first assumption, notice that it is not enough to ask C_j to sign the values it sends to W_i: The client might simply not send a signature, and when W_i complains that C_j did not send a signature, it might be W_i that is lying. In fact, any solution where the client sends something over a secure channel will fall prey to this problem: The client might refuse to send the value, but it might also by the worker lying about not having received that value. We therefore need a solution where clients only sends public values. Furthermore, since all but one worker might be corrupted, any public value not sent to all workers, might still fall prey to the above attack: the $m-1$ workers seeing the value might refuse to have received it. Hence we might essentially restrict ourselves to solutions where the client sends one public values and sends it to all workers.

We describe one such solution. Assume that each W_i has a public encryption key e_i for a public-key encryption scheme and that only W_i knows the decryption key d_i. We need that this encryption scheme is secure against chosen-ciphertext attack (IND-CCA2) and that decryption yields the message plus the randomness used to encrypt, and that IND-CCA2 security holds even if the decryption oracle returns this randomness. In the random oracle model, RSA-OAEP is such an encryption scheme, assuming that the RSA function is one-way. Our solution then proceeds as follows: Use v_i^j to denote the value that a client C_j should send secretly to W_i. The client will compute an encryption $\gamma_i^j = E_{e_i}((i, j, v_i^j); s_i^j)$, where s_i^j is the randomness used by the encryption algorithm. Then C broadcasts $(\gamma_1^j, \ldots, \gamma_n^j)$ to all workers. Then W_i

computes $(i', j', v_i^j, s_i^j) = D_{d_i}(\gamma_i^j)$. If $i' \neq i$ or $j' \neq j$, then W_i broadcasts (i', j', v_i^j, s_i^j) and all workers check that $\gamma_i^j = E_{e_i}((i', j', v_i^j); s_i^j)$ and that $i' \neq i$ or $j' \neq j$. If this is the case, use a default input for C_j. If it is not the case, then W_i is cheating. In that case, terminate the protocol. If W_i is later asked to reveal v_i^j, W_i broadcasts (v_i^j, s_i^j) and all workers check that $\gamma_i^j = E_{e_i}((i, j, v_i^j); s_i^j)$.

It is also possible to get a solution not using the random oracle model. Each key e_i is the parameters for an identity-based encryption scheme and d_i is the master secret key. The client will compute an encryption $\gamma_i^j \leftarrow E_{e_i,(i,j,sid)}(v_i^j)$, i.e., encrypt v_i^j under the identity (i, j, sid), where sid is a session identifier which is fresh for each run of the protocol. If W_i is later asked to reveal v_i^j, W_i generates and broadcasts the secret key $d_{i,j,sid}$ for identity (i, j, sid). Then all workers compute $v_i^j = D_{d_{i,j,sid}}(\gamma_i^j)$.

We are then left with the problem of how the client broadcasts to the servers. Note that we can use a standard authenticated broadcast protocol like Dolev-Strong broadcast [DS83], as this protocol require the sender to send just a single message to each of the other participant, here the workers.

					Client		
Protocol	n	m	Security	Based on	Work	Interaction	Notes
This	any	> 1	Active	any	$m \cdot size_{inp}$ (fo)	N	
PTK [PTK13]	any	2	Passive	Paillier	$size_{inp}$ (pk)	N	
Whitewash [CLT14]	1	2	Active	GC	$size_{inp}$ (sk)	N	
CMTB [CMTB13]	1	2	Active	GC	$size_{inp}$ (pk)	Y	Non-collusion
Salus [KMR12]	any	2	Active	GC	$size_{inp}$ (sk)	Y	Non-collusion

Table 1. Comparison with previous work. n, m are the number of allowed clients and workers respectively. In the Client Work column, (fo/sk/pk) indicate whether the client needs to perform public key operations, secret key operations or simple field operations. The Interaction column states N if clients interaction is limited to sending/receiving one message to/from the workers or Y otherwise. The work of the client in our solution is linear in m, but in all other rows m is a constant.

Note that the above solution requires that the client be able to sign messages. We would ideally like a solution which works under the sole assumption that the clients have an authenticated channel to each of the workers, as this is strictly more general and better models practice, where clients might authenticate themselves towards the computation providers using a simple password mechanism on top of a server authenticated secure transport layer. However, such a solution is not possible. It is well-known that broadcast among m parties without the use of signatures, requires that $> m/2$ of the parties are honest, and we want to tolerate that all but one worker is corrupted. We must therefore settle for a solution where clients need to have public keys for a signature scheme.

We finally note that the security property that we are trying to achieve in this section i.e., how to make sure that a single faulty client cannot make the computation stall, cannot be captured in the UC framework: since the adversary fully controls the network and the delivery of messages, in the UC framework the adversary can make the computation stall even if *no parties* are corrupted.

The resulting protocol, dealing with malicious clients, can be found in Appendix A.

4 Relationship To Previous Work

Here we discuss the relationship with some of the most relevant work in this area. Kamara et al. [KMR11] studied the problem of server-aided secure computation with relaxed security guarantees (i.e., non-collusion between corrupted parties). In their solution one of the clients also acts as a worker (i.e., it performs computation linear in the circuit size), therefore we interpret this as a setting with two workers (where one of the two happens to have an input as well). They also show how to transform protocols for secure delegation of computation into protocols in the server-aided model, but also this requires clients interacting with

each other. The server-aided model with non-colluding parties was also studied in [KMR12], which gives an efficient protocol based on Yao garbled circuits in this setting. The protocol still requires some interaction between the clients. Subsequently Carter et al. [CMTB13] claimed a number of improvements over Kamara et al., but also in their solution the clients need to interact further than the simple *upload input/download output* pattern of our scheme, and also require additional non-collusion assumptions. A recent work known as Whitewash [CLT14] improves this by presenting a protocol that is secure when the client and one of the workers collude. We do not see how to modify Whitewash to allow for multiple clients or workers. In addition, the output message of Whitewash has an extra factor k overhead, whereas in our protocol the output message from the workers to the clients is of the same size as the output of the function. Finally Peter et al. [PTK13] propose a protocol based on a variant of Paillier encryption for the setting of many clients and two workers. However the PTK protocol only offers passive security, and it is not clear whether it can be extended to more than 2 workers.

Compared to all above mentioned protocols, we find our solution to be 1) more elegant, as it decouples the problem of clients providing inputs to the problem of workers performing the computation 2) more flexible, as it supports any number of client and workers, and it allows the workers to chose the best possible protocols to perform the secure computation at hand, without having to modify the protocol at the client side 3) more efficient for the client, as the interaction pattern is minimal and the clients do not need to perform any cryptographic operations (only simple field operations) 4) more secure, as security holds up to $m-1$ actively corrupted workers and no non-collusion assumptions need to be made (that is, even when the corrupted parties share data among each other).

It is worth noting that these advantages are achieved without sacrificing the overall efficiency of the system: interestingly all previous solutions seem to obtain protocols that are *essentially as efficient* as the underlying protocol. This is also true in our case, as we only increase the size of the secure computation circuit by a single secure multiplication when there is no cheating (and $\log(n)$ additional multiplications if detection of corrupted clients is desired).

A summary of this comparison can be found in Table 1.

5 Acknowledgements

The authors are partially supported by the European Research Commission Starting Grant 279447 and the Danish National Research Foundation and The National Science Foundation of China (grant 61061130540) for the Sino-Danish Center for the Theory of Interactive Computation. The first author is supported by The Danish Council for Independent Research Starting Grant 10-081612. The third author is supported by The Danish Council for Independent Research (DFF) Grant 11-116416/FTP.

References

[BCD+09] Peter Bogetoft, Dan Lund Christensen, Ivan Damgård, Martin Geisler, Thomas P. Jakobsen, Mikkel Krøigaard, Janus Dam Nielsen, Jesper Buus Nielsen, Kurt Nielsen, Jakob Pagter, Michael I. Schwartzbach, and Tomas Toft. Secure multiparty computation goes live. In *Financial Cryptography*, pages 325–343, 2009.

[BDO14] Carsten Baum, Ivan Damgård, and Claudio Orlandi. Publicly auditable secure multi-party computation. SCN 2014. Available as Cryptology ePrint Archive, Report 2014/075, 2014. `http://eprint.iacr.org/`.

[BDOZ11] Rikke Bendlin, Ivan Damgård, Claudio Orlandi, and Sarah Zakarias. Semi-homomorphic encryption and multiparty computation. In *EUROCRYPT*, pages 169–188, 2011.

[BFR13] Michael Backes, Dario Fiore, and Raphael M. Reischuk. Verifiable delegation of computation on outsourced data. In *ACM Conference on Computer and Communications Security*, pages 863–874, 2013.

[BLW08] Dan Bogdanov, Sven Laur, and Jan Willemson. Sharemind: A framework for fast privacy-preserving computations. In *ESORICS*, pages 192–206, 2008.

[BSCG+13] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, Eran Tromer, and Madars Virza. Snarks for c: Verifying program executions succinctly and in zero knowledge. In *CRYPTO (2)*, pages 90–108, 2013.

[BV11] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) lwe. In *FOCS*, pages 97–106, 2011.

[Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *FOCS*, pages 136–145, 2001.

[CDF+08] Ronald Cramer, Yevgeniy Dodis, Serge Fehr, Carles Padró, and Daniel Wichs. Detection of algebraic manipulation with applications to robust secret sharing and fuzzy extractors. In *EUROCRYPT*, pages 471–488, 2008.

[CLT14] Henry Carter, Charles Lever, and Patrick Traynor. Whitewash: Outsourcing garbled circuit generation for mobile devices. Cryptology ePrint Archive, Report 2014/224, 2014. `http://eprint.iacr.org/`.

[CMF+14] Koji Chida, Gembu Morohashi, Hitoshi Fuji, Fumihiko Magata, Akiko Fujimura, Koki Hamada, Dai Ikarashi, and Ryuichi Yamamoto. Implementation and evaluation of an efficient secure computation system using 'r' for healthcare statistics. *Journal of the American Medical Informatics Association*, pages amiajnl–2014, 2014.

[CMTB13] Henry Carter, Benjamin Mood, Patrick Traynor, and Kevin Butler. Secure outsourced garbled circuit evaluation for mobile devices. In *Proceedings of the 22nd USENIX conference on Security*, pages 289–304. USENIX Association, 2013.

[DKL+13] Ivan Damgård, Marcel Keller, Enrique Larraia, Valerio Pastro, Peter Scholl, and Nigel P. Smart. Practical covertly secure mpc for dishonest majority - or: Breaking the spdz limits. In *ESORICS*, pages 1–18, 2013.

[DO10] Ivan Damgård and Claudio Orlandi. Multiparty computation for dishonest majority: From passive to active security at low cost. In *CRYPTO*, pages 558–576, 2010.

[DPSZ12] Ivan Damgård, Valerio Pastro, Nigel P. Smart, and Sarah Zakarias. Multiparty computation from somewhat homomorphic encryption. In *CRYPTO*, pages 643–662, 2012.

[DS83] Danny Dolev and H. Raymond Strong. Authenticated algorithms for byzantine agreement. *SIAM J. Comput.*, 12(4):656–666, 1983.

[Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, pages 169–178, 2009.

[GGP10] Rosario Gennaro, Craig Gentry, and Bryan Parno. Non-interactive verifiable computing: Outsourcing computation to untrusted workers. In *CRYPTO*, pages 465–482, 2010.

[GHRW14] Craig Gentry, Shai Halevi, Mariana Raykova, and Daniel Wichs. Outsourcing private RAM computation. *IACR Cryptology ePrint Archive*, 2014:148, 2014.

[HL10] Carmit Hazay and Yehuda Lindell. *Efficient Secure Two-Party Protocols: Techniques and Constructions*. Springer-Verlag, 2010.

[HLP11] Shai Halevi, Yehuda Lindell, and Benny Pinkas. Secure computation on the web: Computing without simultaneous interaction. In *CRYPTO*, pages 132–150, 2011.

[KMR11] Seny Kamara, Payman Mohassel, and Mariana Raykova. Outsourcing multi-party computation. *IACR Cryptology ePrint Archive*, 2011:272, 2011.

[KMR12] Seny Kamara, Payman Mohassel, and Ben Riva. Salus: a system for server-aided secure function evaluation. In *ACM Conference on Computer and Communications Security*, pages 797–808, 2012.

[KMRS14] Seny Kamara, Payman Mohassel, Mariana Raykova, and Saeed Sadeghian. Scaling private set intersection to billion-element sets. In *Financial Crypto*, 2014.

[KS08] Vladimir Kolesnikov and Thomas Schneider. Improved garbled circuit: Free xor gates and applications. In *ICALP (2)*, pages 486–498, 2008.

[LTV12] Adriana López-Alt, Eran Tromer, and Vinod Vaikuntanathan. On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In *STOC*, pages 1219–1234, 2012.

[NNOB12] Jesper Buus Nielsen, Peter Sebastian Nordholt, Claudio Orlandi, and Sai Sheshank Burra. A new approach to practical active-secure two-party computation. In *CRYPTO*, pages 681–700, 2012.

[PHGR13] Bryan Parno, Jon Howell, Craig Gentry, and Mariana Raykova. Pinocchio: Nearly practical verifiable computation. In *IEEE Symposium on Security and Privacy*, pages 238–252, 2013.

[PTK13] Andreas Peter, Erik Tews, and Stefan Katzenbeisser. Efficiently outsourcing multiparty computation under multiple keys. *IEEE Transactions on Information Forensics and Security*, 8(12):2046–2058, 2013.

A The Protocol

Figure 5 shows our protocol, dealing with malicious clients as explained in Section 3.4. It does not include the eventual output consensus of Section 2.3, but adding this should be straightforward. Figure 5 assumes that a public encryption scheme (E, D) has been set up such that each worker W_i holds a private decryption

Algorithm FIND: Input: $([a_1], [a_2], \ldots, [a_q])$, a list of q secret shared values.

1. $[s] \leftarrow \mathsf{Rand}()$;
2. $[\beta] = [s] \cdot [\sum_{j \in 1,\ldots,q} a_q]$;
3. $\beta \leftarrow \mathsf{Open}([\beta])$;
4. If $\beta = 0$ return $\{\}$;
5. Else, if $q = 1$ return $\{[a_1]\}$;
6. Else, let $r = \lfloor q/2 \rfloor$ and return $\mathsf{FIND}([a_1], \ldots, [a_r]) \cup \mathsf{FIND}([a_{r+1}], \ldots, [a_q])$.

Figure 4. Algorithm to identify potentially inconsistent client input.

key d_i while all other parties, workers as well as clients, hold the corresponding public encryption key e_i. Furthermore, Figure 5 assumes a broadcast primitive; we use $\mathsf{BROADCAST}(id, msg, \{R_i\})$ to denote that a party P broadcasts msg to a set of receivers $\{R_i\}$ while $msg \leftarrow \mathsf{RECEIVE}(id, P)$ is used by a party to receive msg broadcast by P. As noted, broadcast can be realized using Dolev-Strong [DS83]. A recursive algorithm FIND, listed in Figure 4, is used for identifying potentially inconsistent client input. FIND may result in $O(n \log n)$ multiplications. In order to have $O(n)$ multiplications, let e be any constant, e.g. $e = 2$, and modify FIND such that a global counter δ is increased for every multiplication (i.e., $\delta := \delta + 1$ each time Step 2 of FIND is executed). If at some point $\delta > e \cdot n$, abort the entire recursion and compute $\cup_{j=1}^{n} \mathsf{FIND}([\alpha^{j,w}])$ as result instead. The default value used for inconsistent client input is 0.

All algebraic notation denotes operations in K.

Clients Input Phase: Each client C^j with input x^j does:

1. Pick random $y^{j,1}$. Compute $y^{j,2} = x^j - y^{j,1}$;
2. For $w \in \{1,2\}$ do the following:
   (a) Pick random $\{y_i^{j,w}\}_{i=1}^m$ from K s.t. $\sum_{i=1}^m y_i^{j,w} = y^{j,w}$;
   (b) Pick random $\{k_i^{j,w}\}_{i=1}^m$ from K; let $k^{j,w} := \sum_{i=1}^m k_i^{j,w}$;
   (c) Pick random $\{r_i^{j,w}\}_{i=1}^m$ from K; let $r^{j,w} := \sum_{i=1}^m r_i^{j,w}$;
   (d) Compute $t^{j,w} = \mathrm{Tag}_{k^{j,w}}(x^{j,w}) = k^{j,w} \cdot x^{j,w}$;
   (e) Pick random $\{t_i^{j,w}\}_{i=1}^m$ from K s.t. $\sum_{i=1}^m t_i^{j,w} = t^{j,w}$;
   (f) Compute the values $v_i^{j,w} := (y_i^{j,w}, t_i^{j,w}, k_i^{j,w}, r_i^{j,w})$ for $i \in 1, \ldots, m$;
   (g) Compute $\gamma_i^{j,w} := E_{e_i}((i, j, v_i^{j,w}); s_i^{j,w})$ for $i \in 1, \ldots, m$;
   (h) Broadcast $\gamma_i^{j,w}$ to the workers, i.e., invoke $\mathsf{BROADCAST}(\texttt{input}, (\gamma_i^{j,w})_{i=1}^m, \{W_i\}_{i=1}^m)$.

Workers Input Phase: Each worker W_i does:

1. Upon $(\gamma_1^{j,w}, \gamma_2^{j,w}, \ldots, \gamma_m^{j,w}) \leftarrow \mathsf{RECEIVE}(\texttt{input}, C^j)$ do:
   (a) Compute $(i', j', v_i^{j,w}, s_i^{j,w}) = D_{d_i}(\gamma_i^{j,w})$;
   (b) If $i' \neq i$ or $j' \neq j$, $\mathsf{BROADCAST}(\texttt{bad-input}, (i', j', v_i^{j,w}, s_i^{j,w}), \{W_l\}_{l=1}^m)$.
2. Upon $(l', j', v_l^{j,w}, s_l^{j,w}) \leftarrow \mathsf{RECEIVE}(\texttt{bad-input}, W_l)$, if $\gamma_l^{j,w} = E_{e_l}((l', j', v_l^{j,w}); s_l^{j,w})$ and if it is also the case that $l' \neq l$ or $j' \neq j$, set $[x^j] \leftarrow \mathsf{Input}(W_i, 0)$. Otherwise, abort the protocol.

Workers Computation Phase: Upon (*eval*) each worker W_i does:

1. For $w \in \{1,2\}$ do:
   (a) For all $j \in \{1, \ldots, n\}$ do:
      i. Compute $[y_i^{j,w}] \leftarrow \mathsf{Input}(W_i, y_i^{j,w})$; $[t_i^{j,w}] \leftarrow \mathsf{Input}(W_i, t_i^{j,w})$; $[k_i^{j,w}] \leftarrow \mathsf{Input}(W_i, k_i^{j,w})$; $[r_i^{j,w}] \leftarrow \mathsf{Input}(W_i, r_i^{j,w})$;
      ii. Compute: $[y^{j,w}] := \sum_{i=1}^m [x_i^{j,w}]$;
      iii. Compute: $[t^{j,w}] := \sum_{i=1}^m [t_i^{j,w}]$;
      iv. Compute: $[k^{j,w}] := \sum_{i=1}^m [k_i^{j,w}]$;
      v. Compute: $[r^{j,w}] := \sum_{i=1}^m [r_i^{j,w}]$;
      vi. $k^{j,w} \leftarrow \mathsf{Open}([k^{j,w}])$;
      vii. $[\alpha^{j,w}] := [t^{j,w}] - k^{j,w} \cdot [y^{j,w}]$;
   (b) Compute $I^w := \mathsf{FIND}([\alpha^{1,w}], [\alpha^{2,w}], \ldots, [\alpha^{n,w}])$.
2. Let $I := I^1 \cup I^2$. For $j \notin I$, set $[x^j] := [y^{j,1}] + [y^{j,2}]$. For $j \in I$, do as follows:
   (a) If $j \in I^1$, let $w := 1$, else $w := 2$. W_i now reveals $v_i^{j,w} := (y_i^{j,w}, t_i^{j,w}, k_i^{j,w})$ as follows:
   (b) W_i invokes $\mathsf{BROADCAST}(\texttt{reveal}, (v_i^{j,w}, s_i^{j,w}), \{W_l\}_{l=1}^m)$;
   (c) Upon $(v_i^{j,w}, s_i^{j,w}) \leftarrow \mathsf{RECEIVE}(\texttt{reveal}, W_i)$, a worker W_l aborts if $\gamma_i^{j,w} \neq E_{e_i}((i, j, v_i^{j,w}); s_i^{j,w})$;
   (d) Upon receiving all $v_l^{j,w} := (y_l^{j,w}, t_l^{j,w}, k_l^{j,w})$ for $l \in \{1, \ldots, m\}$, compute $y^{j,w} := \sum_l y_l^{j,w}$, $t^{j,w} := \sum_l t_l^{j,w}$, $k^{j,w} := \sum_l k_l^{j,w}$;
   (e) W_i then computes $y'^{j,w} \leftarrow \mathsf{Open}([y^{j,w}])$; $t'^{j,w} \leftarrow \mathsf{Open}([t^{j,w}])$; $k'^{j,w} \leftarrow \mathsf{Open}([k^{j,w}])$;
   (f) If $(y'^{j,w}, t'^{j,w}, k'^{j,w}) \neq (y^{j,w}, t^{j,w}, k^{j,w})$, W_i aborts, otherwise $[x^j] \leftarrow \mathsf{Input}(W_i, 0)$.
3. Compute: $([z^1], \ldots, [z^n]) := f([x^1], \ldots, [x^n])$;
4. Compute: $[c^j] = [z^j] + [r^j]$ for all $j \in \{1, \ldots, n\}$;
5. $c_i^j \leftarrow \mathsf{Open}([c^j])$ for all $j \in \{1, \ldots, n\}$;
6. For $j \in \{1, \ldots, n\}$ send c_i^j to C^j.

Client Output Phase: Each client C^j does:

1. Once c_i^j have been received from all workers $\{W_i\}_{i=1}^m$, if $c_1^j, c_2^j, \ldots, c_m^j$ are not all equal, output $\texttt{abort}$;
2. Else, let $c^j := c_1^j$ and output $z^j := c^j - r^j$.

Figure 5. The protocol coping with malicious clients.