Faster Maliciously Secure
Two-Party Computation Using the GPU
Full version

Tore Kasper Frederiksen, Thomas P. Jakobsen, and Jesper Buus Nielsen *
July 11, 2014

Department of Computer Science, Aarhus University
{jot2re|tpj|jbn}@cs.au.dk

Abstract We present a new protocol for maliciously secure two-party computation based on cut-and-choose of garbled circuits using the recent idea of "forge-and-loose", which eliminates around a factor 3 of garbled circuits that needs to be constructed and evaluated. Our protocol introduces a new way to realize the "forge-and-loose" approach, which avoids an auxiliary secure two-party computation protocol, does not rely on any number theoretic assumptions and parallelizes well in a same instruction, multiple data (SIMD) framework.

With this approach we prove our protocol universally composable-secure against a malicious adversary assuming access to oblivious transfer, commitment and coin-tossing functionalities in the random oracle model.

Finally, we construct, and benchmark, a SIMD implementation of this protocol using a GPU as a massive SIMD device. The findings compare favorably with all previous implementations of maliciously secure, two-party computation.

1 Introduction

Background. Secure two-party computation (2PC) is the area of cryptography concerned with two mutually distrusting parties who wish to securely compute an arbitrary function with private output based on their independent and private input. A bit more formally Alice has the input x, Bob the input y, and they wish to compute the function $f(x, y) = z$ while ensuring such properties as correctness, privacy, and input independence.[1]

2PC was introduced in 1982 by Andrew Yao [Yao82, Yao86], specifically for the *semi-honest* case where both parties are assumed to follow the prescribed protocol and only try to compromise security by extracting information from their own views of the protocol execution. Yao showed how to construct a semi-honestly secure protocol using a technique referred to as the *garbled circuit approach*. This approach involves having one party (*the constructor*), say, Alice, encrypt, or "garble", a Boolean circuit computing the desired functionality. The "garbling" involves first choosing two random keys for each wire in the circuit, one representing a value of 0 and another representing a value of 1. Each gate in the garbled circuit is then constructed such that Bob (*the evaluator*), given only one key for each input wire, can compute only one key for the output wire, namely the key corresponding to the bit that the gate is supposed to output, for example the logical AND of the two input bits. Alice sends the garbled circuit to Bob and makes sure that for each input wire, Bob learns only one of that wire's keys. For Alice's own input, she simply sends these keys to Bob. For Bob's input, they use an oblivious transfer (OT) protocol such that Bob learns one key for each input wire corresponding to his own input, without Alice learning Bob's input. Now, given one key for each input wire, Bob can then evaluate the whole garbled circuit, gate by gate, while at the same time, he cannot learn which bits flow

---

[1] Who should learn the output z can differ and in the most general case there might be a specific output for Alice and one for Bob, i.e., $f(x, y) = z = (z_A, z_B) = (f_A(x, y), f_B(x, y))$ where Alice should only learn z_A and Bob only z_B. In this paper we only consider the case where Bob learns the entire output z. If output for Alice is also needed then any of several general approaches can be used. We consider some of these approaches in Appendix H.

on the wires. Only when he reaches the output wires, he uses some auxiliary information to learn which bits the keys encode. See [LP09] for a thorough description of Yao's scheme.

A major reason why Yao's original protocol is only secure against a semi-honest adversary is that Bob cannot be sure that the garbled circuit he receives from a Alice has been garbled correctly. One way to ensure correct garbling is with a *cut-and-choose* approach: Instead of sending one garbled circuit, Alice sends several independently garbled versions of the circuit to Bob. Bob then randomly selects some of these, called the *check circuits*, where all the keys on the input wires are revealed to Bob, allowing him to verify that the check circuits do indeed compute the correct function f. If enough circuits are checked and all found to be correct, a large fraction of the remaining circuits, called *evaluation circuits*, will also be correct.

Doing cut-and-choose on several garbled circuits introduces some other issues that have to be dealt with in order to obtain malicious security. First, there may still be a few incorrect circuits among the evaluation circuits. Also, since there are now many circuits, some mechanisms must ensure *input consistency* amongst the circuits: Alice and Bob must give the same input to all the circuits. Information about Alice's input will leak to Bob if he gets to evaluate the circuit on different inputs. But information might also leak to Alice, depending on the function the circuit is computing, if she gives inconsistent input. Another perhaps more subtle issue that arises is a certain kind of *selective failure attack*: Since Alice will supply Bob with the keys in correspondence with his input bits through an OT Alice can simply input garbage for one of the keys, e.g, the 0-key for the first bit of his input. If Bob aborts the protocol, because he cannot evaluate a garbled circuit when one of the keys is garbage, Alice will know that the first bit of his input is 0! On the other hand, if he does not abort the protocol then his first bit must be 1!

Cut-and-choose of Garbled Circuits. In recent years several protocols have been presented which obtain malicious security by combining Yao's original protocol with cut-and-choose on the circuits [LP07,PSSW09,LP11,sS11,KSS12, Bra13, FN13, HKE13, Lin13, MR13, SS13]. They mainly differ in the way they handle the above issues and, as a consequence, in how many circuits are needed to send from Alice to Bob. We call the amount of circuits needed to be sent from Alice to Bob for the *replication factor*. Most of these approaches are based on an analysis saying that if the fraction of check circuits is high enough and if the check circuits are all found to be correct, then a *majority* of the remaining evaluation circuits will also be correct except with negligible probability. This allows Bob to compute the correct output as the majority of the outputs from the evaluation circuits.

Recent results [Bra13, HKE13, Lin13] manage to ensure that Alice only succeeds in cheating if *all* check circuits are honestly constructed and *all* evaluation circuits are maliciously constructed. This happens with noticeably less probability than only the majority of the evaluation circuits are honestly constructed, thus making it possible to run a protocol with noticeably less garbled circuits. The main idea of [Bra13] and [Lin13] (called "forge-and-loose") is to make sure that if Alice cheats, that is, if two evaluation circuits yield different results, then this enables Bob to learn Alice's private input x. He can then use x to compute locally, unencrypted, the correct input $f(x, y)$. In [Lin13] Bob learns Alice's input through an auxiliary secure protocol execution, using as input part of the transcript of the "original" protocol execution, as proof that Alice is corrupt. In [Bra13] a special kind of commitments are used which leaks some auxiliary information if the outputs of two garbled circuits differ. This auxiliary information can then be used to learn Alice's input. The idea of [HKE13] is to have both parties play the role of the circuit constructor *and* circuit evaluator respectively in two simultaneous executions of a cut-and-choose of garbled circuits protocol.

Other Approaches to Malicious Security. Several other approaches to maliciously secure two party computation exist, based on Yao's garbled circuits as well as other novel constructions.

In [NO09] Nielsen and Orlandi introduced the notion of cut-and-choose at the gate level in order to achieve malicious security for Yao's garbled circuits, popularly referred to as the LEGO approach. LEGO gave an asymptotic increase in efficiency of $O(\log(|C|))$ where $|C|$ is the amount of gates in the circuit to compute. However, their approach relied on heavy group based computations (as the protocol required homomorphic commitments) for each gate in the circuit.

In the later work [FJN+13] a new approach for constructing xor-homomorphic commitments based on error correcting codes was introduced. These xor-homomorphic commitments made it possible to use the LEGO approach without doing group operations for each gate, unfortunately at the price of large constants.

Another approach, *not* using garbled gates, was introduced in [NNOB12] where heavy usage of oblivious transfers was used to construct verified and authenticated Boolean gates. To achieve efficiency the authors constructed a highly efficient maliciously secure OT extension, which took a few random OTs as "seeds" and then used a random oracle to extend these to polynomially many random OTs.

Yet another approach, not necessarily using Yao's garbled circuits, is the idea of "MPC-in-the-head" from [IKOS07, IPS08] where the parties run a semi-honestly secure two-party protocol, but within this they emulate $n > 2$ parties that run a virtual second protocol secure against a malicious minority. The final protocol is then maliciously secure. Another idea is to base the protocol on preprocessed constrained randomness as in [DZ13].

If one is willing to accept a weaker kind of "malicious" security one can get significant improvements by using the "dual execution" approach where only a single garbled circuit is constructed and evaluated by *both* parties. [MF06, HKE12, MR13].

Besides all of the above approaches a whole other subfield of secure multi-party computation exists where the function to be evaluated is arithmetic (which can clearly also be used to evaluate Boolean circuits when the field for the arithmetic operations is $\mathbb{F}_2$). Some of the more recent works in this area include [BDOZ11] (semi-honest security) and [DPSZ12, DKL$^+$13] (malicious security).

Motivation. The area of 2PC and multi-party computation, MPC, (when more than two parties supply input) is very interesting as efficient solutions yield several practical applications. The first "real world" usage of MPC was described in [BCD$^+$09] where MPC was used to decide the market clearing price of sugar beet quotas in Denmark. Still, several other applications of 2PC and MPC exists such as voting, anonymous identification, privacy preserving database queries etc. Thus we believe that it is highly relevant to find practically efficient protocols for 2PC and MPC.

Our Results. We present a new maliciously secure two-party protocol based on cut-and-choose of garbled circuits. The protocol combines previous ideas, mainly from [Bra13, Lin13, HKE13] and [FN13, SS13] with new optimizations to a protocol with the following benefits:

1. A small replication factor for the same reasons as [Bra13, Lin13].[2] But unlike [Lin13] our protocol eliminates the need for running an auxiliary protocol and unlike [HKE13] we eliminate the need of both parties both constructing *and* evaluating $O(s)$ garbled circuits.
2. A very lightweight approach to ensure that Alice gives consistent input to all garbled circuits used for evaluation.
3. Reliance only on lightweight primitives (in practice a hash function) and few OTs.
4. A large degree of same instruction, multiple data (SIMD) parallelization, and therefore direct benefits from access to a large amount of computation cores.
5. A security proof in the Universally Composable (UC) model [Can01].[3] We also show how to realize a highly efficient version of the protocol in the Random Oracle Model (ROM) without number theoretic assumptions, only assuming access to an OT functionality.

Finally, we also provide an implementation of the protocol written in C and CUDA, using a consumer grade GPU for exploiting the inherent parallelism in the protocol.

Outline. We start by going through the overall ideas of our scheme, along with the primitives we need in Section 2. We then continue with a high level description of our protocol in Section 3. A discussion of our implementation along with experimental results are given in Section 4 and we finally conclude the paper in Section 5 by giving a discussion of the findings from our implementation and directions for future work. In Appendix A we give two tables describing the variable and symbols used throughout the paper. Detailed description of the primitives we need are given in Appendix B and Appendix C, then in Appendix D a fully detailed description of our protocol is given and in Appendix E we give a full proof of security with supplementing details in Appendix F, Appendix G and Appendix H.

---

[2] We do not strictly obtain the same replication factor as [Bra13, Lin13], where the replication factor is s, when s is the statistical security parameter. Our replication factor is slightly larger as in [HKE13]. See the full version of this paper for details [FJN14].

[3] Specifically we provide a UC-proof of our protocol in a hybrid/random oracle model with OT, commitments and coin tossing, and a secure garbling scheme with free-xor [KS08].

Notation. We assume Alice is the constructor and Bob the evaluator and that the functionality they wish to compute is $f(x, y) = z$, where Alice gives input x, Bob gives input y and only one party, Bob, receives the output z. We denote the amount of bits of x as $|x| = m$, the amount of bits of y as $|y| = n$ and the amount of bits of z as $|z| = o$. We denote the Boolean circuit computing the functionality $f(\cdot, \cdot)$ by C and assume it consist of gates with fan-in 2 and fan-out 1 such as AND, XOR, OR etc. Furthermore, we say the size of C, $|C|$, is the amount of Boolean gates. When considering a garbled version of C we add a tilde, i.e., $\tilde{C}$. In a garbled circuit we define the *semantic* value of a key to be the bit it represents. We furthermore define the *plain* values or functions to be the unencrypted, unhidden values or functionalities, e.g., the actual function $f(x, y) = z$ with inputs x and y and output z. We use s to represent a statistical security parameter, and we let κ be a computational security parameter. Technically, this means that for any fixed s and any polynomial time bounded adversary, the advantage of the adversary is $2^{-s} + \mathrm{negl}(\kappa)$ for a negligible function negl. I.e., the advantage of any adversary goes to 2^{-s} faster than any inverse polynomial in the computational security parameter. If $s = \Omega(\kappa)$ then the advantage is negligible. If $s = O(\log(\kappa))$ or $s = O(1)$ we only have covert security (see [HL10] for a detailed description of this concept).

Let $[n]$ denote the set of integers $\{1, 2, \ldots, n\}$ and let $H(\cdot)$ denote a hash function modeled as a random oracle giving κ bits of output. We use subscript to denote index, i.e. x_i denotes the i'th bits of a vector x and $m_{i,j}$ denotes the i'th row and j'th column of the matrix m. Finally we let $x \in_R S$ denote that x is sampled uniformly at random from the set S. An overview of the various variables and parameters along with their meaning is given in Appendix A.

2 The Big Picture

We assume the reader is familiar with free-xor Yao circuits [KS08] for maliciously secure computation, otherwise see Appendix B. The overall strategy is as follows:

1. *Garbling.* Alice garbles a bunch of circuits and sends them to Bob.
2. *Oblivious Transfer.* The parties then engage in an OT protocol, secure against malicious adversaries, such that Bob, for each garbled circuit, learns the input keys corresponding to his input to the function f.
3. *Cut-and-Choose.* The parties then use a coin tossing protocol to select around half of the garbled circuits as *check* circuits and the remaining as *evaluation* circuits. Bob then verifies that the check circuits have been constructed correctly and receives from Alice the keys corresponding to her input to the evaluation circuits.
4. *Evaluation.* Bob will then use his own input keys, which he learned from the OT, along with Alice's input keys to evaluate all the evaluation circuits. Using the results from the evaluated circuits he finds the correct output.

This generic scheme leaves open how to handle the problems that arise due to the fact that the parties might act maliciously and due to multiple circuits being garbled and evaluated instead of just one circuit. These problems are (1) selective failure attacks, (2) consistency of the parties' input and (3) how to know which output is the correct one in the evaluation phase. Many different approaches exist to solve these problems with different levels of efficiency and security assumptions. In the following sections we shortly discuss various solutions and in particular, how we solve these issues.

Avoiding Selective Failure Attacks. To ensure privacy, a corrupted party should not be able to deduce anything about the other party's private input based on the other party's behavior. If care is not taken, 2PC based on cut-and-choose of circuits can easily lead to a vulnerability of this kind: If, e.g., Alice inputs another 1-key in the OT for a certain input wire than she uses when garbling the corresponding circuit, and that circuit happens to be an evaluation circuit, she can deduce that Bob's input on that wire is 1 if he aborts. Such an attack is known as a *selective failure attack* and was first pointed out by [KS06, MF06],

Mohassel and Franklin [MF06] suggested thwarting the attack by means of *committing OT*'s, which makes it possible for Bob to verify Alice's input to the OT against the keys in the check-circuits. Lindell and Pinkas [LP07] thwarted the attack by increasing the amount of input wires of Bob, letting Bob replace his original input with constrained randomness, where certain input wires xor'ed together would give his real input. This eliminates the attack

since it is hard for Alice to "guess" Bob's choice of randomness, which is needed in order for her to successfully complete her attack. Unfortunately, this also increases the size of Bob's input from n bits to $\max(4n, 8s)$ bits. Lindell and Pinkas' technique was later refined by shelat and Shen to give a smaller increase in the amount of input bits needed by Bob [SS13].

Another approach works by introducing a tighter coupling between the OTs and the rest of the protocol. Lindell and Pinkas [LP11] introduced the *single-choice batch cut-and-choose oblivious transfer* based on Diffie-Hellman pseudo-random synthesizers, which were used to "glue" together the OTs and the garbled circuits. The pseudo-random synthesizers were also used to ensure consistency of the parties' inputs to the different garbled circuits.

In our solution we handle the selective failure problem in the same manner as in [LP07]. The approach involves increasing the amount of input bits of the evaluator from a string y of n bits to a string $\bar{y}$ of $\bar{n} = \max(4n, 8s)$ bits. More specifically, what we do is to have Bob choose a random binary matrix $M^{\mathsf{Sec}} \in_R \{0,1\}^{\bar{n}} \times \{0,1\}^n$ and his input randomly, i.e., $\bar{y} \in_R \{0,1\}^{\bar{n}}$, but under the constraint that $M^{\mathsf{Sec}} \cdot \bar{y} = y$ where y is his "true" input. Thus the augmented functionality computes exactly the same as the original. Still, the idea is that if a selective failure attack is done, using the augmented function will not leak any useful information: as the entire vector $\bar{y}$ is random, learning a few bits of $\bar{y}$ will only give the adversary a negligible advantage in learning one of the constructor's true input bits. This follows from the fact that the other bits of $\bar{y}$ will be used to hide each of the actual bits of y. The details and a full proof of security of this approach can be found in [LP07].

Ensuring Consistency of Bob's Input. Making just multiple calls to a basic protocol for OT, Bob could input different choice bits for each of the circuits, which would allow him to evaluate the function on different inputs. That is, he would learn both $f(x, y)$ and $f(x, y')$ for two different inputs y and y' of his own choice, which obviously leaks too much information to Bob.

To avoid this we must make sure that for each of the bits y_i in Bob's input, he gives the same choice bit for all of the garbled circuits in the OTs. This can easily be achieved by doing fewer OTs, but of longer bit strings: For each OT Alice inputs a concatenation of the 0-, respectively 1-keys for a given input wire for all circuits, Bob inputs only one choice bit and receives the keys for all the circuits corresponding to his choice bit on that given input wire. This approach is used in most previous results [LP07, LP11, sS11, FN13, Lin13, SS13]. In the ROM this can be very efficiently realized using black box access to a one-out-of-two OT of strings of length s [FN13, SS13]. For more details and its realization in the ROM, see Appendix C.

Ensuring Consistency of Alice's Input. To see why there is also a problem if Alice is inconsistent in her inputs, assume the functionality to be computed is the inner product, that Alice inputs the strings $\langle 100 \ldots 0 \rangle$, $\langle 010 \ldots 0 \rangle, \ldots, \langle 000 \ldots 1 \rangle$ and that she is the party who is supposed to learn the output of the computation. In this case what she will learn is the majority of bits in Bob's input string, which is clearly too much information.

One could perhaps think that the problems arising by Alice giving inconsistent inputs can be solved by having Bob abort if the results of the circuit evaluations diverge. But aborting based on diverging output may leak information about Bob's secret input to Alice. To see this, note that the function f might give the same outcome for several different inputs from Alice, but *only* if Bob's input comes from some particular subset of his possible inputs. Alice would learn whether Bob's input belongs to this subset based on whether Bob aborts or not. See [LP07] for a concrete example of such a function.

A similar attack is possible if Bob instead of aborting, randomly picks one of the different outputs from the circuits: For certain functions f, Alice would be able to control what the majority of outputs would be and to let this depend on Bob's secret input.[4]

---

[4] If Bob aborts we have an example of a *selective failure* attack. If he instead takes majority or picks a random result when the outputs diverge it is sometimes called a *selective input* attack. There seems to be some confusion about these terms. Selective input attacks are used to describe the attacks where Alice can deduce information about Bob's private input by submitting different inputs to different circuits. So, enforcing consistent input from Alice is the same as preventing all selective input attacks, and – if Bob aborts when results diverge – not enforcing consistent input from Alice would make the protocol vulnerable to a specific selective failure attack.

Since our aim is to allow secure evaluation of *any* function (and since it is not obvious which functions allow leakage when Alice gives inconsistent inputs) we therefore must ensure that all circuit evaluations are based on the same inputs from Alice.

Lindell and Pinkas [LP07] embedded a consistency check of Alice's input into the cut-and-choose step itself, resulting in $O(ms^2)$ commitments. Shelat and Shen [sS11] ensures input consistency using another approach based on claw-free functions.

We follow a different approach, introduced independently by Frederiksen and Nielsen [FN13] and shelat and Shen [SS13]. The idea is to augment the functionality computed by the garbled circuits such that Alice's input, apart from being used in the original computation of f, is also fed into a universal hash function decided by Bob. The output of both the original functionality f and the hash digest of Alice's input is then learned by Bob. We can now safely let Bob abort if any of these digests diverge without leaking anything to Alice about his input. Intuitively, this approach is secure if the output of the augmented functionality does not leak information about Alice's input and Alice cannot find two inputs whose digests collide. We achieve the first property by having Alice give some auxiliary random input, which one-time pads the output of the hash function. The second property will follow directly from the nature of a universal hash function if Alice is committed to her input choices before potentially learning the specific hash function used.

A specific, but simple way to construct such an augmentation is as follows [FN13]: Assume the functionality we wish to compute is defined by f as $f(x, y) = z$. We then define a new function f' as $f'(x', y') = f'(x\|a, y\|b) = z\|c$ where $a \in_R \{0,1\}^s$, $b \in_R \{0,1\}^{m+s-1}$ and $c \in \{0,1\}^s$. To compute c we define a matrix $M^{\mathsf{In}} \in \{0,1\}^s \times \{0,1\}^{|x|}$ where the i'th row is the first $m = |x|$ bits of $b << i$ where $<<$ denotes the bitwise left shift. Specifically the j'th bit of the i'th row is the $i + j$'th bit of the binary vector b, i.e., $M^{\mathsf{In}}_{i,j} = b_{i+j-1}$. The computation of c is then defined as $c = (M^{\mathsf{In}} \cdot x) \oplus a$, assuming all binary vectors are in column form. Now see that f computes the same function as f', but requires s extra random bits of input from Alice and $m + s - 1$ extra random bits from Bob. The s extra output bits will work as digest bits (since $(M^{\mathsf{In}} \cdot x) \oplus a$ is actually a universal hash function) and can be used to check that Alice is consistent with her inputs to the circuits, by verifying that they are the same in all the garbled circuits which Bob evaluates.

Next, [SS13] observed that it is in fact not needed to have Bob specify the specific hash function as part of his input. Instead he can send Alice the vector b defining the matrix M^{In} and have her garble the circuits to compute the hash function this vector specifies. Evaluation of the garbled universal hash function can now be done using only xor operations, which we can do for "free" with the free-xor approach. Thus the binary string $b \in \{0,1\}^{m+s-1}$ defines a family of universal hash functions mapping m bits to s bits. In turn, a sampling of such a function is simply a random choice of the string b. We will call a specific function from this family for $\mathbb{H}^{\mathsf{In}}$ and call the actual circuit augmentation for $\mathbb{H}^{\mathsf{In}} \oplus a$ where a will be Alice's auxiliary input sampled from $\{0,1\}^s$.

A remaining problem with the construction is that Alice can now try to find a collision for the hash function before she gives her input keys to Bob. This is a significant problem as collisions for universal hash functions can be easy to find. However, if we let Alice commit to her input and *then* let Bob reveal his choice of hash function to her then she will only have negligible probability in the size of the digest of finding a collision.

It should be noted that in [SS13] another universal hash function was used, which required $2s + \log(s)$ auxiliary random bits of input from Alice, unlike only s in [FN13]. However, the authors of [FN13] did not observe that it was enough to send M^{In} in plain, nor did they actually *prove* their construction secure. In Appendix E we prove that the universal hash function used in [FN13] is secure in our setting when M^{In} is sent in plain after Alice commits to her inputs.

We follow the approach of circuit augmentation outlined above. Concretely we implement the consistency check by having Alice commit to her input before the cut-and-choose phase. Then we have Bob give her the specification of the hash function. Alice then garbles the circuit with the hash function augmentation. When Bob later has evaluated the evaluation circuits, he aborts if the hash values from any two of the augmentations of these circuits differ.

Computing the Correct Output. Despite the input consistency of Alice being enforced as discussed above, Bob may still end up with more than one distinct output when he evaluates the evaluation circuits. This is due to the fact that

the cut-and-choose technique does not ensure that all the remaining evaluation circuits are good. With non-negligible probability some of them may not evaluate at all and some may evaluate, but result in diverging output.

Given several different outputs, Bob cannot just abort the computation or choose a random result without becoming vulnerable to the same kind of selective failure attack as we discussed in connection with the input consistency. However, Lindell and Pinkas [LP07] showed that for a statistical security parameter s, if the number of garbled circuits, i.e., the replication factor, is greater than cs for some constant c and if half are chosen as check circuits, then, except with negligible probability, the majority of the outputs are well-defined and are in fact the correct output.

For realistic circuit sizes, the replication factor has proven to be very important for the performance of the overall protocol [PSSW09, KSS12, FN13]. Thus recent effort has been made to reduce the replication factor. As a start, [PSSW09] claimed that [LP07] only requires a replication factor $\ell \geq 4s$ (to achieve failure probability at most 2^{-s}). The analysis was later improved by [LP11] showing that one only needed a replication factor $\ell \geq 3.22s$. Finally, [sS11] refined the analysis further by showing that taking 60% of the circuits to be check circuits, a replication factor $\ell \geq 3.13s$ is sufficient.

All of these results are based on taking the majority of the outputs from the evaluation circuits. As mentioned in the introduction, recently Brandão [Bra13] and Lindell [Lin13] independently showed how to avoid taking majority by instead enabling Bob to learn Alice's private input if the results diverge, and thereby letting Bob compute the correct result in "clear". Doing so allows for a replication factor of only s to achieve an error probability of 2^{-s}, which is a considerable improvement. Independently Huang *et al.* [HKE13] showed the same, but with a slightly higher replication factor. In particular, doing the protocol with a replication factor ℓ gives an error probability of at most $2^{-\ell+O(\log \ell)}$. For comparison, note that for a typical application with $s = 40$, [LP11] and [sS11] need 129 and 125 circuits respectively, while [Bra13, Lin13] requires only 40 circuits to be garbled.

Our protocol is based on the same idea as [Bra13, Lin13], namely to make it possible for Bob to recover Alice's plain input x if any inconsistency arrives in the evaluation circuits and, in turn, let Bob evaluate the functionality in plain. To do this, [Lin13]relies on a secondary secure computation "inside" the protocol in order for Bob to learn x. While being independent of the size of the original circuit, the secondary secure computation still introduces a significant overhead. Furthermore, [Lin13] requires a number of modular exponentiations linear in the input size and reliance on the DDH assumption in order to ensure Alice inputs the correct x to the secondary computation. In [Bra13] Bob gets to recover Alice's input by having her commit to her input keys and the keys on the circuit output wires using trapdoor commitments where a the trapdoor can be efficiently computed if two different output keys are achieved for the same wire, but in different garbled circuits.

Our main theoretical contribution arrives in the way we allow Bob to find the correct output based on the evaluation of the garbled circuits. We do so without using an extra secure computation as in [Lin13] or computationally heavy trapdoor commitments as in [Bra13], furthermore, we do not strictly require any specific computational assumption. [5]

Let $t = \ell/2$, that is, t is the amount of check circuits. Also, let $k'^0_{i,j}$ and $k'^1_{i,j}$ be the zero and one keys for the i'th output wire on the j'th circuit. We use the free-xor technique [KS08] and let Alice garble each circuit with a different Δ. That is, for circuit j and all wires k Alice chooses the wire keys such that $k^0_{k,j} \oplus k^1_{k,j} = \Delta_j$.

We select the replication factor ℓ such that, except with probability 2^{-s}, if Alice was not caught cheating during the cut-and-choose phase, *at least one* of the evaluation circuits that remain is correct. Our strategy is then to make sure that for each pair of garbled circuit evaluations whose semantic output differ, Bob will be able to efficiently compute the Δ values for one of those circuits. Knowing the Δ value used to garble a circuit allows Bob to learn the input x that Alice submitted for that circuit.

Using Polynomials to Compute Δ's. We now explain how we let Bob compute Δs for those circuits whose outputs diverge.

As part of the garbling phase, Alice associates a polynomial of degree at most t, P_i with each output wire in the circuit, where $t = \ell/2$. Now for each garbled circuit $j \in [\ell]$ and each output wire $i \in [o]$ Alice associates a point on the corresponding polynomial, i.e., a value $P_i(j)$. Thus we have a polynomial associated with each output wire and for each of these polynomials we have a point associated with each garbled circuit. Next, for the 0-key on each output

---

[5] We do require access to OT, but any one-out-of-two OT can be used, including OT extensions.

wire in each circuit $\left\{ \left\{ k'^{0}_{i,j} \right\}_{i \in [o]} \right\}_{j \in [\ell]}$ and point $P_i(j)$ we associate the "link" $L_{i,j}$ (see Appendix B for details). For simplicity assume that the link is realized as follows $L_{i,j} = (r_{i,j}, s_{i,j}, h_{i,j}, g_{i,j})$ where $r_{i,j}, s_{i,j} \in_R \{0,1\}^{\kappa}$ are uniformly random sampled elements, $h_{i,j} = \mathrm{H}\left(P_i(j), r_{i,j}\right) \oplus k'^{0}_{i,j}$ and $g_{i,j} = \mathrm{H}\left(k'^{0}_{i,j}, s_{i,j}\right) \oplus P_i(j)$. Alice sends all these links to Bob.

Next notice, that after the cut-and-choose step, Bob will learn all the 0-keys on the output wires for t garbled circuits and in turn be able to learn t points on all the polynomials (by using the output 0-keys, $k'^{0}_{i,j}$ along with the corresponding links $L_{i,j}$). Thus by learning just one more 0-key for a given output wire, Bob will be able to learn one more point on one of the polynomials. This will give him a total of $t+1$ points on a polynomial of degree at most t and he will thus be able to do polynomial interpolation and learn all the points on that polynomial. Knowing all the points on a given polynomial will make it possible for Bob to learn the 0-key on a given output wire in *all* the garbled circuits (by using the points, $P_i(j)$ along with the links $L_{i,j}$).

A bit more specifically, suppose that Bob ends up with two garbled circuits where there is an output wire that outputs the 0-key in one of the circuits and the 1-key in the other. Say, w.l.o.g. that this is the case for output wire 1 in circuit 1 and circuit 2, such that Bob learns both $k'^{0}_{1,1}$ and $k'^{1}_{1,2} = k'^{0}_{1,2} \oplus \Delta_2$. Using the link $L_{1,1} = (r_{1,1}, s_{1,1}, h_{1,1}, g_{1,1})$ Bob can compute the point $P_1(1)$ as $P_1(1) = g_{1,1} \oplus \mathrm{H}\left(k'^{0}_{1,1}, s_{1,1}\right)$. This gives him a total of $t+1$ points on the polynomial P_1 (remember that he already knew t points from the check part of the cut-and-choose phase). Using these points he can then do polynomial interpolation, which in turn will make it possible for him to easily compute the point $P_1(2)$. Finally, he can then compute the following:

$$ h_{1,2} \oplus \mathrm{H}\left(P_1(2), r_{1,2}\right) \oplus k'^{1}_{1,2} = k'^{0}_{1,2} \oplus k'^{1}_{1,2} = \Delta_2 \ . $$

Using Δ_2 he can completely degarble the second garbled circuit, and in turn, also learn Alice's plain input to circuit 2. This is so, as knowing the Δ value and one key for all wires allows to compute both keys for all wires and thus completely open up a garbled circuit.

The approach of using the links to learn the Δ values generalizes to each output wire in two circuits where one is the 0-key and other the 1-key. Thus, assuming we have enough output wires with different values it is possible to find Δ for all garbled circuits, and in turn Alice's plain input to all garbled circuits.

We want to emphasize, however, that Bob does not use the Δ values to completely open up any garbled circuit, as this would be too expensive (it would essentially count as an extra garbling towards the computational complexity). Instead Bob only opens up the input layer to learn the semantic value of the input keys of Alice to the circuit. He then evaluates the universal hash function on Alice's input in plain and compare it with the semantic meaning of the garbled evaluation of the universal hash function. Then Bob discards the circuits where the plain and semantic meaning of the output of the universal hash function is discrepant, which cannot lead to selective errors, as the correctness of the augmented part of the circuit is known by Alice already. Now, for all the remaining circuits, the inputs of Alice are the same, except with a small probability 2^{-s} that a collision for the universal hash function occurred. This is so, as the circuits computed the hash function correctly, and the hash function was chosen by Bob after Alice committed to her inputs. Hence Bob can safely abort if there are different inputs left and otherwise just evaluate f in plain on his own input and the unique input of Alice.


Figure 1. Four garbled circuits with three output wires. Each 0-key on these output wires is linked to a particular point on a polynomial of degree at most 2.

Ensuring Consistency of Polynomials. We made a few assumptions in the strategy described above. That is, it only works if the polynomials are consistent, that is, if they are at most degree t *and* we have enough output wires that diverge in value to make sure Bob can learn the Δ values for each of the evaluation circuits. Since Alice might be malicious, and thus can deviate from the prescribed protocol in an arbitrary manner, we cannot be sure that she constructs the polynomials of degree at most t. In particular she might construct these to be degree $2t$ and thus make it impossible for Bob to use polynomial interpolation to find the Δ values. We could use techniques like Kate *et al.* [KZG10] to ensure consistency of the polynomials, but this would be inefficient and impose certain number theoretic

assumptions. We instead introduce an additional cut-and-choose phase in which Alice prepares more polynomials than actually needed, and then Bob randomly chooses half of the polynomials as *check* polynomials, which Alice then reveals and Bob verifies to have degree at most t. We say that a polynomial of degree greater than t is *inconsistent* and if it has degree at most t we say it is *consistent*. Thus for Bob to accept the check polynomials they must all be consistent. If all the check polynomials are consistent then a large quantity of the remaining polynomials are guaranteed to be consistent. In fact, we select half of the polynomials as check polynomials, and therefore get that a majority of the remaining circuits are guaranteed to be correct except with negligible probability in the amount of challenge polynomials, for the same reasons as in the analysis of circuit cut-and-choose in [LP07].

Ensuring consistency of the polynomials with cut-and-choose is efficient and does not rely on any specific number theoretic assumptions. However, it leaves us with the issue that some of the polynomials that are not checked might still be inconsistent and thus not usable for interpolation in our context, along with the problem that we still need enough output wires which diverge in value to make sure Bob can learn the Δ values for each of the evaluation circuits.

To cope with these problems we introduce another circuit augmentation, which is a universal hash function, taking the output of the original computation as input. We then let the the wire keys used in the links be the wire keys of the output of this new universal hash function instead of the wire keys of the actual output of the garbled circuit. Intuitively, if there is any divergence on one bit in the original output then the digests of the universal hash function will diverge in many bits. Thus the hash function ensures that we have enough output wires which will differ in their keys, since the output of the hash function is uniformly distributed. Thus, if a garbled circuit is maliciously constructed then the hashed value of the malicious output will differ in many bits compared with the hashed value of the correct output. This ensures, except with negligible probability in the amount of output bits of the augmentation, that we will have enough output wires with different keys to learn Δ for all the evaluation circuits. In the full version [FJN14] we prove that with only a majority of the polynomials guaranteed to be consistent, a hash function with at least $\lceil 4.82s + 4.82 \rceil$ output bits is sufficient in order to guarantee, except with probability 2^{-s}, that Bob will be able to learn all the Δ values of the evaluation circuits if he ends up with two or more evaluation circuits where at least one output wire diverge.

The augmentation on the output wires is done in almost the same efficient way as the augmentation to ensure input consistency of Alice. That is, we exploit the free-xor technique and let Bob send the specification of the hash function to Alice in clear. Specifically, let z be the binary output of the computation. Then we have Bob choose two random binary strings $b^{\mathsf{Out1}} \in_R \{0,1\}^{\lceil o+4.82s+3.82 \rceil}$ and $b^{\mathsf{Out2}} \in_R \{0,1\}^{\lceil 4.82s+4.82 \rceil}$. Using the first string we define a $\lceil 4.82s + 4.82 \rceil \times o$ binary matrix M^{Out} such that the j'th bit of the i'th row is the $i+j-1$'th bit of b^{Out1}. That is, $M_{i,j}^{\mathsf{Out}} = b_{i+j-1}^{\mathsf{Out1}}$. The augmentation should then compute $(M^{\mathsf{Out}} \cdot z) \oplus b^{\mathsf{Out2}}$ assuming z is a binary vector in column form.

Figure 2. Illustration comparing the original circuit (right) and the augmented circuit (left). Inputs and output are shown in light grey, whereas actual computation is shown with a darker grey.

Thus the binary strings $b^{\mathsf{Out1}} \in \{0,1\}^{\lceil o+4.82s+4.82 \rceil}$ and $b^{\mathsf{Out2}} \in \{0,1\}^{\lceil 4.82s+4.82 \rceil}$ defines a family of universal hash functions mapping o bits to $\lceil 4.82s + 4.82 \rceil$ bits. In turn a sampling of such a function is simply a random choice of the strings b^{Out1} and b^{Out2}. We will call a specific function from this family for $\mathbb{H}^{\mathsf{Out}}$.

An illustration of the final circuit with both augmentations is given in Fig. 2 and an informal version of the protocol itself is given in Section 3.

1. *Setup*
   (a) Assume the function to compute is $f(x, y) = z$ where Alice inputs $x \in \{0, 1\}^m$, Bob inputs $y \in \{0, 1\}^n$ and Bob learns the results $z \in \{0, 1\}^o$.
   (b) Bob then samples a random binary matrix $M^{\mathsf{Sec}} \in_R \{0, 1\}^{\bar{n}} \times \{0, 1\}^n$ and a random input $\bar{y}$ of $\bar{n} = \max(4n, 8s)$ bits such that $M^{\mathsf{Sec}} \cdot \bar{y} = y$. He sends M^{Sec} to Alice and they both agree on a function, f' with M^{Sec} embedded such that $f'(x, \bar{y}) = f(x, y)$. Denote the Boolean circuit computing f by C.
2. *Polynomial Setup*
   (a) Alice then randomly selects $p = 6s + 7$ polynomials, P_i for $i \in [p]$, of degree $\leq t = \ell/2$, over a finite field with 2^κ elements.
   (b) Alice computes ℓ points on each of the polynomials, achieving a total of $\ell \cdot (6s + 7)$ points. Alice commits to each point j on each polynomial i using a verifiable commitment scheme and sends these commitments to Bob.
   (c) Alice and Bob complete a cut-and-choose procedure on the random polynomials to select $\lfloor 1.18s + 2.18 \rfloor$ of these for *checking* and $\lceil 4.82s + 4.82 \rceil$ for *evaluation*. For each of the check polynomials Alice sends the openings to the points committed to in the previous step. Bob then uses the points to interpolate the polynomials and verifies that they are in fact all polynomials of degree at most t. If not, he aborts.
3. *Oblivious Transfer*
   (a) For each $j \in \ell$ Alice constructs a free-xor garbling of C, denoted by $\tilde{C}_j$, which uses a uniformly random chosen global difference, denoted by Δ_j.
   (b) Alice and Bob then engage in batch OT such that Bob learns the ℓ keys in correspondence with each of his input bits $\bar{y}$.
4. *Commitment* Alice sends the ℓ garbled circuits to Bob and makes a verifiable commitment to the 0- and 1-output keys of each of the garbled circuits. She then commits to her own choice of input keys and also makes verifiable commitments to the concatenation of the 0- and 1-keys on each input wire in each of the garbled circuits.
5. *Augmentation*
   (a) Bob randomly samples a universal hash function $\mathbb{H}^{\mathsf{In}}$ from a family of universal hash functions mapping $|x| = m$ bits to s bits and a universal hash function $\mathbb{H}^{\mathsf{Out}}$ mapping $|z| = o$ bits to $\lceil 4.82s + 4.82 \rceil$ bits from another family of hash functions as described in the previous section.
   (b) Bob sends a description of these hash functions to Alice who augments all of the garbled circuits with $\mathbb{H}^{\mathsf{In}} \oplus a$ and $\mathbb{H}^{\mathsf{Out}}$ using only xor operations of respectively her input keys and the output keys of the garbled circuits.
   (c) Alice then makes verifiable commitments to the 0- and 1-output keys of the $\mathbb{H}^{\mathsf{In}} \oplus a$ augmentation. Alice also constructs, and sends, a "link" of each of the output keys of $\mathbb{H}^{\mathsf{Out}}$ to the corresponding polynomial points. That is, for each output wire i of $\mathbb{H}^{\mathsf{Out}}$ in each garbled circuit j she sends to Bob the link $L_{i,j}$.
6. *Cut-and-Choose*
   (a) Alice and Bob uses a coin-tossing protocol to agree on a random subset $\mathcal{C}$ of size $t = \ell/2$ of the garbled circuits, called the *check circuits*. Alice then sends both the 0- and 1-key for each input wire in the set of check circuits.
   (b) Bob then checks the input keys against the verifiable commitments and verifies that the check circuits were correctly constructed. He does so by iteratively checking the construction of each gate and at the end checking that the verifiable commitments to the output keys were correctly constructed.
   (c) Bob then verifies the $\mathbb{H}^{\mathsf{In}} \oplus a$ augmentation by using the input keys to compute the output keys of the augmentation locally and then checking the output keys against the verifiable commitments. Finally, he computes the output keys of $\mathbb{H}^{\mathsf{Out}}$, using the output keys of the garbled circuit. He then uses these keys on the "links" from the augmentation phase to learn the t points on each of the $\lceil 4.82s + 4.82 \rceil$ polynomials remaining from the set-up phase. He then checks these points using the verifiable commitments. If any of the checks above fail he aborts.
7. *Evaluation* Alice opens the input keys in correspondence with her input x for all the evaluation circuits to Bob. Bob then evaluates $\mathbb{H}^{\mathsf{In}} \oplus a$. If the semantic value of any of the evaluations of $\mathbb{H}^{\mathsf{In}} \oplus a$ are discrepant, or if any of the verifiable commitments on the output keys of this augmentation is incorrectly constructed, he aborts. Otherwise he evaluates the garbled circuits. He then uses the verifiable commitments on the output wires of the garbled circuits to find the semantic meaning of the output keys he learn. If any of the semantic values on the output wires of the evaluated garbled circuits are inconsistent then he proceeds with the reconstruction phase. Otherwise he "simulates" the reconstruction phase using garbage data (to avoid timing attacks) and in the end outputs the only semantic output value he learned.

Figure 3. Protocol Overview - Part 1

1. *Reconstruction*
   (a) Bob computes the augmented output $\mathbb{H}^{\mathsf{Out}}$ for each of the evaluated garbled circuits. Since he now knows the semantic value on each garbled circuit's output wires, along with the specification of $\mathbb{H}^{\mathsf{Out}}$, he can compute, in plain, the semantic values he expects to find after evaluating $\mathbb{H}^{\mathsf{Out}}$. Any augmented circuit where this is not the case he discards. From the remaining circuits he computes the points on the remaining polynomials from the *Polynomial Setup* phase using the links from the augmentation phase. He then uses the t polynomial points from the cut-and-choose phase along with the discrepant keys of an augmented output wire in two circuits to learn $t + 1$ points on a polynomial for a given wire. He then does polynomial interpolation on these points to find all the ℓ points of this wire, i.e. one for each garbled circuit. He then uses the links to find the 0-key on that wire in all of the evaluation circuits. Hence Bob can learn Δ_j for at least one garbled circuit (because the output of that wire is discrepant and thus at least one garbled circuit will have 1-key output on that wire). He continues in this manner for each discrepant output wire until he learns Δ_j for all the non-discarded evaluation circuits.
   (b) Using these Δ's along with Alice's input keys and the verifiable commitments to the keys of the input wires Bob extracts the plain values of Alice's input in all the evaluation circuits. For each input, Bob evaluates $\mathbb{H}^{\mathsf{In}}$ on the plain input and discards the circuits where the plain output from $\mathbb{H}^{\mathsf{In}} \oplus a$ does not match the semantic output previously computed by the garbled circuit. He then uses Alice's plain input for one of the remaining circuits and computes the function f in plain to learn the correct output. If there are no remaining circuits, Bob terminates.

Figure 4. Protocol Overview - Part 2

3   The Protocol

Assume we have access to a free-xor garbling scheme, an OT functionality, a coin-tossing functionality along with a commitment scheme supporting both standard commitments (computationally hiding/binding) and commitments where one opens a commitment by revealing *exactly* the message committed to. We call the last type of commitment scheme a *verifiable commitment scheme* as it allows a party to verify if a particular message is the opening of a particular commitment without learning the opening if this is not the case.

We let Alice garble ℓ versions of a Boolean circuit C using the free-xor technique [KS08], we call these garbled circuits $\tilde{C}_j$ where $j \in [\ell]$. For each of these garbled circuits we associate a distinct global value $\Delta_j \in \{0,1\}^\kappa$, which is used to enforce the constraint that for any wire, i, in the garbled circuit $\tilde{C}_j$, we have that $k_{i,j}^1 = k_{i,j}^0 \oplus \Delta_j$ in correspondence with a free-xor garbling scheme.

At an informal level our protocol can then be summarized with the phases *Setup*, *Polynomial Setup*, *Oblivious Transfer*, *Commitment*, *Augmentation*, *Cut-and-Choose*, *Evaluation*, and *Reconstruction*. Of these, *Setup*, *Oblivious Transfer* and *Commitment*, are very similar to other protocol based on cut-and-choose of Yao garbled circuit, e.g., [LP07,LP11,sS11,FN13,SS13]. In Fig. 3 we sketch what happens in each of these steps and leave the specific details to Appendix D and the proof of security against a static, malicious, computationally bounded adversary to Appendix E.

4   Implementation

As our protocol is designed to work well in the SIMD model we did our implementation in CUDA, which is a platform that supports both explicit programming of SIMD execution, along with cheap hardware, GPUs, which supports SIMD execution. More specifically we implemented our protocol in ANSI C using the CUDA extension by NVIDIA. Some parts of our implementation is based on the code from [FN13] which gave very efficient construction and evaluation of garbled circuits (including handling of the possible selective failure attack on Bob's input), along with an "OT extension", using a consumer GPU. We direct the reader to [FN13] for details on these aspects of the implementation.

It should be noted that a recent result by Husted *et al.* [HMSG13] presents another scheme for constructing and evaluation garbled circuits using the GPU. Their scheme is up to a factor 3 faster than the scheme in [FN13] in computation on the same hardware, but increases the size of each gate with another ciphertext (both xor and non-xor gates). Unfortunately their result was not published until after we had completed our implementation and thus we do not know whether their scheme would be faster when used with our protocol. However, a study of this would be interesting for future work.

Trading Assumptions for Efficiency. Since the description of our protocol is given in the hybrid setting, assuming access to OTs, commitments, coin-tossing and garbling, we need to make some decisions on the security assumptions needed in order to realize the functionalities in our implementation. We have chosen to be liberal in the security assumptions in order to achieve a more efficient solution. In particular this means that our implementation assumes the existence of random oracles and SHA-1 is used to implement random oracle queries with a 160 bit output, thus we set $\kappa = 160$. We use the ROM garbling scheme of [PSSW09], which only requires 4 calls to the oracle to construct a garbled gate, and only a single call to the oracle in order to evaluate a garbled gate, see Appendix B for details.

We realize a coin-tossing functionality along with a commitment scheme (that fits our needs) highly efficiently in the ROM. We also base the OTs we need on an efficient OT extension, which only assume access to $O(s)$ "seed" OTs in the ROM. See Appendix C for details on these realizations. Finally, we also implement the verification of "circuit seeds" [GMS08] to eliminate the need of sending the garbled computation tables of the check circuits.[6]

SIMD Implementation Optimizations.

$\mathbb{H}^{\mathsf{In}} \oplus \mathsf{a}$ *and* $\mathbb{H}^{\mathsf{Out}}$. First notice that $\mathbb{H}^{\mathsf{Out}}$ is very similar to $\mathbb{H}^{\mathsf{In}} \oplus \mathsf{a}$ and they can both be computed using only xor operations on particular keys. Furthermore, each output key of both functions can be computed independently of the other output keys, since each output key only depends on the input keys to $\mathbb{H}^{\mathsf{In}} \oplus \mathsf{a}$, respectively $\mathbb{H}^{\mathsf{Out}}$, and the matrix M^{In}, respective M^{Out}. Thus the output of each function can clearly be done in a SIMD manner, up to the amount of outputs of $\mathbb{H}^{\mathsf{In}}$, respectively $\mathbb{H}^{\mathsf{Out}}$, i.e. s, respectively $\lceil 4.82s + 4.82 \rceil$.[7] However, the xor operation is so light that it is in general considered overkill to have a GPU thread do nothing but a few xors, because of the overhead of copying data from the RAM of the host system to the RAM of the GPU and back. Thus we simply implemented the functionalities in a sequential manner using the CPU. However, using the CPU might not always be the best case as we will discuss, based on our experiments, in Section 4.

Commitments. The commitments, both the computationally hiding/binding and verifiable types are also done in the same manner as in [FN13] i.e. by using SHA-1 to hash the string we wish to commit to (concatenated with an ID and some pseudo-randomness if it is supposed to be computationally hiding). For commitments of long strings a reduction approach is used: Concatenate some random salt to the string we need to commit to and assume it is now l bits long, then we can simply hash all independent and continues ck bits pieces in parallel, for some constant c. This results in $\lceil l/ck \rceil$ digests. We concatenate these digests together to a string and again hash all independent and continues ck sized pieces of this string in parallel. We continue in this manner until a single digest remains, which will be the commitment. The overhead, and sequential complexity, of this approach is only logarithmic in the size of the string we hash.

Polynomial Representation. In the protocol description we explained that the polynomials would have points in $\mathrm{GF}\left(2^{160}\right)$ as they need to have the same size as the wire keys. However, our initial implementation showed, which was also expected, that doing polynomial interpolation on elements in $\mathrm{GF}\left(2^{160}\right)$ was quite slow. This was expected since multiplying two elements of $\mathrm{GF}\left(2^{\kappa}\right)$ requires $O(\kappa^2)$ operations using the "trivial" approach. The multiplication can of course be optimized, in particular when one of the numbers only has a few bits set. However, in our setting at

---

[6] In the "circuit seeds" approach Alice constructs each specific gabled circuit using a small seed of entropy and a pseudo-random generator (in our case SHA-1). Then, instead of sending the garbled computation tables of all the garbled circuits she sends a hash digest of each of the garbled circuits, which is much smaller than the actual garbled computation tables). If a circuit is selected as a check circuit Alice simply sends the seed for that circuit and Bob can construct the garbled computation tables himself and verify them against the digest. For the evaluation circuits Alice sends the actual garbled computation tables, which Bob verifies against the commitments to make sure Alice did not try to cheat.

[7] The computation of $\mathbb{H}^{\mathsf{In}}$ and $\mathbb{H}^{\mathsf{Out}}$ can be parallelized even further: Since xor is commutative we can use a classic reduction approach for computing each output key. That is, have several threads working on each output key by xor'ing together a few of the input keys and storing the intermediate result. When these threads terminate new threads take over and xor together the intermediate results, and so on, until a single thread remains which then contains the output key for a single output bit of the function. This makes it possible to parallelize up to the logarithm of the input to $\mathbb{H}^{\mathsf{In}}$, respectively $\mathbb{H}^{\mathsf{Out}}$.

least one of the numbers will always be random in $GF(2^\kappa)$ and thus it does not seem possible to hope for better complexity than $O(\kappa)$. However, when κ is small enough, multiplication and inversion in $GF(2^\kappa)$ can be realized highly efficiently, with a good choice of reduction polynomial, in κ time, or in constant time through a lookup table. This is in particular true in the implementation of the S-box in AES for $\kappa = 8$. Now notice that in our setting we will need to use polynomials with random coefficients evaluated on points in the range $[2t] = [\ell] = [O(s)]$. Since s is a statistical security parameter and the garbled table are based on a computation security parameter of $\kappa = 160$ it will not make sense to have $s > 160$. Thus, in practice the range of points we will need to evaluate polynomials on will always be strictly less than 200. Now see that $2^8 > 200$. So instead of using polynomials in $GF(2^{160})$ we use 20 polynomials in $GF(2^8)$ in place of a single polynomial in $GF(2^{160})$ – i.e., we use polynomials over the ring $GF(2^8)^{20}$, which is still a secure secret sharing scheme, c.f. [CFIK03]. Thus that we can use lookup tables to do multiplication and inversion in constant time, in a SIMD manner for each of the 20 polynomials.[8]

Polynomial Construction. The $20p$ polynomials with elements from $GF(2^8)$ are generated in a SIMD manner using the GPU based on a 160 bit seed of entropy. Specifically what we do is to have $p \times (t+1)$ threads generate coefficients consisting of 160 bits of pseudo-randomness by hashing a seed along with a unique identifier. Using these coefficients we have $p \times 2t$ threads computing $2t$ points on each of the $20p$ polynomials. More specifically $P_i(j)$ for all $j \in [2t]$ and $i \in [20p]$. This is done by a loop of $t+1$ iterations where an element $c_{i,j}j^d \in GF(2^8)$ is computed in the d'th iteration and then added to the result of the previous $d-1$ iterations. The value j^d is computed by taking j^{d-1} from the previous round and multiplying it with j. This is done in order to avoid computing any exponentiations. Notice that the loop could also be parallelized in a SIMD manner by having a multiple of $t+1$ more threads, each computing a single value of $c_{i,j}j^d$. Finally, these values could be added together using a reduction approach. However, since benchmarks showed that the construction of the polynomials only accounted for a very small part of the execution time, we choose to optimize other aspects of the protocols contributing more to the overall execution time.

Polynomial Interpolation. Our polynomial interpolation (part of the *Polynomial Setup*) is based on Lagrange interpolation and consists of 4 subphases, *denominator, numerator, combination* and *reduction* in order to optimize the SIMD parallelization. Before we go through these phases remember that in Lagrange interpolation we assume knowledge of $t+1$ data point pairs $\{(x_j, y_j)\}_{j \in [t+1]}$ where all x_j are distinct. These define the interpolation polynomials in Lagrange form:

$$L(x) := \sum_{j \in [t+1]} y_j l_j(x) \,,$$

where the Lagrange basis polynomials are defined as

$$l_j(x) := \prod_{1 \le g \le t+1, g \ne j} \frac{x - x_g}{x_j - x_g} \,,$$

where $1 \le j \le t+1$. Using $l_j(\cdot)$ an arbitrary point x can then be computed. In our case we will know $t+1$ points from $[2t]$ and want to compute the remaining $t-1$ points in $[2t]$.

When we verify that the $2t$ points Bob got from Alice, for each of the $\lfloor 1.18s + 2.18 \rfloor$ check polynomials, constitutes polynomials of degree at most $t+1$, we do polynomial interpolation of the points $]t+1; 2t]$ using the points $[1; t+1]$ and verify that the newly interpolated points are the same as Alice sent us. Now the phases for which we interpolate these points in a SIMD manner goes as follows:

Denominator: We have $t+1$ threads computing the denominator of each $l_j(x)$. Notice that this is enough as the choice of value of x does not come into play here, but only in the numerator. Finally, each thread computes the inverse so that each denominator can be directly multiplied onto a numerator later.

---

[8] In the actual implementation we do not use a look up table for multiplication as it turned out to be around two times slower than simply doing 8 iterations of bit shifts and if statements. Furthermore, as an artifact of the implementation framework we do not do 20 polynomials in a SIMD manner, but instead sequentially.

Numerator: We use $(t+1) \times (t-1)$ threads to compute the numerator of each $l_j(i)$. That is $l_j(i)$ will be computed by thread $j \cdot (t-1) + i$ so that we compute the numerator of each of the $t+1$ Lagrange polynomials using a different value for $i \in]t+1; 2t]$. Finally, each thread multiply its result with the appropriate denominator from the previous step. That is, computing $l_j(i)$ for each $i \in]t+1; 2t]$.

Combination: We use $\lfloor 1.18s + 2.18 \rfloor \times (t+1) \times (t-1)$ threads to compute each term of the $t+1$ terms of $L(i)$ for the $20\lfloor 1.18s + 2.18 \rfloor$ polynomials for each $i \in]t+1; 2t]$. This is simply done by having each thread multiply the appropriate values y_j for each for the $20\lfloor 1.18s + 2.18 \rfloor$ polynomials with the corresponding Lagrange basis polynomial (computed in the "Numerator" step).

Reduction: To continue with the maximal level of SIMD parallelization up to $\lfloor 1.18s + 2.18 \rfloor \times (t+1) \times (t-1)$ threads will complete a reduction approach, i.e. any given thread will xor together two terms of $L(i)$. A new thread will take over two of the intermediate results and xor these together. This will continue until $\lfloor 1.18s + 2.18 \rfloor \times (t-1)$ threads remain, containing the result $L(i)$ for all $i \in]t+1; 2t]$ and all $20\lfloor 1.18s + 2.18 \rfloor$ polynomials

All the steps are done with all generated data staying in the GPU's RAM and only at the end the final results are copied back to the host's RAM. Because all the data is on the GPU it is sensible to do the reduction step instead of a sequential computation on the host system, even though the operations done will only be xors.

Polynomial Reconstruction. The polynomial reconstruction (part of the *Reconstruction* phase) is done almost in the same manner as the polynomial interpolation. The main difference being that in the reconstruction the set of $t+1$ points we know might be different for each of the $\lceil 4.82s + 4.82 \rceil$ sets of 20 polynomials. However, at most one of these points might be different since Bob will always learn the same t points for all the polynomials as an effect of the cut-and-choose of garbled circuits. Thus we can implement the *numerator* step in the following two substeps:

- First we use $t \times (t-1)$ threads to compute $\prod_{1 \le g \le t, g \ne j} (x - x_g)$ for each point $j \in [t]$ where x_j is an index of a check circuit and for each of the $t-1$ points, x being an index of an evaluation circuit.
- We then use $\lceil 4.82s + 4.82 \rceil \times (t+1) \times (t-1)$ threads to multiply the last point, which might vary from polynomial to polynomial, onto the result of the previous step. This will give us the numerator in all the $t+1$ Lagrange basis polynomials for each of the $t-1$ points Bob needs to learn, i.e. the $t-1$ indices of the evaluation circuits.

The same idea applies to the computation of the *denominator* step, but with a factor $t-1$ less, as the denominator remains the same for all possible values of x we need to learn.

With the above approach we avoid having $\lceil 4.82s + 4.82 \rceil \times (t+1) \times (t-1)$ threads doing a multiplication loop of $t+1$ iterations and instead only use $t \times (t-1)$ threads with a loop of t iterations and $\lceil 4.82s + 4.82 \rceil \times (t+1) \times (t-1)$ threads with a constant amount of multiplications to achieve the same result.

A Note on Multi-Threading. In order to limit the time each party is idle we employ multi-threading to the protocol in a non-SIMD manner. That is, when several distinct steps of the protocol can be carried out independently of each other we fork the host process to carry out these computations in parallel. In particular we fork when some data needs to be sent/received while other operations can be carried out with the data already known.

A Note on Parallel Complexity. The parallel complexity of cut-and-choose of garbled circuits is bounded by the depth of the circuit to compute times the garbling of a single gate along with the parallel complexity of the handling of selective failure attacks, consistency of inputs, OTs and commitments. In our case, as well as in [FN13], the complexity of handling selective failure attacks is bounded by $O(\log(\max(n, s)))$ as it only consists of xor'ing (because of the free-xor approach) at most $\max(4n, 8s)$ keys of κ bits with each other, $O(sn)$ times in parallel, using a reduction approach. The same argumentation goes for ensuring consistency of Alice's input, i.e. computing $\mathbb{H}^{\mathsf{In}} \oplus \mathsf{a}$: Using a reduction approach it is basically $O(\log(m))$ SIMD xor operations using $O(s^2 \kappa m)$ SIMD processors. Regarding handling input recovery in case of cheating, first notice that the parallel computation of $\mathbb{H}^{\mathsf{Out}}$ is almost the same as for $\mathbb{H}^{\mathsf{In}} \oplus \mathsf{a}$, thus it requires $O(\log(o))$ SIMD xor operations using $O(s^2 o \kappa)$ SIMD processors.

Considering the polynomial generation, notice that we first generate coefficients and then evaluate points based on the coefficients. Now, evaluating points can be done using a reduction approach, that is, each term of a polynomial is

evaluated independently, then two terms are added together by one thread, these results are then added together two at a time, and so on until all the terms for a given point in a given polynomial has been added together. This implies that the complexity of polynomial generation is bounded by the logarithm of the degree of the polynomial, that is $O(\log(s))$ SIMD operations using $O(s^3\kappa)$ SIMD processors. Regarding polynomial interpolation, assuming we use Lagrange interpolation, the straight forward approach would be to make an evaluation of a Lagrange basis polynomial for each point we wish to interpolate. These basis polynomials will consist of $1 + s/2$ factors. Again we can compute each factor independently and combine the factors using a reduction approach. These basis polynomials will be the same for each of the polynomials we wish to interpolate, with the exception of at most one factor (depending on which augmented output wires contains 0-values for each evaluation circuits). These basis polynomials are then used to find the actual points by multiplying each of them with a known point (learned from the links) and then adding the terms together. There will be $1 + s/2$ terms and again these can be added together using a reduction approach. Thus, the interpolation complexity is also limited to $O(\log(s))$ SIMD operations using $O(s^3\kappa)$ SIMD processors.

Notice that in the above analysis we assume that addition and multiplication is constant time, which makes sense since they can be implemented in a loop with 8 iterations or through a lookup table following the implementation idea in Section 4.

s		IO	Comp.	Comm.	Idle	Party total	Total	Total idle
9	A.	4.073 ± 0.036	53.19 ± 0.47	45.77 ± 0.17	81.70 ± 3.5	184.7 ± 3.6	211.0 ± 3.6	160.1 ± 7.0
	B.	4.040 ± 0.0012	82.81 ± 0.60	45.77 ± 0.17	78.38 ± 3.5	211.0 ± 3.6		
19	A.	4.053 ± 0.020	65.81 ± 0.89	71.97 ± 0.20	86.75 ± 4.5	228.6 ± 4.3	260.1 ± 4.3	171.3 ± 8.7
	B.	4.039 ± 0.00097	99.49 ± 0.71	71.97 ± 0.20	84.56 ± 4.3	260.1 ± 4.3		
30	A.	4.046 ± 0.0080	119.9 ± 0.90	108.1 ± 0.26	112.3 ± 3.8	344.4 ± 3.9	398.1 ± 4.1	231.0 ± 7.6
	B.	4.066 ± 0.031	167.3 ± 0.70	108.1 ± 0.26	118.7 ± 3.8	398.1 ± 4.1		
40	A.	4.055 ± 0.017	137.3 ± 0.79	132.3 ± 0.30	118.5 ± 4.1	392.1 ± 4.2	455.7 ± 4.2	244.7 ± 8.1
	B.	4.049 ± 0.016	193.1 ± 1.0	132.3 ± 0.30	126.2 ± 4.1	455.6 ± 4.2		
60	A.	4.043 ± 0.0069	170.1 ± 1.9	178.1 ± 0.38	130.3 ± 4.5	482.6 ± 4.3	583.0 ± 4.3	263.7 ± 8.5
	B.	4.093 ± 0.088	266.2 ± 0.77	178.1 ± 0.38	133.4 ± 4.3	581.8 ± 4.4		
80	A.	4.055 ± 0.017	247.1 ± 1.6	231.2 ± 0.37	168.1 ± 4.1	650.5 ± 4.2	810.4 ± 3.5	340.2 ± 6.9
	B.	4.069 ± 0.038	398.8 ± 1.1	231.2 ± 0.37	172.0 ± 3.0	806.1 ± 3.5		
119	A.	4.055 ± 0.020	377.9 ± 2.9	343.8 ± 0.37	215.1 ± 6.9	940.8 ± 4.2	1220 ± 5.2	431.0 ± 12
	B.	4.040 ± 0.0026	641.5 ± 1.1	343.8 ± 0.37	215.9 ± 5.3	1205 ± 5.2		

Table 1. Timing in milliseconds of our protocol computing oblivious AES-128 for both Alice (A.) and Bob (B.) under different statistical security parameters (s). Column "Comp." represents wall-clock time where at least one thread does computation. Columns "IO", "Comm." and "Idle " represents wall-clock time where the protocol execution is single-threaded and does disk loading, network communication or is completely idle, respectively. Column "Party total" represents the total wall-clock time of protocol execution for each party. Column "Total" represents the total wall-clock execution time of the entire protocol. Column "Total idle" represents the total wall-clock time when one of the parties is idle.

Experimental Results. All of our experiments are based on the same, commonly used, circuit for oblivious 128 bit AES encryption.[9] This circuit is used as benchmark in [HEKM11, LP11, NNOB12, HKS⁺10, FN13], and many more implementations of 2PC for functions expressed as a Boolean circuit. We ran experiments on the circuit with several different statistical security parameters on two consumer grade desktop computers connected to Aarhus University's gigabit local area network. At the time of purchase (2012) each of these machines had a price of less than $1600. Both

---

[9] We thank Benny Pinkas, Thomas Schneider, Nigel P. Smart and Stephen C. Williams for supplying the circuit.

machines have similar specifications: an Intel Ivy Bridge i7 3.5 GHz quad-core processor, 8 GB DDR3 RAM, an Intel series-520 180 GB SSD drive, a MSI Z77 motherboard with gigabit LAN and a MSI GPU with an NVIDIA GTX 670 chip and 2 GB GDDR5 RAM. The machines ran up-to-date versions of Linux Mint 14 and CUDA 5.5. Each of the experiments was repeated 50 times and with no front end applications running on either of the machines. The timings are summarized in Table 1 and Table 2. These timings include loading circuit description and randomness along with communication between the host and device and communication between the parties.[10] However, in the same manner as done in [NNOB12] and [FN13] the timings of the seed OTs have not been included as it is a computation that is needed once between two parties and thus will get amortized out in a practical context. The time it takes to initialize the GPU device (driver related overhead) has not counted either, and generally would constitute between 50 and 60 milliseconds on our test systems when the GPU is set to "persistence mode".

The timings are in milliseconds and represent "wall-clock" times. However, since some aspects of the execution are multi-threaded we have chosen to count those parts as computation time, even though one thread might not be doing computation, but rather be idle or doing communication. Thus what is counted in the communication, respectively idle columns, is the time where the party is only doing communication, respectively is completely idle.

Table 2. Timing in miliseconds of the parts of our protocol responsible for cheating recovery. The timings are taken from the computation of oblivious AES-128 for both Alice (A.) and Bob (B.) under different statistical security parameters (s). Column "Comp." represents wall-clock time where at least one thread of a given party does computation. Columns "IO", "Comm." and "Idle " represents wall-clock time where the protocol execution is single-threaded and does disk loading, network communication or is completely idle, respectively. Column "Party total" represents the total wall-clock time of protocol execution.

s		Comp.	Comm.	Idle	Party total
9	A.	4.836 ± 0.17	1.592 ± 0.046	< 1	6.467 ± 0.17
	B.	14.04 ± 0.43	1.592 ± 0.046	4.808 ± 2.4	20.44 ± 2.5
19	A.	7.373 ± 0.93	3.885 ± 0.066	< 1	11.28 ± 0.94
	B.	20.10 ± 0.52	3.885 ± 0.066	10.49 ± 3.9	34.47 ± 3.9
30	A.	33.24 ± 1.4	8.244 ± 0.079	< 1	41.86 ± 1.4
	B.	34.56 ± 0.78	8.244 ± 0.079	12.53 ± 2.5	55.34 ± 2.7
40	A.	42.60 ± 0.46	12.64 ± 0.079	< 1	55.34 ± 0.44
	B.	51.63 ± 1.1	12.64 ± 0.079	18.40 ± 2.9	82.67 ± 3.1
60	A.	47.35 ± 0.62	24.38 ± 0.11	< 1	71.94 ± 0.58
	B.	116.1 ± 0.83	24.38 ± 0.11	23.67 ± 2.5	164.2 ± 2.7
80	A.	66.48 ± 0.40	40.91 ± 0.089	< 1	107.4 ± 0.37
	B.	185.4 ± 2.7	40.91 ± 0.089	39.13 ± 0.28	265.4 ± 2.7
119	A.	102.4 ± 0.59	84.71 ± 0.11	< 1	187.6 ± 0.51
	B.	315.1 ± 1.8	84.71 ± 0.11	71.28 ± 0.22	471.1 ± 1.7

Data Analysis. From Table 1 we see that idle time takes up a significant portion of the total execution time of the protocol, i.e. between 18% and 44% for each party and up to 76% of the wall-clock time one party sits idle. Taking Table 2 into account, we see that at most 30% of the overhead of recovery is idle time. The reason is probably that most messages that needs to be sent, as part of our approach to input recovery, can be batched together with messages that needs to be sent as part of the generic structure of cut-and-choose of garbled circuits.

It should be noted that we have spent quite a bit of time trying to limit the idle time of the implementation by using multi-threading, batching messages and restructuring steps within a given party's execution. Unfortunately, it remains unknown how much this has limited idle time compared to other implementations of cut-and-choose of garbled circuits, as other authors with comparable protocols have not included measurements of idle time.

Comparing the total times of Table 1 and Table 2 we see that input recovery constitutes from 10% up to 39% for of the execution time for Bob, whereas for Alice it goes from 3.5% up to 20%. Looking further into the different steps of input recovery it turns out that around half of the computation time Bob spends on *Reconstruction* is actually spent doing the "free" computation of the output keys of $\mathbb{H}^{\mathsf{Out}}$ (since it is all xor operations on keys and we use free-xor).

---

[10] In the tests the entropy used was sampled from `/dev/urandom`, which is a non-blocking source of pseudo-randomness. This source was used in order to avoid high variance in the execution time caused by waiting for "true" randomness from `/dev/random`, which is a blocking source of randomness.

He spend even more time on "free" computations before we optimized the *Reconstruction* part of our implementation to limit the amount of cache-misses.

	Security	s	Model	Rounds	Time (s)	Equipment
[HEKM11]	Semi honest	-	ROM	$O(1)$	0.20	Desktop
This work	Malicious	2^{-9}	ROM	$O(1)$	0.21	Desktop w. GPU
This work	Malicious	2^{-40}	ROM	$O(1)$	0.46	Desktop w. GPU
[FN13]	Malicious	2^{-39}	ROM	$O(1)$	1.1	Desktop w. GPU
This work	Malicious	2^{-60}	ROM	$O(1)$	0.58	Desktop w. GPU
This work	Malicious	2^{-80}	ROM	$O(1)$	0.81	Desktop w. GPU
[NNOB12]	Malicious	2^{-58}	ROM	$O(d)$	1.6	Desktop
[KSS12]	Malicious	2^{-80}	SM	$O(1)$	1.4	Cluster, 512 nodes
[FN13]	Malicious	2^{-79}	ROM	$O(1)$	2.6	Desktop w. GPU
[SS13]	Malicious	2^{-80}	SM	$O(1)$	40.6	Cluster, 8 nodes

Table 3. Timing comparison of secure two party computation protocols evaluating oblivious 128-bit AES. d is the depth of the circuit to be computed. Notice that our implementation is run on the same hardware and using the same garbling scheme as [FN13].

5   Conclusion and Future Work

We have presented a new approach to the recent idea of "forge-and-loose" which does not need any number theoretic assumptions, avoids an auxiliary secure computation and parallelizes well in a SIMD manner. We have implemented a random oracle realization of this protocol in CUDA to take advantage of its SIMD friendly structure. The implementation shows that our protocol is up to a factor 3 faster than the most comparable implementation, showing the greatest improvements for the larger the statistical security parameters. The benchmark shows that this is not because of the extra computations associate with our forge-and-loose approach, but rather because of the asymmetry of the overall cut-and-choose of garbled circuits approach, which takes up a significantly larger percentage of the total execution time for the smaller statistical security parameters.

As future work it would be interesting to look into ways of making the overhead, needed for recovering Alice's input, smaller. In particular for larger statistical security parameters. Furthermore, making implementations (based on the same garbling, OT, etc.) of the other protocols using the forge-and-loose approach would be interesting, in particular to see which forge-and-loose technique introduces the least amount of overhead in different contexts. Finally, as we see that a significant portion of the wall-clock execution time of the protocol, even when using multi-threading, is idle time it would be interesting to look into ways of making the general structure of cut-and-choose of garbled circuits more symmetric.[11]

References

[App13]   Benny Applebaum. Bootstrapping obfuscators via fast pseudorandom functions. *IACR Cryptology ePrint Archive*, 2013:699, 2013.

[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 Roger Dingledine and Philippe Golle, editors, *Financial Cryptography*, volume 5628 of *Lecture Notes in Computer Science*, pages 325–343. Springer, 2009.

---

[11] More symmetric cut-and-choose protocols of garbled circuits has already been done in [MR13,HKE13]. However, their protocols require both parties to execute the protocol both as circuit constructor *and* as circuit evaluator. Thus, even though their protocols will probably have less idle time, it does not directly seem to be helpful in making the wall-clock execution time smaller.

[BDOZ11] Rikke Bendlin, Ivan Damgård, Claudio Orlandi, and Sarah Zakarias. Semi-homomorphic encryption and multiparty computation. In Kenneth G. Paterson, editor, *EUROCRYPT*, volume 6632 of *Lecture Notes in Computer Science*, pages 169–188. Springer, 2011.

[Blu82] Manuel Blum. Coin flipping by telephone - a protocol for solving impossible problems. In *COMPCON*, pages 133–137. IEEE Computer Society, 1982.

[Bra13] Luís T. A. N. Brandão. Secure two-party computation with reusable bit-commitments, via a cut-and-choose with forge-and-lose technique - (extended abstract). In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT (2)*, volume 8270 of *Lecture Notes in Computer Science*, pages 441–463. Springer, 2013.

[Can00] Ran Canetti. Security and composition of multiparty cryptographic protocols. *J. Cryptology*, 13(1):143–202, 2000.

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

[CFIK03] Ronald Cramer, Serge Fehr, Yuval Ishai, and Eyal Kushilevitz. Efficient multi-party computation over rings. In Eli Biham, editor, *EUROCRYPT*, volume 2656 of *Lecture Notes in Computer Science*, pages 596–613. Springer, 2003.

[CKKZ12] Seung Geol Choi, Jonathan Katz, Ranjit Kumaresan, and Hong-Sheng Zhou. On the security of the "free-XOR" technique. In *TCC*, pages 39–53, 2012.

[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 Jason Crampton, Sushil Jajodia, and Keith Mayes, editors, *ESORICS*, volume 8134 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2013.

[DPSZ12] Ivan Damgård, Valerio Pastro, Nigel P. Smart, and Sarah Zakarias. Multiparty computation from somewhat homomorphic encryption. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO*, volume 7417 of *Lecture Notes in Computer Science*, pages 643–662. Springer, 2012.

[DZ13] Ivan Damgård and Sarah Zakarias. Constant-overhead secure computation of boolean circuits using preprocessing. In *TCC*, pages 621–641, 2013.

[FJN⁺13] Tore Kasper Frederiksen, Thomas Pelle Jakobsen, Jesper Buus Nielsen, Peter Sebastian Nordholt, and Claudio Orlandi. Minilego: Efficient secure two-party computation from general assumptions. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT*, volume 7881 of *Lecture Notes in Computer Science*, pages 537–556. Springer, 2013.

[FJN14] Tore Kasper Frederiksen, Thomas P. Jakobsen, and Jesper Buus Nielsen. Faster maliciously secure two-party computation using the GPU. *IACR Cryptology ePrint Archive*, 2014:270, 2014.

[FN13] Tore Kasper Frederiksen and Jesper Buus Nielsen. Fast and maliciously secure two-party computation using the GPU. *IACR Cryptology ePrint Archive*, 2013:46, 2013.

[GMS08] Vipul Goyal, Payman Mohassel, and Adam Smith. Efficient two party and multi party computation against covert adversaries. In Nigel P. Smart, editor, *EUROCRYPT*, volume 4965 of *Lecture Notes in Computer Science*, pages 289–306. Springer, 2008.

[Gol04] Oded Goldreich. *The Foundations of Cryptography - Volume 2, Basic Applications*. Cambridge University Press, 2004.

[HEKM11] Yan Huang, David Evans, Jonathan Katz, and Lior Malka. Faster secure two-party computation using garbled circuits. In *USENIX Security Symposium*. USENIX Association, 2011.

[HKE12] Yan Huang, Jonathan Katz, and David Evans. Quid-pro-quo-tocols: Strengthening semi-honest protocols with dual execution. In *IEEE Symposium on Security and Privacy*, pages 272–284. IEEE Computer Society, 2012.

[HKE13] Yan Huang, Jonathan Katz, and Dave Evans. Efficient secure two-party computation using symmetric cut-and-choose. *IACR Cryptology ePrint Archive*, 2013:81, 2013.

[HKS⁺10] Wilko Henecka, Stefan Kögl, Ahmad-Reza Sadeghi, Thomas Schneider, and Immo Wehrenberg. TASTY: Tool for automating secure two-party computations. In Ehab Al-Shaer, Angelos D. Keromytis, and Vitaly Shmatikov, editors, *ACM Conference on Computer and Communications Security*, pages 451–462. ACM, 2010.

[HL10] C. Hazay and Y. Lindell. *Efficient Secure Two-party Protocols: Techniques and Constructions*. Information security and cryptography. Springer Berlin Heidelberg, 2010.

[HMQ04] Dennis Hofheinz and Jörn Müller-Quade. Universally composable commitments using random oracles. In Moni Naor, editor, *TCC*, volume 2951 of *Lecture Notes in Computer Science*, pages 58–76. Springer, 2004.

[HMSG13] Nathaniel Husted, Steven Myers, Abhi Shelat, and Paul Grubbs. GPU and CPU parallelization of honest-but-curious secure two-party computation. In Charles N. Payne Jr., editor, *ACSAC*, pages 169–178. ACM, 2013.

[IKOS07] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge from secure multiparty computation. In David S. Johnson and Uriel Feige, editors, *STOC*, pages 21–30. ACM, 2007.

[IPS08] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. Founding cryptography on oblivious transfer - efficiently. In Wagner [Wag08], pages 572–591.

[KS06] Mehmet S. Kiraz and Berry Schoenmakers. A protocol issue for the malicious case of Yao's garbled circuit construction. In *In Proceedings of 27th Symposium on Information Theory in the Benelux*, pages 283–290, 2006.

18

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

[KSS12]   Benjamin Kreuter, Abhi Shelat, and Chih-Hao Shen. Towards billion-gate secure computation with malicious adversaries. *IACR Cryptology ePrint Archive*, 2012:179, 2012.

[KZG10]   Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In Masayuki Abe, editor, *ASIACRYPT*, volume 6477 of *Lecture Notes in Computer Science*, pages 177–194. Springer, 2010.

[Lin13]   Yehuda Lindell. Fast cut-and-choose based protocols for malicious and covert adversaries. *IACR Cryptology ePrint Archive*, 2013:79, 2013.

[LP07]   Yehuda Lindell and Benny Pinkas. An efficient protocol for secure two-party computation in the presence of malicious adversaries. In Moni Naor, editor, *EUROCRYPT*, volume 4515 of *Lecture Notes in Computer Science*, pages 52–78. Springer, 2007.

[LP09]   Yehuda Lindell and Benny Pinkas. A proof of security of Yao's protocol for two-party computation. *J. Cryptology*, 22(2):161–188, 2009.

[LP11]   Yehuda Lindell and Benny Pinkas. Secure two-party computation via cut-and-choose oblivious transfer. In Yuval Ishai, editor, *TCC*, volume 6597 of *Lecture Notes in Computer Science*, pages 329–346. Springer, 2011.

[MF06]   Payman Mohassel and Matthew K. Franklin. Efficiency tradeoffs for malicious two-party computation. In Moti Yung, Yevgeniy Dodis, Aggelos Kiayias, and Tal Malkin, editors, *Public Key Cryptography*, volume 3958 of *Lecture Notes in Computer Science*, pages 458–473. Springer, 2006.

[MNT90]   Yishay Mansour, Noam Nisan, and Prasoon Tiwari. The computational complexity of universal hashing. In *Structure in Complexity Theory Conference*, page 90, 1990.

[MR13]   Payman Mohassel and Ben Riva. Garbled circuits checking garbled circuits: More efficient and secure two-party computation. In Ran Canetti and Juan A. Garay, editors, *CRYPTO (2)*, volume 8043 of *Lecture Notes in Computer Science*, pages 36–53. Springer, 2013.

[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.

[NO09]   Jesper Buus Nielsen and Claudio Orlandi. LEGO for two-party secure computation. In Omer Reingold, editor, *TCC*, volume 5444 of *Lecture Notes in Computer Science*, pages 368–386. Springer, 2009.

[NPS99]   Moni Naor, Benny Pinkas, and Reuban Sumner. Privacy preserving auctions and mechanism design. In *ACM Conference on Electronic Commerce*, pages 129–139, 1999.

[PSSW09]   Benny Pinkas, Thomas Schneider, Nigel P. Smart, and Stephen C. Williams. Secure two-party computation is practical. In Mitsuru Matsui, editor, *ASIACRYPT*, volume 5912 of *Lecture Notes in Computer Science*, pages 250–267. Springer, 2009.

[PVW08]   Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A framework for efficient and composable oblivious transfer. In Wagner [Wag08], pages 554–571.

[Sha79]   Adi Shamir. How to share a secret. *Commun. ACM*, 22(11):612–613, 1979.

[sS11]   shelat abhi and Chih-Hao Shen. Two-output secure computation with malicious adversaries. In *EUROCRYPT*, pages 386–405, 2011.

[SS13]   Abhi Shelat and Chih-Hao Shen. Fast two-party secure computation with minimal assumptions. *IACR Cryptology ePrint Archive*, 2013:196, 2013.

[Wag08]   David Wagner, editor. *Advances in Cryptology - CRYPTO 2008, 28th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2008. Proceedings*, volume 5157 of *Lecture Notes in Computer Science*. Springer, 2008.

[Yao82]   Andrew Chi-Chih Yao. Protocols for secure computations (extended abstract). In *FOCS*, pages 160–164. IEEE Computer Society, 1982.

[Yao86]   Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *FOCS*, pages 162–167. IEEE Computer Society, 1986.

A Overview of Variables and Parameters

A list of variable names and their meaning is given in Table 4 and Table 5.

Symbol	Meaning
s	Statistical security parameter.
κ	Computational security parameter.
C	The plain description of the Boolean circuit to compute.
f	The function computed by C.
x	A bit string representing the constructor's (Alice's) input to the circuit.
y	A bit string representing the evaluator's (Bob's) input to the circuit.
z	The circuit output destined for the evaluator.

Table 4. Overview of the free parameters of the protocol along with their meaning.

Symbol	Meaning
ℓ	Circuit replication factor; defined as the smallest even integer satisfying $\ell - \frac{1}{2}\log(\ell) + \log\left(2\frac{\sqrt{2\pi}}{e^2}\right) \geq s$.
t	The amount of check circuits; $t = \ell/2$.
p	The number of polynomials; defined as the smallest integer satisfying $p \geq 6s + 7$.
$|C|$	Amount of gates in C.
m	Amount of input bits to the circuit from the constructor, $m = |x|$.
n	Amount of input bits to the circuit from the evaluator, $n = |y|$.
o	Total amount of output bits from the non-augmented circuit, $o = |z|$.
$\bar{m}$	The length of the constructor's extended input; $\bar{m} = n + s$.
$\bar{n}$	The length of the evaluator's extended input; $\bar{n} = \max\{4n, 8s\}$.
ψ	The check set for cut-and-choose on polynomials.
π, ϕ	The check and evaluation sets, respectively, for cut-and-choose on circuits.
Δ_j	The global key for the j'th garbled circuit.
$P_i(j)$	The j'th values of the i'the polynomial in the field $\mathbb{F}_{2^\kappa}$.

Table 5. Overview of "internal" variables and parameters derived from the free parameters.

B Preliminaries

Garbled Circuits. Our protocol requires a free-xor garbling scheme. We now describe a garbling scheme that combines the state of the art optimizations for Yao garbled gates i.e., free-xor [KS08], permutation bits [NPS99] and garbled row-reduction [NPS99]. The scheme is secure in the ROM and is due to [PSSW09].

In the RO scheme of [PSSW09] one can garble a gate by doing 4 evaluations of SHA-1, and a garbled gate consists of only 3 ciphertexts (therefore saving on communication complexity). The evaluation of the gate can be done using a single SHA-1 evaluation.

When considering keys in the following we do so only for a single garbled circuit and thus eliminate the circuit index, i.e., the 0-key on wire i is called k_i^0 and the 1-key k_i^1. When it is unknown if a key represents a 0 or 1 bit we discard the superscript, i.e., we simply call the key k_i. We will sometimes also abuse notation and call the index of the respectively left, right and output key of a given gate for k_L, k_R and k_O respectively, adding the superscript 0 or 1 as appropriate.

- We have a (possibly randomized) algorithm $\mathsf{Yao}(g, k_L^0, k_R^0, \Delta, id)$ with g a description of a Boolean function, a left input *0-key* $k_L^0 \in \{0,1\}^\kappa$, a right input 0-key $k_R^0 \in \{0,1\}^\kappa$, a global difference $\Delta \in \{0,1\}^\kappa$ and a unique gate identifier id, outputs a garbled gate $\tilde{g}$ and a output 0-key $k_O^0 \in \{0,1\}^\kappa$.
- We have an algorithm $\mathsf{Eval}(\tilde{g}, k_L', k_R')$ that on input a garbled gate $\tilde{g}$, a left key $k_L' \in \{0,1\}^\kappa$ and a right key $k_R' \in \{0,1\}^\kappa$ outputs an output key $k_O' \in \{0,1\}^\kappa \cup \{\bot\}$.
- We define the *1-keys* k_L^1, k_R^1, k_O^1 s.t. $k_L^0 \oplus k_L^1 = k_R^0 \oplus k_R^1 = k_O^0 \oplus k_O^1 = \Delta$.

The idea is that a garbled gate $\tilde{g}$ has a 0- and a 1-key associated with each of its wires (left input, right input and output wire), and that these keys represent the bit values on those wires. E.g., if $\tilde{g}$ is a garbled AND gate generated as $(\tilde{g}, k_O^0) \leftarrow \mathsf{Yao}(\wedge, k_L^0, k_R^0, \Delta, id)$ then $\mathsf{Eval}(\tilde{g}, k_L^a, k_R^b)$ for any $a, b \in \{0,1\}$ should output $k_O^{a \wedge b}$.

Note that if Alice samples Δ and a 0-key, say k_L^0, at random and gives the key k_L^a to Bob then there is no way for Bob to infer the bit a from k_L^a. Furthermore, even if Bob learns a he cannot guess the key k_L^{1-a} with better probability than guessing Δ. For a garbling scheme to be secure we intuitively want that even if Bob learns $\tilde{g}$ and keys k_L^a and k_R^b for $a, b \in \{0,1\}$, and thus is able to evaluate $k_O^{g(a,b)} \leftarrow \mathsf{Eval}(\tilde{g}, k_L^a, k_R^b)$, then he cannot learn anything about k_L^{1-a}, k_R^{1-b} or $k_O^{1-g(a,b)}$, even if he knows a and/or b.

Furthermore, Bob cannot evaluate the gate on any other inputs. Thus if Bob sends back $k_O^{g(a,b)}$ to Alice, Alice can learn $g(a, b)$ (as she knows k_O^0 and Δ) and be confident that this is the correct result.

Definition 1. *We say that* $(\mathsf{Yao}, \mathsf{Eval})$ *is a* Yao free-xor gate garbling scheme *if the following holds:*
Correctness: *Let* $(\tilde{g}, k_O^0) \leftarrow \mathsf{Yao}(g, k_L^0, k_R^0, \Delta, id)$ *where* k_L^0, k_R^0 *and* Δ *are random bit strings, id is a uniques gate identifier and g is a Boolean function such that* $g : \{0,1\} \times \{0,1\} \to \{0,1\}$ *then for all* $b_L, b_R \in \{0,1\}$ *we demand that* $Pr\left[\mathsf{Eval}(\tilde{g}, k_L^{b_L}, k_R^{b_R}) = k_O^{g(b_L, b_R)}\right] = 1$.

From a garbled gate scheme we can construct a circuit garbling scheme. Let C be a Boolean circuit with m input wires, n output wires and c gates. First sample a uniformly random $\Delta \in \{0,1\}^\kappa$ and a random key for each input wire, denoted by $k_{I,1}^0, \ldots, k_{I,m}^0 \in \{0,1\}^\kappa$. Then inductively, for all gates which has both input keys defined, proceed as follows: For an XOR-gate with input wires L and R and output wire O, let $k_O^0 = k_L^0 \oplus k_R^0$. To later evaluate the gate on k_L^a and k_R^b, compute $k_O^{a \oplus b} = k_L^a \oplus k_R^b$. For a non-XOR-gate with identifier id, input wires L, R, output wire O and truth table g, sample $(\tilde{g}_{id}, k_O^0) \leftarrow \mathsf{Yao}(g, k_L^0, k_R^0, \Delta, id)$. This will eventually associate a key to each wire. Denote the output keys as $k_{O,1}^0, \ldots, k_{O,n}^0$. Let ID be the set of identifiers of gates. The output of the garbling is then Δ, the input and output keys along with the garbled gates, i.e. $\left(\Delta, \{k_{I,i}^0\}_{i \in [m]}, \{k_{O,i}^0\}_{i \in [n]}, \{\tilde{g}_{id}\}_{id \in \mathrm{ID}}\right)$. We use $\tilde{C}$ to denote $\left(\{\tilde{g}_{id}, \}_{id \in \mathrm{ID}}, C\right)$, and we call $\tilde{C}$ the garbled circuit. We use $\left(\Delta, \{k_{I,i}^0\}_{i \in [m]}, \{k_{O,i}^0\}_{i \in [n]}, \tilde{C}\right) \leftarrow \mathsf{CYao}(C)$ to denote a random circuit garbling. Let $x \in \{0,1\}^m$, then to evaluate $\tilde{C}$ on $\{k_{I,i}^{x_i}\}_{i \in [m]}$, inductively compute $k_w^{x_w}$ for each wire w in the circuit, where x_w is the value that wire w would have in $C(x)$. This is possible because of the correctness of the garbled gates and the fact that the XOR evaluation described above is correct too. Let $y = C(x)$. We denote the output of an evaluation of a garbled circuit by $\{k_{O,i}^{y_i}\}_{i \in [n]} \leftarrow \mathsf{CEval}\left(\tilde{C}, \{k_{I,i}^{x_i}\}_{i \in [m]}\right)$. The evaluation is deterministic if the evaluation of each garbled gate is deterministic, which we will assume for simplicity.

Definition 2. *We say that* $(\mathsf{CYao}, \mathsf{CEval})$ *is a* Yao free-xor circuit garbling scheme *if the following holds:*
Correctness: *Let* $\left(\Delta, \{k_{I,i}^0\}_{i \in [m]}, \{k_{O,i}^0\}_{i \in [n]}, \tilde{C}\right) \leftarrow \mathsf{CYao}(C)$. *Then for all* $x \in \{0,1\}^m$ *and* $y = C(x)$ *it holds that* $\mathsf{CEval}\left(\tilde{C}, \{k_{I,i}^{x_i}\}_{i \in [m]}\right) = \{k_{O,i}^{y_i}\}_{i \in [n]}$.

Figure 5. Scheme for garbling of gates and circuits with free xor.

Secrecy: *For a value $\Delta \in \{0,1\}^\kappa$, let $\mathcal{O}_\Delta(\cdot)$ be an oracle which on input Q outputs a bit b where $b = 1$ if and only if (iff) $Q = \Delta$. Consider the following indistinguishability under chosen input attack game for a stateful adversary $\mathcal{A}$.*

$$\underline{\mathsf{indcia}_{(\mathsf{CYao}, \mathsf{CEval})}^{\mathcal{A}}(\kappa)}$$

$C \leftarrow \mathcal{A}(1^\kappa)$, where $C : \{0,1\}^m \to \{0,1\}^n$ is a circuit as above.
$(x^0, x^1) \leftarrow \mathcal{A}(1^\kappa)$, where $x^0, x^1 \in \{0,1\}^m$.
$\left(\Delta, \left\{k_{I,i}^0\right\}_{i \in [m]}, \left\{k_{O,i}^0\right\}_{i \in [n]}, \tilde{C}\right) \leftarrow \mathsf{CYao}\left(C\right)$.
$c \leftarrow \{0,1\}$.
$b \leftarrow \mathcal{A}^{\mathcal{O}_\Delta(\cdot)}\left(\tilde{C}, \left\{k_{I,i}^{x_i^c}\right\}_{i \in [m]}\right)$.
If $\mathcal{O}_\Delta(\cdot)$ ever returned 1, then output 0.
Otherwise, output $b \oplus c$.

The game says that the adversary gets to choose a Boolean circuit C with m input bits and n output bits, along with two bit vectors of m bits, x^0 and x^1. The game then samples a random garbling of the circuit C, denoted by $\tilde{C}$ and flips a coin to choose a bit c. It then returns the input keys in correspondence with the bit vector x^c to the adversary along with the garbled circuit $\tilde{C}$. In the end the adversary returns a bit b and wins the game if $b = c$ or if he ever made the oracle query $\mathcal{O}_\Delta(\Delta)$.

We say that the scheme is IND-CIA *if for all PPT $\mathcal{A}$s, it holds that* $\left| \Pr\left(\text{indcia}^{\mathcal{A}}_{(\text{CYao},\text{CEval})}(\kappa) = 0\right) - \frac{1}{2} \right| = \text{negl}(\kappa)$.

The definition of secrecy captures two requirements at the same time: Given a garbled circuit and one key for each wire, including input keys corresponding to either x^0 or x^1, (1) one does not learn anything about the input, and (2) one cannot compute Δ. Assume namely that $\mathcal{A}$ could distinguish the two inputs x_0 and x_1 with advantage ε. Then he could output $c = b$ with probability $\frac{1}{2} + \varepsilon$, and then $b \oplus c = 0$ with probability $\frac{1}{2} + \varepsilon$, and hence $\left| \Pr\left(\text{indcia}^{\mathcal{A}}_{(\text{CYao},\text{CEval})}(\kappa) = 0\right) - \frac{1}{2} \right| \geq \varepsilon$. Assume on the other hand that $\mathcal{A}$ could compute Δ with probability ε. Then he can input Δ to the oracle when he has it. Furthermore, if he always outputs a uniformly random b. Then $\left| \Pr\left(\text{indcia}^{\mathcal{A}}_{(\text{CYao},\text{CEval})}(\kappa) = 0\right) - \frac{1}{2} \right| = \varepsilon$.

In Fig. 5, a free-xor garbled gates construction, with optimizations, is presented. We briefly sketch its properties. For correctness, remember that the requirement is: Let $(\tilde{g}, k_O^0) \leftarrow \text{Yao}(g, k_L^0, k_R^0, \Delta, id)$, then for all $a, b \in \{0, 1\}$ we have

$$\Pr\left[\text{Eval}\left(\tilde{g}, k_L^a, k_R^b\right) = k_O^{g(a,b)}\right] = 1 \ .$$

This is the case because, if we let $k_L' = k_L^a, k_R' = k_R^b$, then $\text{lsb}(k_L') = a \oplus p_L$ and $\text{lsb}(k_R') = b \oplus p_R$. Then by construction $\tau^j = g(a, b)$ for $j = 2p_L + p_R$ and

$$\alpha_j = k_O^{\tau^j} \oplus \text{KDF}(k_L^a, k_R^b, id)$$

and therefore

$$k_O' = k_O^{\tau^j} = k_O^{g((p_L \oplus a),(p_R \oplus b))} \ .$$

This garbling scheme can be proven secure under the assumption that $\text{KDF}(\cdot)$ behaves like a random function. Note that, as with any other free-xor based construction, we need to assume that the hash function we use is at least circular 2-correlation robust [CKKZ12].

To find the plain value of the output keys one use an output decryption table. An output decryption table for a pair of keys, (k_w^0, k_w^1), on wire w can for example be implemented as the pair $\left(\text{H}(k_w^0), \text{H}(k_w^1)\right)$ in the ROM. In particular we view the table as a verifiable commitment to the 0-, respectively 1-key, in that particular order. Our garbling construction will not make direct use of decryption tables and instead implement them with a black-box call to a verifiable commitment functionality, see Appendix C for details.

Secret Sharing. For our protocol we use a variation of Shamir's Secret Sharing scheme [Sha79] in order to make it possible to reconstruct keys the evaluator is not in possession of. We will secret share a key of κ bits in blocks of length e bits (we use $e = 8$ in the implementation). With this in hand we describe the simple polynomial interpolation methods we will use in our protocol in Fig. 6.

Links. We use a technical tool called *links*. A *link scheme* is a tuple (G, F) of PPT algorithms. The *link generator G* takes as input two values $x, y \in \{0, 1\}^\kappa$, called the *joints*, and outputs a *link $L \in \{0, 1\}^*$*. The *follower* allows to take one joint of a link and compute the other joint. To be more precise, for all $x, y \in \{0, 1\}^\kappa$ and $L \leftarrow G(x, y)$ it holds with probability 1 that $F(\texttt{forwards}, L, x) = y$ and $F(\texttt{backwards}, L, y) = x$.

As for security we want that a link leaks no other information than the ability to compute one joint from the other. In particular, to a party which does not know how to compute x nor y, the link should not reveal any information on x not y. We will model this by requiring that an adversary who cannot compute x nor y cannot distinguish a link of x and y from a link of any two other values.

Initialization

Initialized with the same parameters (t, ℓ, κ, e) from both Alice and Bob, with $2^e > \ell$ and e dividing κ, define the following polynomial manipulation methods. In the description we assume for notational convenience that $e = \kappa$. If $e < \kappa$, the values will be secret shared, interpolated and checked in κ/e blocks.

Construct

1. Alice construct a random polynomial in GF (2^e) of degree t by choosing $t + 1$ random values in $\{0, 1\}^e$ which will be the coefficients. Call this polynomial P.
2. Using P compute and return the following ℓ pairs $(i, P(i))_{i \in [\ell]}$.

Reconstruct$(\psi, (i, P(i))_{i \in \psi})$

1. If ψ is not a subset of $[\ell]$ with $|\psi| \geq t + 1$ then return `abort`. Otherwise take the $t + 1$ first pairs of $(i, P(i))_{i \in \psi}$ (call this set ψ') and do interpolation on these points. Call the polynomial that results from this interpolation P'.
2. Return $(i, P'(i))_{i \in [\ell]}$

Verify$(\psi, (i, P(i))_{i \in \psi})$

1. If ψ is not a subset of $[\ell]$ of at least $t+1$ entries then output `abort`. Otherwise take the $t+1$ first pairs of $(i, P(i))_{i \in \psi}$ (call this set ψ') and do interpolation on these points. Call the polynomial that results from this interpolation P'.
2. For the values in ψ not used in the interpolation use the interpolated polynomial P to verify the values lie on this polynomial. That is, for all $i \in \psi \backslash \psi'$ check if $P(i) =^? P'(i)$. If this is not true for some i return `reject` otherwise return `accept`.

Figure 6. The Methods for Polynomial Manipulation

Definition 3. *We call a distribution $(x_1, y_1, x_2, y_2, \ldots, x_\ell, y_\ell, z) \leftarrow D$ on $(\{0, 1\}^\kappa)^{2\ell} \times \{0, 1\}^*$ hard if it holds for all PPT B that it wins the following game with negligible probability: First sample $(x_1, y_1, x_2, y_2, \ldots, x_\ell, y_\ell, z) \leftarrow D$ and give z to B. Run B to produce $X \subset \{0, 1\}^*$, where the set is represented by a list of its elements. Then B wins iff $x_i \in X$ or $y_i \in X$ for some i. We call a class $\mathcal{D}$ of distributions hard if all $D \in \mathcal{D}$ are hard. We call (G, F) a secure link scheme for $\mathcal{D}$ if it holds for all PPT A and all distributions $D \in \mathcal{D}$ that A wins the following game with probability negligibly close to $\frac{1}{2}$: Sample $(x_1, y_1, x_2, y_2, \ldots, x_\ell, y_\ell, z) \leftarrow D$ and give z to A. Sample uniformly random $(x_1', y_1', x_2', y_2', \ldots, x_\ell', y_\ell', z) \in (\{0, 1\}^\kappa)^{2\ell}$. Sample a uniformly random bit $b \in \{0, 1\}$. If $b = 0$, let $(x_1'', y_1'', x_2'', y_2'', \ldots, x_\ell'', y_\ell'') = (x_1, y_1, x_2, y_2, \ldots, x_\ell, y_\ell)$. Otherwise, let $(x_1'', y_1'', x_2'', y_2'', \ldots, x_\ell'', y_\ell'') = (x_1', y_1', x_2', y_2', \ldots, x_\ell', y_\ell')$. For $i = 1, \ldots, \ell$, sample $L_i \leftarrow G(1^\kappa, x_i'', y_i'')$. Input $(L_1, \ldots, L_\ell)$ to A. Run A to produce a guess g. The adversary wins iff $g = b$.*

We can build a link scheme from a hash function H : $(\{0, 1\}^\kappa)^2 \to \{0, 1\}^\kappa$ as follows: On input $G(x, y)$, sample uniformly random salts $r, s \in \{0, 1\}^\kappa$. Then output $L = (r, s, h, g) = (r, s, \text{H}(x, r) \oplus y, \text{H}(y, s) \oplus x)$. On input $F(\texttt{forwards}, (r, s, h, g), x)$, output $\text{H}(x, r) \oplus h = y$. Similarly for the other direction, i.e. $F(\texttt{backwards}, (r, s, h, g), y) = \text{H}(y, s) \oplus h = y$.

It is straight forward to prove the above scheme secure in the random oracle model for the class of *all* hard distributions. Namely, until an adversary A queries some $\text{H}(\cdot, r_i)$ or $\text{H}(\cdot, s_i)$, the value $(r_i, s_i, \text{H}(x_i, r_i) \oplus y_i, \text{H}(y_i, s_i) \oplus x_i)$ is uniformly random in its view. So, turn it into an adversary B by giving it uniformly random values (r_i, s_i, g_i, h_i) as input, and then add each x_i queried to some $\text{H}(\cdot, r_i)$ or $\text{H}(\cdot, s_i)$ to the set X, and then at the end output X.

It appears, however, to be a very strong assumption to assume that the scheme is secure for all hard distributions in the plain model. Assume namely that one could black box obfuscate a function $f_{x,y}$ which takes as input a function F and a value L and then outputs 1 if $F(\texttt{forwards}, L, x) = y$ and $F(\texttt{backwards}, L, y) = x$. Let the obfuscation of $f_{x,y}$ be called $O_{x,y}$. Then sample (x, y) uniformly at random and output $(x, y, O_{x,y})$. This is a hard distribution, as it is hard to find x or y from just $O_{x,y}$. However, given $L = G(x, y)$ for a link scheme (G, F) one can run $O_{x,y}$ on F and L and will get back 1. Running it on a random link will return 0 except with negligible probability. Hence the link scheme is not secure when the side information $O_{x,y}$ is given. So, assuming that the link scheme is secure is also an assumption that $f = f_{x,y}$ cannot be black-box obfuscated. This is certainly a strong assumption. Note, however, that this might not be too unreasonable an assumption: According to the recent result [App13], black-box obfuscating any complexity class which allows to implement a pseudo-random function would allow to black-box obfuscate all poly-sized circuits, which is impossible. It appears hard to black-box obfuscate a sufficiently strong hash function and not being able to obfuscate a pseudo-random function. It also seems hard to imagine that one can obfuscate f but not

obfuscate the function H on which it is based. Hence it might just be impossible to obfuscate f. However, this is a rather fragile argument, so we would like to not rely on this faint hope.

Another way around the problem of an obfuscater as side information is to require that there is high min-entropy about either x or y given z. In that case z cannot contain an obfuscation containing x and y, like in the argument above. Indeed, it seems a reasonable assumption that if each x_i or y_i hash high min-entropy in the view of the adversary, then the above construction is a secure link scheme in the standard model for a sufficiently strong hash function H. Assume, e.g., that x_i has high min-entropy in the view of the adversary. Then it is reasonable to assume that $r_i, \mathrm{H}(x_i, r_i)$ acts as a strong extractor and that $r_i, \mathrm{H}(x_i, r_i,) \oplus y_i$ hides all partial information about y_i to a computationally bounded adversary. In that case y_i is still hard to compute given z and $r_i, \mathrm{H}(x_i, r_i) \oplus y_i$, and then it is reasonable to assume that $s_i, \mathrm{H}(y_i, s_i)$ acts as an extractor of the computational entropy of y_i and produces a pseudo-random value $s_i, \mathrm{H}(y_i, s_i)$, and that $s_i, \mathrm{H}(y_i, s_i) \oplus x_i$ hides all partial information about x_i to a computationally bounded adversary.

We will call a distribution super-hard if for each pair (x_i, y_i) one of the two values x_i or y_i has min-entropy $\Theta(k)$ bits in the view of the adversary. We will assume that our link scheme is secure for all super-hard distributions.

C   Some Sub-Functionalities and Their Implementations

In this appendix we describe the ideal functionality for secure two-party function evaluation that we realize with our protocol. We then describe some of the ideal functionalities that are used to realize our protocol in the hybrid model, and show how to realize them.

Secure Two-Party Function Evaluation.   In Fig. 7 the ideal functionality for secure function evaluation is presented.

Initialization

    Initialized by a circuit $\mathcal{C}$ and a statistical security parameter s.

Input Alice

    On input $(\texttt{input}, x)$ from Alice, where x is a possible first input to f and where no input was given by Alice before, store x, and output $\texttt{input}$ to Bob.

Input Bob

    On input $(\texttt{input}, y)$ from Bob, where y is a possible second input to f and where no input was given by Bob before, and at a point after $\texttt{input}$ was output to Bob, compute $z = f(x, y)$ and returns z value to Bob.

Corruption

    The adversary may input $(\texttt{corrupt}, \mathsf{Alice})$ or $(\texttt{corrupt}, \mathsf{Bob})$ to corrupt Alice or Bob respectively. Any corrupted party may then at any time input $\texttt{abort}$ in which case $\texttt{abort}$ is output to both Alice and Bob and the ideal functionality halts. If Alice is corrupted and does not abort, then Alice may input $(\texttt{cheat}, z')$. In that case the ideal functionality samples a bit $d \in \{0, 1\}$ with $d = 0$ with probability 2^{-s}. If $d = 0$, then output y to Alice and output z' to Bob. If $d = 1$, then output $\texttt{cheating}$ to Alice and Bob.

Figure 7. The ideal functionality, $\mathcal{F}_{\mathrm{SFE}}$, for secure function evaluation for two parties. The box reflects the structure of garbling and that only Bob gets output.

Commitments.   We also need a commitment scheme as defined in Fig. 8. `commit` and `open` define a standard commitment scheme while `vc`, `check` and `reject` define a verifiable commitment scheme.

In the ROM commitments can simply be implemented using a hash function. For standard commitments (`commit`), the comitter sends a hash of the message to be comitted concatenated with a random string (the random string can even be omitted if the value is chosen randomly from a big domain). Verifiable commitments can be done similarly, but without concatenating the message with a random string. For details on universally composable commitments in the ROM see [HMQ04].

Initialization

Both Alice and Bob input (`init`).

Input

- Upon receiving the command (`commit`, id, x) from either Alice or Bob, if id is a new unique identifier, then the functionality stores ($id, 0, x$, name) where name = Alice if Alice is the party sending the command and name = Bob otherwise. The functionality then outputs (`commit`, id) to both parties.
- Upon receiving the command (`vc`, id, x) from either Alice or Bob, if id is a new unique identifier, then the functionality stores ($id, 1, x$, name) where name = Alice if Alice is the party sending the command and name = Bob otherwise. The functionality then outputs (`vc`, id) to both parties.

Output

Upon receiving (`open`, id) from Alice if a commitment ($id, \cdot, \cdot$, Alice) is stored then the functionality outputs (`open`, id, x) to Bob. Similarly upon receiving (`open`, id) from Bob if a commitment ($id, \cdot, \cdot$, Bob) is stored then the functionality outputs (`open`, id, x) to Alice.

Query

Upon receiving (`check`, id, x) from any party the box outputs `accept` if a commitment ($id, 1, x, \cdot$) is stored. Otherwise it returns `reject`.

Figure 8. The ideal functionality $\mathcal{F}_{\text{COM}}$ for commitments.

Single Choice Batch Oblivious Transfer. For each circuit that Alice garbles, oblivious transfers must be carried out that lets Bob receive one key for each of his input wires without revealing to Alice if he asked for a 0- or 1-key.

Like in the rest of the paper we will use ℓ to denote the number of circuits to garble (that is, the replication factor) and n to denote the number of bits in Bob's input while κ will denote the amount of bits in each key.

We here consider a functionality that allows $n\ell$ OTs to be carried out in one "batch". In addition, while the functionality does not make any restrictions on the keys Alice inputs except that they are of κ bits, it does force Bob to use the same selection bit for all of the ℓ circuits. For this reason we will refer to the functionality as *batch oblivious transfer with consistent choice*. We denote it $\mathcal{F}_{\text{BOTCC}}^{\bar{n},\ell,\kappa}$ and it is outlined in Fig. 9. Note that $\mathcal{F}_{\text{BOTCC}}^{\bar{n},\ell,\kappa}$ should not be confused with the batch single choice cut-and-choose functionality used in [LP11] that also performs the cut-and-choose on the circuits.

Initialization

Initialized by ($\bar{n}, \ell, \kappa$).

Transfer

On input $\left(\texttt{transfer}, \{\{a_{i,j}^0, a_{i,j}^1\}_{i\in[\bar{n}]}\}_{j\in[\ell]}\right)$ from Alice where $a_{i,j}^0, a_{i,j}^1 \in \{0,1\}^\kappa$ for some non-negative integer constant κ and $\left(\texttt{transfer}, \{b_i\}_{i\in[\bar{n}]}\right)$ from Bob where $b_i \in \{0,1\}$, the functionality outputs $\left(\texttt{transfer}, \{\{a_{i,j}^{b_i}\}_{i\in[\bar{n}]}\}_{j\in[\ell]}\right)$ to Bob.

Corruption

A corrupted player can input `abort` at any time, in which case $\perp$ is output to the honest player and the functionality no longer responds to input.

Figure 9. The ideal functionality $\mathcal{F}_{\text{BOTCC}}^{\bar{n},\ell,\kappa}$ (with ℓ circuits, $\bar{n}$ input bits from Bob and keys of κ bits).

Given access to a random oracle, Fig. 11 shows how to realize $\mathcal{F}_{\text{BOTCC}}^{\bar{n},\ell,\kappa}$ efficiently in the $\mathcal{F}_{\text{OT}}$-hybrid model in the presence of malicious and static adversaries. The ideal functionality of $\mathcal{F}_{\text{OT}}$ is shown in Fig. 10 and can for example be realized as described in [PVW08].[12]

---

[12] We could always realize the functionality by $\bar{n}$ repeated executions of $\mathcal{F}_{\text{OT}}$ where each message was a string of $\ell \cdot \kappa$ bits. However, depending on the protocol used to realize $\mathcal{F}_{\text{OT}}$ this could be quite inefficient as group operations on group elements of at least $\ell \cdot \kappa$ bits might be needed.

<div style="border:1px solid black; padding:10px;">

Initialization

Both Alice and Bob input (init).

Transfer

- Upon receiving $(\mathsf{sid}, \mathsf{sender}, a^0, a^1)$ where $a^0, a^1 \in \{0,1\}^\kappa$ from the first party, if no tuple $(\mathsf{sid}, \mathsf{sender}, \cdot, \cdot)$ is stored, then store the tuple $(\mathsf{sid}, \mathsf{sender}, a^0, a^1)$.
- Upon receiving $(\mathsf{sid}, \mathsf{receiver}, b)$ where $b \in \{0,1\}$ from the other party, if a tuple $(\mathsf{sid}, \mathsf{sender}, \cdot, \cdot)$ is stored where none of the elements is the symbol $\perp$ then output (sid, a^b) to the second party and (sid) to the adversary, otherwise output nothing. Finally, update the tuple such that what is now stored is $(\mathsf{sid}, \mathsf{sender}, \perp, \perp)$.

Corruption

A corrupted player can input abort at any time, in which case $\perp$ is output to the honest player and the functionality no longer responds to input.

</div>

Figure 10. The ideal functionality $\mathcal{F}_{\mathrm{OT}}$.

<div style="border:1px solid black; padding:10px;">

Input

Alice holds ℓ sets of $\bar{n}$ key pairs, $\{\{(a_{i,j}^0, a_{i,j}^1)\}_{i \in [\bar{n}]}\}_{j \in [\ell]}$, with $a_{i,j}^0, a_{i,j}^1 \in \{0,1\}^\kappa$. Bob holds a set of $\bar{n}$ choice bits $\{b_i\}_{i \in [\bar{n}]}$, i.e., $b_i \in \{0,1\}$.

The Protocol

1. Alice and Bob initialize $\mathcal{F}_{\mathrm{OT}}$ by calling (init).
2. For each $i \in [\bar{n}]$ Alice calls $(i, \mathsf{sender}, r_i^0, r_i^1)$ where $r_i^0, r_i^1 \in_R \{0,1\}^\kappa$.
3. For each $i \in [\bar{n}]$ Bob calls $(i, \mathsf{receiver}, b_i)$ where $\{b_i\}_{i \in [\bar{n}]}$ is the set of his choice bits.
4. Thus, for each $i \in [\bar{n}]$ Bob learns $(i, r_i^{b_i})$ with $r_i^{b_i} \in \{0,1\}^\kappa$ from $\mathcal{F}_{\mathrm{OT}}$.
5. Bob now uses a hash function with κ bits output (modeled as a random oracle) to extend each of the $\bar{n}$ κ-bit strings to ℓ strings of κ bits. He computes:
$$r_{i,j}'^{b_i} = \mathrm{H}\left(r_i^{b_i} \| id_{i,j}\right) \in \{0,1\}^\kappa$$
for all $i \in [\bar{n}]$ and all $j \in [\ell]$ where $id_{i,j}$ is a unique identifier. Thus he ends up with a total of $\bar{n} \cdot \ell$ strings each of κ bits.
6. In a similar manner Alice uses the same hash function and IDs to extend each of her $2\bar{n}$ κ-bit inputs to ℓ strings of κ bits. She computes the following pairs for all $i \in [\bar{n}]$ and all $j \in [\ell]$:
$$\left(r_{i,j}'^0, r_{i,j}'^1\right) = \left(\mathrm{H}\left(r_i^0 \| id_{i,j}\right), \mathrm{H}\left(r_i^1 \| id_{i,j}\right)\right) \in \{0,1\}^\kappa \times \{0,1\}^\kappa.$$
Thus she ends up with a total of $2\bar{n} \cdot \ell$ strings each of κ bits; half of these are equal to the ones Bob ended up with in the previous step.
7. Alice then sends the differences needed to change these values to her original input set. That is, she computes and sends the following pairs to Bob:
$$\left(\delta_{i,j}^0, \delta_{i,j}^1\right) = \left(r_{i,j}'^0 \oplus a_{i,j}^0, r_{i,j}'^1 \oplus a_{i,j}^1\right).$$
8. Bob uses these pairs to reconstruct the input keys in correspondence with his input set $\{b_i\}_{i \in [\bar{n}]}$. Thus he simply computes and outputs:
$$a_{i,j}^{b_i} = \delta_{i,j}^{b_i} \oplus r_{i,j}'^{b_i}.$$

</div>

Figure 11. A protocol efficiently realizing $\mathcal{F}_{\mathrm{BOTCC}}^{\bar{n},\ell,\kappa}$ in the $\mathcal{F}_{\mathrm{OT}}$-hybrid model with access to a random oracle.

Theorem 1. *Assuming the hash function $H(\cdot)$ has κ bits output and is a random oracle then the protocol in Fig. 11 UC-securely realizes the $\mathcal{F}_{\mathrm{BOTCC}}^{\bar{n},\ell,\kappa}$ functionality in the $\mathcal{F}_{\mathrm{OT}}$-hybrid model in the presence of static, malicious adversaries.*

If Bob has many bits of input to the functionality he wishes to compute with Alice, i.e., if $\bar{n}$ is larger than cs for some constant c, then we can use *OT-extension* to limit the amount of heavy cryptographic operations needed to realize the $\mathcal{F}_{\mathrm{BOTCC}}^{\bar{n},\ell,\kappa}$ functionality. A very efficient way of doing this is described in [NNOB12].

Coin-tossing. In order to prove security by simulation we need to have some of the random choices made in the protocol be done using a coin-tossing functionality. We define an ideal functionality for coin-tossing in Fig. 12.

Initialization

Upon receiving inputs $(\texttt{init}, \kappa)$ from both Alice and Bob the functionality leaks κ to the adversary.

Output

The functionality outputs the same uniformly random bit string of κ bits to both Alice and Bob, i.e., $\rho \in \{0, 1\}^\kappa$. If Alice is corrupted, then Alice gets to see ρ first, and if she then inputs $\texttt{abort}$, then output $\texttt{abort}$ to Bob instead of ρ.

Figure 12. The ideal functionality $\mathcal{F}_{\text{CT}}$ for maliciously secure coin tossing of κ bits.

The Protocol

1. Bob selects a random string $\rho_B \in_R \{0, 1\}^\kappa$ and sends a hash digest of it to Alice.
2. Alice also chooses a random string, $\rho_A \in_R \{0, 1\}^\kappa$ and sends ρ_A to Bob.
3. Bob then sends the string ρ_B to Alice, who computes a hash digest of it and verifies that it equal to the digest she received in the first step. Furthermore, if $|\rho_B| \neq \kappa$ Alice outputs $\texttt{abort}$.
4. Alice and Bob now defines $\rho = \rho_A \oplus \rho_B$ and both output ρ.

Figure 13. A protocol realizing $\mathcal{F}_{\text{CT}}$ in the ROM.

UC-secure coin-tossing can be based UC-secure commitment in the plain model using the protocol in [Blu82] or in the ROM as described in Fig. 13.

D   The Full Protocol

The full protocol is described in the figures on the following pages and the proof is given afterwards in Appendix E.

We let Alice garble ℓ versions of a Boolean circuit C using the free-xor technique [KS08], we call these garbled circuits $\tilde{C}_j$ where $j \in [\ell]$. For each of these garbled circuits we associate a distinct global value $\Delta_j \in \{0, 1\}^\kappa$, which is used to enforce the constraint that for any wire, i, in the garbled circuit $\tilde{C}_j$, we have that $k_{i,j}^1 = k_{i,j}^0 \oplus \Delta_j$ in correspondence with a free-xor garbling scheme.

Now, remember that Alice is supposed to input $x \| a \in \{0, 1\}^{m+s}$ in the augmented circuit, we then say that the wires from 1 to m represent her input x and the wires from $m + 1$ to $m + s$ represent her auxiliary random input. Next, the wires from $m + s + 1 = \bar{m} + 1$ to $\bar{m} + \bar{n}$ will represent Bob's input. The wires from $\bar{m} + \bar{n} + 1$ to $\bar{m} + \bar{n} + |C|$ will represent the internal wires of the circuit, i.e., the output wires of the $|C|$ gates and in turn the wires from $\bar{m} + \bar{n} + |C| - o + 1$ to $\bar{m} + \bar{n} + |C| + o$ will represent the output z.

In the following we enumerate the output keys resulting from the augmentations from 1 to $\lceil 5.82s + 4.82 \rceil$. The first s will correspond to the digest on Alice's input and the last $\lceil 4.82s + 4.82 \rceil$ will correspond to the hash function applied to the output z. In particular notice that the non-augmented garbled circuit will not use the $m + 1$ to $m + s$'th input wires, as they will only be used "after the fact" to construct the augmentation on Alice's input.

To distinguish these keys from keys used in the actual circuit we call these keys $k_{i,j}'$ for $i \in [\lceil 5.82s + 4.82 \rceil]$ and $j \in [\ell]$. In a similar manner we will let $z_{i,j}'$ denote the bit value of the i'th augmented key in the j'th circuit.

Setup

1. Let $\bar{n} = \max\{4n, 8s\}$ and let ℓ be the smallest even integer satisfying $\ell - \frac{1}{2}\log(\ell) + \log\left(\frac{2\sqrt{2\pi}}{e^2}\right) \geq s$. Let p be the smallest integer satisfying $p \geq 6s + 7$. Let Alice and Bob initialize $\mathcal{F}_{\text{COM}}$, $\mathcal{F}_{\text{BOTCC}}^{\bar{n},\ell,\kappa}$, and $\mathcal{F}_{\text{CT}}$.

2. Bob chooses $M^{\text{Sec}} \in \{0,1\}^{n\times\bar{n}}$ and a random bit vector $\bar{y} \in \{0,1\}^{\bar{n}}$ under the constraint that $M^{\text{Sec}} \cdot \bar{y} = y$ and sends M^{Sec} to Alice.

3. Based on M^{Sec} the parties agree on a Boolean circuit C computing the function $f'(\bar{x}, \bar{y}) = f(x, y) = z$ where Alice gives $x \in \{0,1\}^{\bar{m}}$ as input and Bob gives $y \in \{0,1\}^{\bar{n}}$ as input and Bob receives $z \in \{0,1\}^o$ as output.

Polynomial Setup

1. Alice randomly chooses p polynomials of degree at most $t = \ell/2$. We call these polynomials $P_1, P_2, \ldots, P_p$. For $i \in [p]$ and $j \in [\ell]$ she then computes the points $P_i(j)$. For each $i \in [p]$ the ℓ points $(P_i(j))_{j\in[\ell]}$ thus uniquely define the polynomial P_i of degree at most t.

2. Using $\mathcal{F}_{\text{COM}}$ Alice commits to $P_i(j)$ for all $i \in [p]$ and $j \in [\ell]$. Formally, she calls $\mathcal{F}_{\text{COM}}\left(\text{vc}, \text{ID}_{i,j}^{\text{poly}}, P_i(j)\right)$ and Bob receives $\left(\text{vc}, \text{ID}_{i,j}^{\text{poly}}\right)$ from $\mathcal{F}_{\text{COM}}$ for all $i \in [p]$ and $j \in [\ell]$.

3. Bob chooses a challenge ψ, that is, a random set of distinct elements from $[p]$ such that $|\psi| = \lfloor 1.18s + 2.18 \rfloor$ and sends ψ to Alice.

4. Alice decommits to all the points corresponding to the received challenge. That is, for each $i \in \psi$ and $j \in [\ell]$ she calls $\mathcal{F}_{\text{COM}}\left(\text{open}, \text{ID}_{i,j}^{\text{poly}}\right)$ and in turn Bob learns $P_i(j)$.

5. Bob now verifies that for each $i \in \psi$ the points $(P_i(j))_{j\in[\ell]}$ define a polynomial of degree at most $t = \ell/2$ by computing **Verify**$(\psi, (i, P_i)_{i\in\psi})$. If any of these checks fail Bob outputs `abort` and halts. Otherwise define the set $\chi = [p]\backslash\psi$, which is then the indices of the remaining polynomials.

Oblivious Transfer

1. Alice garbles ℓ copies of C by setting $\left(\Delta_j, \{k_{i,j}^0\}_{i\in[\bar{m}+\bar{n}]}, \{k_{\bar{m}+\bar{n}+|C|+i,j}^0\}_{i\in[o]}, \tilde{C}_j\right) \leftarrow \text{CYao}(C)$ for all $j \in [\ell]$.

2. Alice then inputs to $\mathcal{F}_{\text{BOTCC}}^{\bar{n},\ell,\kappa}$ the set of pairs $\left\{\{k_{\bar{m}+i,j}^0, k_{\bar{m}+i,j}^1\}_{i\in[\bar{n}]}\right\}_{j\in[\ell]}$ and Bob inputs $\{\bar{y}_i\}_{i\in[\bar{n}]}$. Thus Bob learns the set $\left\{\{k_{\bar{m}+i,j}^{\bar{y}_i}\}_{i\in[\bar{n}]}\right\}_{j\in[\ell]}$.

Commitment

1. Alice sends the garbled circuits $\left\{\tilde{C}_j\right\}_{j\in[\ell]}$ to Bob. She also makes a verifiable commitment to both the 0- and 1-output keys of each garbled circuit. That is, she calls $\mathcal{F}_{\text{COM}}\left(\text{vc}, \text{ID}_{i,j}^{\text{out0}}, k_{\bar{m}+\bar{n}+|C|+i,j}^0\right)$ and $\mathcal{F}_{\text{COM}}\left(\text{vc}, \text{ID}_{i,j}^{\text{out1}}, k_{\bar{m}+\bar{n}+|C|+i,j}^1\right)$ for all $i \in [o]$ and $j \in [\ell]$.

2. Alice then uses $\mathcal{F}_{\text{COM}}$ to make verifiable commitments to the concatenation of the 0- and 1-key on each of her augmented input wires for each garbled circuit. She does so by calling $\mathcal{F}_{\text{COM}}\left(\text{vc}, \text{ID}_{i,j}^{\text{A-order}}, k_{i,j}^0 \| k_{i,j}^1\right)$ for all $i \in [\bar{m}]$ and $j \in [\ell]$.

3. She also commits to her augmented input by calling $\mathcal{F}_{\text{COM}}\left(\text{commit}, \text{ID}_{i,j}^{\text{A-in}}, k_{i,j}^{\bar{x}_i}\right)$ for all $i \in [\bar{m}]$ and $j \in [\ell]$.

4. Finally, she make verifiable commitments to the keys for Bob's input wires, that is, for all $i \in [\bar{n}]$ and $j \in [\ell]$ she calls $\mathcal{F}_{\text{COM}}\left(\text{vc}, \text{ID}_{\bar{m}+i,j}^{\text{B-order}}, k_{\bar{m}+i,j}^0 \| k_{\bar{m}+i,j}^1\right)$.

Augmentation

1. Now Bob samples three random bit vectors $b^{\text{In}} \in_R \{0,1\}^{\bar{m}-1}$, $b^{\text{Out1}} \in_R \{0,1\}^{\lceil o+4.82s+3.82 \rceil}$ and $b^{\text{Out2}} \in_R \{0,1\}^{\lceil 4.82s+4.82 \rceil}$. The first two vectors define the matrices $M^{\text{In}} \in \{0,1\}^{m\times s}$ and $M^{\text{Out}} \in \{0,1\}^{\lceil 4.82s+4.82 \rceil \times o}$ as $M_{i,j}^{\text{In}} = b_{i+j-1}^{\text{In}}$, respectively $M_{i,j}^{\text{Out}} = b_{i+j-1}^{\text{Out1}}$.

2. Bob sends the three vectors b^{In}, b^{Out1} and b^{Out2} to Alice who then defines the matrices M^{In} and M^{Out} similarly.

3. Alice augments the garbled circuits to compute $(M^{\text{In}} \cdot x) \oplus a$ where x is Alice's binary inputs (in column form) and a is her auxiliary binary inputs (in column form). That is, she computes the keys $k_{i,j}'^0 = \left(\bigoplus_{l=1}^m M_{i,l}^{\text{In}} \cdot k_{l,j}^0\right) \oplus k_{m+i,j}^0$, for all $i \in [s]$ and all circuits $j \in [\ell]$.

4. She then constructs verifiable commitments to the output keys of M^{In}. That is, she calls $\mathcal{F}_{\text{COM}}\left(\text{vc}, \text{ID}_{i,j}^{\text{out0-aug}}, k_{i,j}'^0\right)$ and $\mathcal{F}_{\text{COM}}\left(\text{vc}, \text{ID}_{i,j}^{\text{out1-aug}}, k_{i,j}'^1\right)$ for all $i \in [s]$ and $j \in [\ell]$.

5. Next, Alice augments the garbled circuits further such that they compute $(M^{\text{Out}} \cdot z) \oplus b^{\text{Out2}}$ where z is the binary output (in column form). That is, she computes the keys $k_{s+i,j}'^0 = \left(\bigoplus_{l=1}^o M_{i,l}^{\text{Out}} \cdot k_{l+\bar{m}+\bar{n}+|C|-o,j}^0\right) \oplus \left(b_i^{\text{Out2}} \cdot \Delta_j\right)$, for all $i \in [\lceil 4.82s+4.82 \rceil]$ and all circuits $j \in [\ell]$.

6. Alice then computes the links $L_{i,j} = G\left(P_{\chi[i]}(j), k_{s+i,j}'^0\right)$ for $i \in [\lceil 4.82s+4.82 \rceil]$ and $j \in [\ell]$. Alice then sends these links to Bob.

Figure 14. The Protocol – Part 1

Cut-and-choose

1. Alice and Bob invoke $\mathcal{F}_{\text{CT}}$ to randomly select $t = \ell/2$ elements from $[\ell]$. Call these t elements π and define $\phi = [\ell]\backslash\pi$. We say that the set π is the *check set* and that ϕ is the *evaluation* set.

2. For the check circuits Alice then opens the commitments to all the input keys along with the polynomial points associated with the augmented output wires. That is, she executes the following commands $\mathcal{F}_{\text{COM}}\left(\texttt{open}, \text{ID}_{i,j}^{\text{A-order}}\right)$ for all $i \in [\bar{m}]$, $j \in [\pi]$ and $\mathcal{F}_{\text{COM}}\left(\texttt{open}, \text{ID}_{i,j}^{\text{B-order}}\right)$ for all $i \in [\bar{n}]$, $j \in [\pi]$. Thus Bob receives from $\mathcal{F}_{\text{COM}}$ the values $k_{i,j}^0 \| k_{i,j}^1$ for all $j \in \pi$ and $i \in [\bar{m} + \bar{n}]$.

3. Bob uses the pairs of input keys to degarble the circuits $\tilde{C}_j$ and checks that they do in fact compute the function agreed upon in the *Setup* phase. Thus, letting $\tilde{g}_{i,j}$ denote the i'th garbled gate in the j'th garbled circuit, he calls $\tilde{k}_{i,j}^{a,b} \leftarrow \texttt{Eval}\left(\tilde{g}_{i,j}, k_{i,j}^a, k_{i,j}^b\right)$ for all $a, b \in \{0, 1\}$, $i \in [\bar{n} + \bar{m}]$ and $j \in \pi$. If he ever learns $\tilde{k}_{i,j}^{a,b} \neq k_{i,j}^{g(a,b)}$ for any $a, b \in \{0, 1\} \times \{0, 1\}$, he outputs $\texttt{abort}$. He also checks the verifiable commitments to the output keys. That is, he calls $\mathcal{F}_{\text{COM}}\left(\texttt{check}, \text{ID}_{i,j}^{\text{out0}}, k_{\bar{m}+\bar{n}+|C|+i}^0\right)$ and $\mathcal{F}_{\text{COM}}\left(\texttt{check}, \text{ID}_{i,j}^{\text{out1}}, k_{\bar{m}+\bar{n}+|C|+i}^0\right)$ for all $i \in [o]$ and $j \in \pi$ and outputs $\texttt{abort}$ if the functionality ever returns $\texttt{reject}$.

4. Bob also checks that the augmentation computing a digest on Alice's input has been constructed correctly by checking the verifiable commitments on the output keys of the check circuits. That is, he computes $\Delta_j = k_{1,j}^0 \oplus k_{1,j}^1$ for all $j \in \pi$. He then computes $k_{i,j}'^0 = \left(\bigoplus_{l=1}^m M_{i,l}^{\text{In}} \cdot k_{l,j}^0\right) \oplus k_{m+i,j}^0$, $k_{i,j}'^1 = k_{i,j}'^0 \oplus \Delta_j$ and then calls $\mathcal{F}_{\text{COM}}\left(\texttt{check}, \text{ID}_{i,j}^{\text{out0-aug}}, k_{i,j}'^0\right)$ and $\mathcal{F}_{\text{COM}}\left(\texttt{check}, \text{ID}_{i,j}^{\text{out1-aug}}, k_{i,j}'^1\right)$ for all $i \in [s]$ and $j \in \pi$. If $\mathcal{F}_{\text{COM}}$ ever returns $\texttt{reject}$ he outputs $\texttt{abort}$.

5. In the same manner Bob then uses all the output keys from the garbled circuits to compute the augmented output. That is, he computes $k_{s+i,j}'^0 = \left(\bigoplus_{l=1}^o M_{i,l}^{\text{Out}} \cdot k_{l+\bar{m}+\bar{n}+|C|-o,j}^0\right) \oplus \left(b_i^{\text{Out2}} \cdot \Delta_j\right)$, $k_{s+i,j}'^1 = k_{s+i,j}'^0 \oplus \Delta_j$ and uses these values to compute $P_{\chi[i]}(j)$ via the link $L_{i,j}$. That is he computes $F(\texttt{backwards}, L_{i,j}, k_{s+i,j}'^0) = P_{\chi[i]}(j)$. He then verifies these values by calling $\mathcal{F}_{\text{COM}}\left(\texttt{check}, \text{ID}_{\chi[i],j}^{\text{poly}}, P_{\chi[i]}(j)\right)$ for all $i \in [\lceil 4.82s + 4.82 \rceil]$ and $j \in \pi$.

6. If any of the above checks fail or if $\mathcal{F}_{\text{COM}}$ returns $\texttt{reject}$ on any of the calls Bob outputs $\texttt{abort}$ and terminates the protocol.

Evaluation

1. Alice now opens her input keys by calling $\mathcal{F}_{\text{COM}}\left(\texttt{open}, \text{ID}_{i,j}^{\text{A-in}}\right)$ for all $i \in [\bar{m}]$ and $j \in \phi$ from which Bob in turn learns $k_{i,j}^{\bar{x}_i}$.

2. Using Alice's input keys Bob evaluates the augmentation for consistency checking of Alice's input for all $j \in \phi$. Thus he computes $k_{i,j}' = \left(\bigoplus_{l=1}^m M_{i,l}^{\text{In}} \cdot k_{l,j}^{\bar{x}_l}\right) \oplus k_{m+i,j}^{\bar{x}_{m+i}}$ for all $i \in [s]$ and $j \in \phi$. He then uses the verifiable commitments to find the bit each key $k_{i,j}'$ represent. Thus he calls $\mathcal{F}_{\text{COM}}\left(\texttt{check}, \text{ID}_{i,j}^{\text{out0-aug}}, k_{i,j}'\right)$ and $\mathcal{F}_{\text{COM}}\left(\texttt{check}, \text{ID}_{i,j}^{\text{out1-aug}}, k_{i,j}'\right)$ and sets the bit $z_{i,j}' = 0$ if the first query returns $\texttt{accept}$, and $z_{i,j}' = 1$ if instead the second query returns $\texttt{accept}$. If both queries return $\texttt{accept}$ or both return $\texttt{reject}$ he then outputs $\texttt{abort}$. He then verifies that $q_{i,j} = q_{i,j'}$ for all $i \in [s]$ and $j, j' \in \phi$ and if any check fails he outputs $\texttt{abort}$.

3. Bob now uses Alice's input keys along with his own input keys which he learned in the *Oblivious Transfer* phase to evaluate all garbled circuits in the set ϕ. That is, he sets $\left\{k_{\bar{m}+\bar{n}+|C|+i}\right\}_{i\in[o]} \leftarrow \texttt{CEval}\left(\tilde{C}_j, \{k_{i,j}\}_{i\in[\bar{m}+\bar{n}]}\right)$ for all $j \in \phi$. For each j where $\tilde{C}_j = \bot$ he discards $\tilde{C}_j$.

4. He finds the output bits the keys represent (i.e., the semantic values) by calling $\mathcal{F}_{\text{COM}}\left(\texttt{check}, \text{ID}_{i,j}^{\text{out0}}, k_{\bar{m}+\bar{n}+|C|+i}\right)$ and $\mathcal{F}_{\text{COM}}\left(\texttt{check}, \text{ID}_{i,j}^{\text{out1}}, k_{\bar{m}+\bar{n}+|C|+i}\right)$ and sets the bit $z_{i,j} = 0$ if the first query returns $\texttt{accept}$, and $z_{i,j} = 1$ if instead the second query returns $\texttt{accept}$. If both queries return $\texttt{accept}$ or both return $\texttt{reject}$ he then sets $z_{i,j} = \bot$.

5. If $z_{i,\phi[1]} = z_{i,j}$ for all $j \in \phi\backslash\phi[1]$ then Bob accepts and returns this output. If not then Bob proceeds with the *Reconstruction* phase.[a]

---

[a] In order to avoid timing attacks Bob should always do the following computations for simulated values.

Figure 15. The Protocol – Part 2

Reconstruction

1. Bob computes the bits $z'_{s+i,j} = \left(\bigoplus_{l=1}^{o} M_{i,l}^{\mathsf{Out}} \cdot z_{l,j}\right) \oplus b^{\mathsf{Out2}}$ for all $i \in [[\lceil 4.82s + 4.82 \rceil]]$ and $j \in \phi$.

2. Similarly he computes the augmented output keys as $k_{s+i,j}^{\prime z'_{s+i,j}} = \left(\bigoplus_{l=1}^{o} M_{i,l}^{\mathsf{Out}} \cdot k_{l,j}^{z_{l,j}}\right)$ for all $i \in [[\lceil 4.82s + 4.82 \rceil]]$ and $j \in \phi$.

3. Now for each augmented output wire, evaluating to a semantic 0, Bob computes the values $P_{\chi[i]}(j)$ for all $i \in [[\lceil 4.82s + 4.82 \rceil]]$ and $j \in \phi$. That is, he uses link $L_{i,j}$ to compute $P_{\chi[i]}(j) = F(\texttt{backwards}, L_{i,j}, k_{s+i,j}^{\prime z'_{s+i,j}})$. for all $i \in [[\lceil 4.82s + 4.82 \rceil]]$ and $j \in \phi$ where $z'_{s+i,j} = 0$. For each of these values, he calls $\mathcal{F}_{\mathsf{COM}}\left(\texttt{check}, \mathsf{ID}_{\chi[i],j}^{\mathsf{poly}}, P_{\chi[i]}(j)\right)$. If $\mathcal{F}_{\mathsf{COM}}$ returns $\texttt{reject}$ on any of these instances he discards the associated circuit j.

4. Bob now takes the remaining garbled circuits that evaluate to different values, and looks at all the augmented output wires that evaluates differently between at least two of these garbled circuits to find a new polynomial point. That is, define $U \subseteq \phi$ to be the set of circuits which has not been discarded and define $W \subseteq [[\lceil 4.82s + 4.82 \rceil]]$ such that for each $i \in W$ there exist a pair of indices $j, j' \in U$ where $z'_{s+i,j} \neq z'_{s+i,j'}$ and let $V = W \times U$ be the Cartesian product of U and W. Now for each $i \in W$ do the following:[a]

   (a) Find the first element $j \in U$ for which it is true that $z'_{s+i,j} = 0$. Now using link $L_{i,j}$ he calls $\mathcal{F}_{\mathsf{COM}}\left(\texttt{check}, \mathsf{ID}_{i,j}^{\mathsf{poly}}, F(\texttt{backwards}, L_{i,j}, k'_{s+i,j})\right)$. If $\mathcal{F}_{\mathsf{COM}}$ returns $\texttt{reject}$ then take the next element (if it exist) $j \in U$ where $z'_{s+i,j} = 0$ and call $\mathcal{F}_{\mathsf{COM}}\left(\texttt{check}, \mathsf{ID}_{i,j'}^{\mathsf{poly}}, F(\texttt{backwards}, L_{i,j}, k'_{s+i,j})\right)$ and continue taking elements (if they exist) until $\mathcal{F}_{\mathsf{COM}}$ returns $\texttt{accept}$.

   (b) When $\mathcal{F}_{\mathsf{COM}}$ returns $\texttt{accept}$ do polynomial interpolation using the set of pairs $S = \{\pi \cup j\} \times \left\{P_i(l)_{l \in \pi} \cup F(\texttt{backwards}, L_{i,j}, k'_{s+i,j})\right\}$ by calling **Reconstruct**$(\pi \cup j, S)$. Then the result should be $(l, P_i(l))_{l \in [\ell]}$. For each $l \in [\ell] \setminus \{\pi \cup j\}$ call $\mathcal{F}_{\mathsf{COM}}\left(\texttt{check}, \mathsf{ID}_{i,l}^{\mathsf{poly}}, P_i(l)\right)$.

   (c) Next, for each element $j' \in U$ for which it is true that $z'_{s+i,j} = 1$ compute the values $\Delta_{j'} = F(\texttt{forwards}, L_{i,j}, P_i(j')) \oplus k'_{s+i,j'}$.

5. Now Bob should know Δ_j for all $j \in \phi$ where the garbled circuit has not bee discarded. If he does not, then he outputs $\texttt{abort}$ and terminates.

6. Using Δ_j and Alice's input keys Bob can find Alice's input for all circuits calling $\mathcal{F}_{\mathsf{COM}}\left(\texttt{check}, \mathsf{ID}_{i,j}^{\mathsf{A-order}}, k_{i,j}^{x_i} \| k_{i,j}^{x_i} \oplus \Delta_j\right)$ and $\mathcal{F}_{\mathsf{COM}}\left(\texttt{check}, \mathsf{ID}_{i,j}^{\mathsf{A-order}}, k_{i,j}^{x_i} \oplus \Delta_j \| k_{i,j}^{x_i}\right)$, if the first call returns $\texttt{accept}$, Alice's bit $x_i = 0$, if the second call returns $\texttt{accept}$ then $x_i = 1$, otherwise he outputs $\texttt{abort}$ and terminates the protocol.

7. If Alice's input bits are the same in all the remaining garbled circuits Bob evaluates the function $f(x, y) = z$ unencrypted and finally outputs $\texttt{ok}$ and uses z as his output of the computation. However, if there is a discrepancy he evaluates $\mathbb{H}^{\mathsf{In}} \oplus \mathsf{a}$ in plain and compares the output with the semantic meaning of the output of $\mathbb{H}^{\mathsf{In}} \oplus \mathsf{a}$ he found in *Evaluation*. He discards any circuit where the plain and the semantic output are discrepant. Finally, he uses Alice's input of any remaining circuit, to evaluate the function f in plain.[b]

---

[a] The following should be done for each j being an index of an evaluation circuit in order to avoid timing attacks.

[b] Alice's input will be unique on each of the remaining garbled circuit unless there is a collision in $\mathbb{H}^{\mathsf{In}} \oplus \mathsf{a}$, which only happens with probability at most 2^{-s}.

Figure 16. The Protocol – Part 3

E   Proof of Security

Before continuing with the actual proof we need the following preliminaries about universal hash functions:

Universal Hash Functions.

Definition 4 ((Binary) universal hash function). *Given a family of functions* $\mathcal{H} = \{h : \{0,1\}^q \to \{0,1\}^w\}$ *is a family of universal hash functions if*

$$\forall x, y \in \{0,1\}^q, x \neq y \ : \ \Pr_{h \in_R \mathcal{H}}[h(x) = h(y)] \leq 2^{-w} \ .$$

Furthermore, we say that a family of universal hash functions have the uniform difference property *if*

$$\forall x, y \in \{0,1\}^q, x \neq y \ : \ \forall z \in \{0,1\}^w \ : \ \Pr_{h \in_R \mathcal{H}}[h(x) \oplus h(y) = z] = 2^{-w} \ .$$

We say that the family is pairwise independent *if*

$$\forall x, y \in \{0,1\}^q, x \neq y \ : \ \forall z_1, z_2 \in \{0,1\}^w \ : \ \Pr_{h \in_R \mathcal{H}}[h(x) = z_1 \wedge h(y) = z_2] = 2^{-2w} \ .$$

Finally, we say that family is uniform *if*

$$\forall x \in \{0,1\}^q, \forall z \in \{0,1\}^w \ : \ \Pr_{h \in_R \mathcal{H}}[h(x) = z] = 2^{-w}$$

From the definition notice that pairwise independence is the strongest type of family as it implies both uniformity and the uniform difference property. Furthermore, notice that if the family is uniform then it also satisfies the uniform difference property, but the converse is not necessarily true.

Definition 5 ($\mathbb{H}^{\mathsf{ln}}$ **(formal)).** *Let* $b \in \{0,1\}^{m+s-1}$ *and* $M^{\mathsf{ln}} \in \{0,1\}^s \times \{0,1\}^m$ *be a matrix with each entry defined as* $M^{\mathsf{ln}}_{i,j} = b_{i+j-1}$. *Now define the function* $h_b : x \to z$ *as* $h_b(x) = M^{\mathsf{ln}}x = z$ *when queried on* b. *Finally define a family of functions* $\mathbb{H}^{\mathsf{ln}} = h_b(\cdot)_{b \in \{0,1\}^{m+s-1}}$.

Lemma 1. $\mathbb{H}^{\mathsf{ln}}$ *is a universal hash function with the uniform difference property.*

Proof. First notice that it is enough to show the uniform difference property by the following:

$$\forall x, y \in \{0,1\}^m, x \neq y \ : \ \forall z \in \{0,1\}^s \ : \ \Pr_{h \in_R \mathcal{H}}[h(x) \oplus h(y) = z] = 2^{-s} \Rightarrow$$
$$\forall x, y \in \{0,1\}^m, x \neq y \ : \ \forall z \in \{0,1\}^s \ : \ \Pr_{h \in_R \mathcal{H}}[h(x) = z \oplus h(y)] = 2^{-s} \ .$$

This must in particular be true when $z = 0^s$ and so $\forall x, y \in \{0,1\}^m, x \neq y \ : \ \Pr_{h \in_R \mathcal{H}}[h(x) = h(y)] = 2^{-s} \leq 2^{-s}$, which meets the definition of a universal hash function.

Now what we need to show is $\Pr[(M^{\mathsf{ln}}x) \oplus (M^{\mathsf{ln}}x') = z] = 2^{-s}$ for any $z \in \{0,1\}^s$ when each entry of M^{ln} is defined as $M^{\mathsf{ln}}_{i,j} = b_{i+j-1}$, $b \in_R z^{m+s-1}$ and $x \neq x'$. First notice that matrix multiplication is linear, so these are equivalent:

$$(M^{\mathsf{ln}}x') \oplus (M^{\mathsf{ln}}x) = z$$
$$M^{\mathsf{ln}}(x \oplus x') = z$$

Thus it is enough to show that $\Pr[M^{\mathsf{ln}}(x \oplus x') = z] = 2^{-s}$ for $x \neq x'$ and any z. We do this by showing that for any choice of $(x \oplus x') \neq 0$ there exist 2^{m-1} choices of b for each of the 2^s possible elements in the range of the function. Thus when b is chosen at random (from $\{0,1\}^{m+s-1}$) there is uniform probability of hitting each of the elements in the function's range, in particular of hitting z.

Now set $\tilde{x} = x \oplus x'$ and see that given b, $M^{\mathsf{In}}\tilde{x}$ can be computed as follows:

$$M^{\mathsf{In}}\tilde{x} = \begin{pmatrix} b_1 & b_2 & \cdots & b_m \\ b_2 & b_3 & \cdots & b_{m+1} \\ \vdots & \vdots & \ddots & \vdots \\ b_s & b_{s+1} & \cdots & b_{m+s-1} \end{pmatrix} \cdot \begin{pmatrix} \tilde{x}_1 \\ \tilde{x}_2 \\ \vdots \\ \tilde{x}_m \end{pmatrix}.$$

That is, as an $s \times m$ binary matrix times a column vector of m bits. Notice that we can rewrite this computation as follows:

$$\begin{pmatrix} b_1 & b_2 & \cdots & b_m \\ b_2 & b_3 & \cdots & b_{m+1} \\ \vdots & \vdots & \ddots & \vdots \\ b_s & b_{s+1} & \cdots & b_{m+s-1} \end{pmatrix} \cdot \begin{pmatrix} \tilde{x}_1 \\ \tilde{x}_2 \\ \vdots \\ \tilde{x}_m \end{pmatrix} = \begin{pmatrix} \tilde{x}_1 \ \tilde{x}_2 \ \tilde{x}_3 \cdots \tilde{x}_m \ 0 \ \cdots \ 0 \\ 0 \ \tilde{x}_1 \ \tilde{x}_2 \cdots \tilde{x}_m \ 0 \ \cdots \ 0 \\ 0 \ 0 \ \tilde{x}_1 \cdots \tilde{x}_m \ 0 \ \cdots \ 0 \\ \vdots \ \vdots \ \vdots \ \ddots \ \vdots \ \vdots \ \vdots \ \vdots \\ 0 \ 0 \ 0 \ \cdots \ 0 \ \tilde{x}_1 \ \cdots \ \tilde{x}_m \end{pmatrix} \cdot \begin{pmatrix} b_1 \\ b_2 \\ b_3 \\ \vdots \\ b_{m+s-1} \end{pmatrix}.$$

A bit more specific we view b as a column vector of $m+s-1$ bits and define a binary matrix $X \in \{0,1\}^s \times \{0,1\}^{m+s-1}$ where

$$X_{i,j} = \begin{cases} \tilde{x}_{j-i+1} & \text{if } j - i + 1 \le m \\ 0 & \text{otherwise} \end{cases}$$

Now see that since $\tilde{x} \ne 0^{m+s-1}$ there must be at least one bit, say at position i, that is 1. Thus, we see that matrix X will be of the following form:

$$\begin{pmatrix} \tilde{x}_1 \ \cdots \ 1 \ \cdots \ \tilde{x}_m \ 0 \ 0 \ \cdots \ 0 \\ 0 \ \tilde{x}_1 \ \cdots \ 1 \ \cdots \ \tilde{x}_m \ 0 \ \cdots \ 0 \\ \vdots \ \vdots \ \vdots \ \ddots \ \vdots \ \vdots \ \vdots \ \vdots \\ 0 \ 0 \ \cdots \ 0 \ \tilde{x}_1 \ \cdots \ 1 \ \cdots \ \tilde{x}_m \end{pmatrix}.$$

That is, the matrix is in row echelon form and does not have any all-0 rows. Thus the matrix has full rank (s) and therefore is surjective and so we can hit each of the 2^s possible s-bit outputs.

Next consider row i and column j. No matter what the value is of all the other $m + s - 2$ entries in row i are set to, the i'th bit in the output vector can be uniquely decided by setting b_i to either 0 or 1. That is, if $b_i = 0$ the $m + s - 2$ other entries in row i and the vector b computes the i'th output bit, if it is set to 1, the negation of this will be the i'th output bit. The point is that b_i can be used to hit either 0 or 1 no matter what the other bits in the matrix and the vector b are set to. This is clearly true individually for all the rows. However, as the output is a vector and we use one bit of b for each row to determine each output bit we see that for a given output vector (of s bits) we need s bits of b to uniquely decide its value, whereas the rest of the $m + s - 1 - s = m - 1$ bits of b remain free no matter what output we desire and what the input is. Thus there will exist at least 2^{m-1} choices of b to give a specific output, in particular z on any non-zero choice of x. But since we have already established, by the fact that our new matrix description of the computation has full rank, it means that this is the case for all 2^s possible outputs, thus there can only be exactly 2^{m-1} choices of b for each output, and in particular z, as $2^s \cdot 2^{m-1} = 2^{m+s-1}$ covers all the possible choices of b. $\square$

Definition 6 ($\mathbb{H}^{\mathsf{Out}}$ **(formal)**). *Let* $b^{\mathsf{Out1}} \in \{0,1\}^{\lceil o + 4.82s + 4.82 \rceil - 1}$, $b^{\mathsf{Out2}} \in \{0,1\}^{\lceil 4.82s + 4.82 \rceil}$ *and* $M^{\mathsf{Out}} \in \{0,1\}^{\lceil 4.82s + 4.82 \rceil} \times \{0,1\}^o$ *be a matrix with each entry defined as* $M^{\mathsf{Out}}_{i,j} = b^{\mathsf{Out1}}_{i+j-1}$. *Now define the function*

$$h'_{b^{\mathsf{Out1}}, b^{\mathsf{Out2}}} : x \to z \text{ as } h'_{b^{\mathsf{Out1}}, b^{\mathsf{Out2}}}(x) = (M^{\mathsf{Out}}x) \oplus b^{\mathsf{Out2}} = z.$$

Finally, define a family of functions

$$\mathbb{H}^{\mathsf{Out}} = h'_{b^{\mathsf{Out1}}, b^{\mathsf{Out2}}}(\cdot)_{b^{\mathsf{Out1}} \in \{0,1\}^{\lceil 4.82s + 4.82 \rceil}, b^{\mathsf{Out2}} \in \{0,1\}^{\lceil o + 4.82s + 4.82 \rceil - 1}}.$$

Lemma 2. $\mathbb{H}^{\mathsf{Out}}$ *is a pairwise independent universal hash function.*

Proof. A proof can be found in [MNT90]. $\square$

Proof of Protocol Security. We prove the protocol secure according to the standard real/ideal simulation paradigm with sequential composition [Can00,Gol04,HL10] under the assumption that our garbling scheme is IND-CIA secure, we have a secure link scheme in accordance with Def. 3 and hybrid access to the $\mathcal{F}_{\text{BOTCC}}$, $\mathcal{F}_{\text{CT}}$ and $\mathcal{F}_{\text{COM}}$ functionalities. In the end we argue how the protocol can be proven secure in the UC model [Can01] and show how to efficiently realize it in the random oracle, $\mathcal{F}_{\text{OT}}$-hybrid model.

Theorem 2. *Let π be the protocol in Fig. 14 to Fig. 16. Assume that the garbling scheme $\mathcal{G} = (\mathsf{CYao}, \mathsf{CEval})$ used in π is* IND-CIA *secure according to Def. 2 and the link scheme (G, F) used is secure according to Def. 3. Then π securely realizes $\mathcal{F}_{\text{SFE}}$ in Fig. 7 in the $(\mathcal{F}_{\text{BOTCC}}, \mathcal{F}_{\text{CT}}, \mathcal{F}_{\text{COM}})$-hybrid model with abort in the presence of static, malicious adversaries.*

Proof. We split the proof up in two lemmas, one considering a corrupted Alice and one where Bob is corrupted. The two lemmas are treated in Appendix E.1 and Appendix E.2, respectively. Theorem 2 follows directly from these two lemmas. □

E.1 Corrupt Alice

Lemma 3. *Let π be the protocol in Fig. 14 to Fig. 16. Assume that the garbling scheme $\mathcal{G} = (\mathsf{CYao}, \mathsf{CEval})$ used in π is* IND-CIA *secure according to Def. 2 and the link scheme (G, F) used is secure according to Def. 3 then π securely realizes $\mathcal{F}_{\text{SFE}}$ in Fig. 7 in the $(\mathcal{F}_{\text{BOTCC}}, \mathcal{F}_{\text{CT}}, \mathcal{F}_{\text{COM}})$-hybrid model with abort in the presence of static, malicious adversary corrupting* Alice.

Proof. Before proceeding with the formal proof we give some high-level intuition. Notice first that if Alice tries to cheat she can only do so before the cut-and-choose phase as there is no interaction (except opening of commitments) with her after this phase.

A common strategy [LP07,LP11,FN13,sS11] is to take the result to be the *majority* of the results of the evaluation circuits. It can be shown that selecting half of the circuits as evaluation circuits, the majority of the evaluation circuits are correct except with probability at most 2^{-s} if $\ell = 3.22s$ [LP11]. It was later shown [sS11] that with the majority approach, selecting 3/5 of the circuits as check circuits allows for a replication factor of $\ell = 3.13s$.

However, in our approach we make sure that two evaluation circuits evaluating to different results enable Bob to recover Alice's plain input for these circuits. Hence, intuitively, a cheating Alice must guess *exactly* the cut-and-choose challenge (in order to ensure that all check circuits are good while all evaluation circuits are bad). In this case the optimal ratio between check and evaluation circuits is 1/2. This is seen as a cheating Alice's success probability is $\binom{s}{cs}^{-1}$ when $c \leq 1$. Now, $\binom{s}{cs}^{-1}$ is clearly minimized when $c = 1/2$ since $\binom{s}{cs}$ for $0 < c \leq 1$ is maximized when $c = 1/2$.

The following lemma gives a lower bound on the replication factor ℓ required in our protocol (that is, how many circuits to garble) in order to ensure that Alice guesses exactly the cut-and-choose challenge with probability at most 2^{-s}.

Lemma 4. *The probability of guessing exactly the elements in a set of size $\ell/2$ sampled randomly from a set of ℓ elements is less than 2^{-s} when $\log\left(\frac{2\sqrt{2\pi}}{e^2}\right) - \frac{1}{2}\log(\ell) + \ell \geq s$.*

Proof. See Appendix F. □

Choosing the replication factor according to Lemma 4, the only problem is therefore that Bob might not know which circuit is the correct one. Thus, if more than one circuit "seemingly" evaluate correctly Bob must be able to discover which of these evaluations are correct. He does that by restoring Alice's input as mentioned in the *Reconstruction* phase and then evaluating the function f in plain. We now proceed with the formal proof.

Let $\mathcal{A}$ be an adversary controlling Alice in the real world execution of protocol π. We construct a simulator $\mathcal{S}$ running in the ideal world with external access to $\mathcal{F}_{\text{SFE}}$. Internally $\mathcal{S}$ runs $\mathcal{A}$, controlling Alice, along with a simulation of the interaction between $\mathcal{A}$, an honest Bob and the ideal functionalities ($\mathcal{F}_{\text{BOTCC}}$, $\mathcal{F}_{\text{CT}}$, $\mathcal{F}_{\text{COM}}$). See Fig. 17 for an illustration of this. The simulator $\mathcal{S}$ works as follows:

Figure 17. The ideal and the real world execution with corrupt Alice ($\mathcal{A}$). We must prove that for all adversaries $\mathcal{A}$ there exists a simulator $\mathcal{S}$ such that for all input x, y and all auxiliary input aug to $\mathcal{A}$ and all values of s it holds that the (joint) distributions $\mathrm{IDEAL}_{\mathcal{S}}(x, y, z, \kappa)$ and $\mathrm{REAL}_{\mathcal{A}}(x, y, z, \kappa)$ are computationally indistinguishable in κ.

- $\mathcal{S}$ starts simulating internally the real world execution of protocol π consisting of $\mathcal{A}$, Bob and the ideal functionalities $\mathcal{F}_{\mathrm{BOTCC}}$, $\mathcal{F}_{\mathrm{CT}}$ and $\mathcal{F}_{\mathrm{COM}}$. If an honest Bob would abort at any time during the simulation, then $\mathcal{S}$ sends `abort` to $\mathcal{F}_{\mathrm{SFE}}$.
- The simulator then simulates the *Setup*, *Polynomial Setup* and *Oblivious Transfer* phases, but simulating his "real" input as the all-0 string, i.e., it sets $y = 0^n$.
- When reaching the *Commitment* phase the simulator extracts the ℓ (possibly different) inputs of $\mathcal{A}$ to the ℓ garbled circuits it has constructed. It does so directly since $\mathcal{A}$ calls $\mathcal{F}_{\mathrm{COM}}$ with the keys of its input wires. We will later determine which of these inputs is the one $\mathcal{A}$ is actually committed to. In the rest of the proof we call the i'th input bit in the j'th circuit given by $\mathcal{A}$ for x'^j_i. When we consider the entire vector of input bits in the j'th circuit we remove the subscript, i.e., x'^j, and when we consider the input to the ideal functionality we simply call it x'. In either case if we consider the augmented input we add a bar, i.e., $\bar{x}'^j_i$, $\bar{x}'^j$ and $\bar{x}'$.
- The simulator continues and chooses two random universal hash functions in the *Augmentation* phase and then runs the *Cut-and-Choose*, *Evaluation* and *Reconstruction* phases.
- At the end of the *Evaluation* phase an honest Bob (and in this case the simulator) will have discarded any inconsistent garbled circuits received by $\mathcal{A}$. Thus what remains are garbled circuits that can be evaluated. If they all give the same output then the simulator will use one of Alice's input vectors extracted during the *Oblivious Transfer* phase to one of the evaluable garbled circuits and input such a vector to $\mathcal{F}_{\mathrm{SFE}}$. Otherwise, it will continue simulating an honest Bob in the *Reconstruction* phase and in turn find Alice's input to each of the garbled circuits. In particular $\mathcal{S}$ will, like honest Bob, then evaluate the function $\mathbb{H}^{\mathsf{In}} \oplus \mathsf{a}$ in plain and compare the plain output with the semantic meaning of the output he learned in *Evaluation* to find an input of Alice given to the correctly constructed circuits. $\mathcal{S}$ then uses such an input of Alice as input to $\mathcal{F}_{\mathrm{SFE}}$. Notice that there can only be one suitable set of inputs (except with probability 2^{-s}) since having more would imply that either the verifiable commitments on the output keys of $\mathbb{H}^{\mathsf{In}} \oplus \mathsf{a}$ were incorrectly constructed, in which case the simulator would have discarded the circuit, or Alice has guessed a collision for $\mathbb{H}^{\mathsf{In}} \oplus \mathsf{a}$, which can only happen with probability 2^{-s}.

Our first observation is that in the real (hybrid) world execution Bob's input $\bar{y}$ is only used as input to the ideal functionality computing $\mathcal{F}^{\bar{n}, \ell, \kappa}_{\mathrm{BOTCC}}$. Consequently, the messages that Bob sends during the remaining part of the protocol and hence the view of $\mathcal{A}$ do not depend on $\bar{y}$. Therefore, the output of $\mathcal{A}$ in the real world execution and $\mathcal{S}$ in the ideal world are identical.

We must, however, prove that the *joint* distribution of $\mathcal{A}$'s and Bob's output in the real world execution is identical to the *joint* distribution of $\mathcal{S}$'s and Bob's output in the ideal world, except with negligible probability in κ. This involves proving the following cases:

1. Consider all steps of the protocol except for the reconstruction phase. For all these steps we must argue that Bob outputs `abort` with the same probability in the ideal and real worlds. In the ideal world Bob outputs `abort` if and only if $\mathcal{S}$ sends `abort` to $\mathcal{F}_{\text{SFE}}$. By construction $\mathcal{S}$ outputs `abort` if and only if Bob aborts in the simulation. But the probability that Bob aborts may in general depend on his input, which is y in the real execution and 0^n in the simulation. So we must argue that the probability that Bob aborts does not depend on whether his input is y or 0^n.

2. In the cases where Bob does not abort before the reconstruction phase we must argue that, except with probability negligible in the security parameter, he does not abort during reconstruction and that his output in the real execution is $f(x', y)$ where x' is the same input for Alice that $\mathcal{S}$ inputs to $\mathcal{F}_{\text{SFE}}$ in the ideal world. Our strategy for this is the following:

   (a) We show that after a certain point in the protocol, even a malicious Alice is committed to her inputs x'^j for each of the ℓ garbled circuits.
   (b) We show how $\mathcal{S}$ can extract x'^j for all $j \in [\ell]$ from $\mathcal{A}$.
   (c) Then we show that both $\mathcal{S}$ and honest Bob will use the x'^j as input for a circuit that is correctly garbled, committed and where the verifiable commitments on the output keys of $\mathbb{H}^{\text{ln}} \oplus \mathsf{a}$ are also correctly constructed.
   (d) We then show that such a circuit exists except with negligible probability in the security parameter, if not he will abort.
   (e) Next we show that there cannot exist several good circuits where the input Alice is committed to is different.
   (f) Finally, we argue that even with a malicious Alice, Bob outputs $f(x', y)$ except with negligible probability, in particular that the real Bob will compute $f(x'^j, y)$ for the same x'^j as the simulator, if not he will abort.

Before we show each of these steps, consider the following characterization: We let the *circuit keys* for a garbled circuit be the symmetric keys corresponding to the input and output wires of that circuit which Alice commits to in the *Commitment* phase. We say that a garbled circuit is *good* if the plaintext circuit that is obtained by decrypting (degarbling/opening) the garbled circuit using its corresponding circuit keys results in the correct circuit (that is, the circuit that correctly computes the augmented circuit). A garbled circuit that decrypts to some other plaintext circuit, or that cannot be decrypted at all, using its circuit keys, is called a *bad* circuit. Note that after the *Commitment* phase all garbled circuits are either good or bad. In a similar manner we call a verifiable commitment to an output key good if it is correctly constructed and the id of the commitment reflects the key committed to. For example, the commitment with id $\text{ID}_{i,j}^{\text{out0}}$ should be made to the i'th output 0-key in the j'th garbled circuit.

Proof of Case 1 We now consider all the steps in which Bob can output `abort`:

Polynomial setup   If the opening of the polynomials, after the cut-and-choose step, is wrong or the polynomials are not of correct degree then Bob aborts. However, which polynomials are checked is completely determined by Bob's challenge and is thus independent of his input.

Cut-and-choose   Bob terminates in this phase if any of the garbled check circuits are incorrect, if Alice has not committed to the correct input keys for each of the garbled circuits, if any commitment to the output keys of the garbled circuit or $\mathbb{H}^{\text{ln}} \oplus \mathsf{a}$ is bad or if there is a problem with the verifiable commitments to the polynomial points, found using the links. All of these cases depend on the cut-and-choose challenge which is chosen by coin tossing[13] and thus independent of Bob's input.

Evaluation   If Alice does not give the correct openings to the commitments of her choice of input keys, if she committed to garbage or if any of the verifiable commitments to the output keys of $\mathbb{H}^{\text{ln}} \oplus \mathsf{a}$ are bad, then Bob will terminate. However, the computation of $\mathbb{H}^{\text{ln}} \oplus \mathsf{a}$ is only based on Alice's input and thus independent of Bob's input. Finally, Bob will terminate if he cannot evaluate any of the garbled evaluation circuits. Whether Bob terminates might be based on his input as he learns only one of the keys for each wire in each circuit in the *Oblivious Transfer* phase. However, if Alice tries to do a selective failure attack during the OTs she cannot learn anything about Bob's true input because what Bob actually inputs to the protocol is $\bar{y}$, which by construction seems uniformly random. Still, Alice can always learn one bit of $\bar{y}$ using a selective failure attack, and even more bits if she is lucky. However,

---

[13] Failure of $\mathbb{H}^{\text{Out}}$ also depends on the cut-and-choose of polynomials but we have already showed that such failure is independent of Bob's input.

as proved in [LP07] if $|\bar{y}| = \max(4n, 8s)$ this does not give Alice any information about Bob's true input except with probability 2^{-s}.

Proof of Case 2.a and 2.b We first observe that in the protocol even a malicious Alice is committed to a value associated with each of her input wires for all the $j \in [\ell]$ garbled circuits after the *Commitment* phase. The simulator can furthermore extract each of these values since Alice commits to these using $\mathcal{F}_{\text{COM}}$. Furthermore, the simulator can find the semantic values (if the circuit can be evaluated) for each of Alice's commitments by extracting Alice's "order commitments" (The *Commitment* phase, step 2). More specifically what $\mathcal{S}$ does is as follows:

- In Step 2 of the *Commitment* phase $\mathcal{S}$ learns $k_{i,j}^0 \| k_{i,j}^1 = k_{i,j}^0 \| (k_{i,j}^0 \oplus \Delta_j)$ for all $i \in [\bar{m}]$ and $j \in [\ell]$ directly by extracting them from Alice's calls to $\mathcal{F}_{\text{COM}}$.
- Then in Step 3 of the same phase $\mathcal{S}$ learns $k_{i,j}^{x_i'}$ for all $i \in [m]$ and $j \in [\ell]$.
- For each $i \in [\bar{m}]$ and $j \in [\ell]$ $\mathcal{S}$ learns the bits Alice committed to by checking if $k_{i,j}^{x_i'} = k_{i,j}^0$ or $k_{i,j}^{x_i'} = k_{i,j}^1$. If the first is true then $\bar{x}_i'^j = 0$, if the second is true then $\bar{x}_i'^j = 1$.

Proof of Case 2.c Next, notice that some of these commitments may not be actual keys to their associated garbled circuit but instead just random garbage. In that case and both $\mathcal{S}$ and Bob will discard their associated garbled circuits during the *Evaluation* phase (or terminate the protocol during *Cut-and-Choose*). So, unless the protocol has terminated, by the end of the *Evaluation* phase, at least one garbled circuit for evaluation will exist where each of the inputs committed to by Alice are actual keys for the corresponding garbled circuit. However, there might be more than one such circuit and Alice might have committed to keys with different semantic value in each of these. Still, if that is the case then the augmented output of $\mathbb{H}^{\text{In}} \oplus \mathsf{a}$ will be different for each of the circuits where she has given different inputs except with probability 2^{-s}.

To see this we must show that it will not be possible for Alice to find a collision on $\mathbb{H}^{\text{In}} \oplus \mathsf{a}$ with probability larger than 2^{-s}. That is that it is not possible for her to give two inputs $x\|a, x'\|a' \in \{0,1\}^{\bar{m}}$ with $x \neq x'$ such that $(M^{\text{In}}x) \oplus a = (M^{\text{In}}x') \oplus a'$ with probability larger than 2^{-s}. Now remember that we have from Lemma 1 that $\mathbb{H}^{\text{In}}$ computes a universal hash function with the uniform difference property, i.e., that $\Pr[M^{\text{In}}x \oplus M^{\text{In}}x' = \bar{a}] = 2^{-s}$ for any $\bar{a} \in \{0,1\}^s$. It turns out that the uniform difference property is enough to ensure Alice cannot find a collision. To see this notice that what we actually need to show is $\Pr[(M^{\text{In}}x) \oplus a = (M^{\text{In}}x') \oplus a'] \leq 2^{-s}$ for all $x\|a, x'\|a' \in \{0,1\}^{\bar{m}}$ with $x \neq x'$ when M^{In} is uniformly sampled. However, setting $\bar{a} = a \oplus a'$ we see

$$M^{\text{In}}x \oplus M^{\text{In}}x' = \bar{a}$$
$$M^{\text{In}}x \oplus M^{\text{In}}x' = a \oplus a'$$
$$(M^{\text{In}}x) \oplus a = (M^{\text{In}}x') \oplus a' ,$$

and thus we have $\Pr[(M^{\text{In}}x) \oplus a = (M^{\text{In}}x') \oplus a'] = 2^{-s}$ and we are done.

Next see that both Bob and the simulator will see if the evaluation of $\mathbb{H}^{\text{In}} \oplus \mathsf{a}$ is correct for the circuits in ϕ during the *Evaluation* phase and abort if that is not the case. However, Alice could commit to different inputs without being detected by corrupting the verifiable commitments of the output keys of $\mathbb{H}^{\text{In}} \oplus \mathsf{a}$. Still, we force Alice to commit to the output keys before the *Cut-and-Choose* phase, so if she chooses to cheat we then have from Lemma 4 that she cannot successfully do this for all evaluation circuits and thus there will be at least one circuit for which the verifiable commitments to the output keys of $\mathbb{H}^{\text{In}} \oplus \mathsf{a}$ are correctly constructed. Furthermore, one of these must be associated with some correct input keys as the protocol would otherwise terminate during *Evaluation*.

We have now established that there must be a vector of input keys from Alice along with good verifiable commitments to the output key of $\mathbb{H}^{\text{In}} \oplus \mathsf{a}$. However, there might still be instances of input keys that fit their associated circuit but where $\mathbb{H}^{\text{In}} \oplus \mathsf{a}$ is incorrectly constructed, e.g., so that it gives the same digest as the correctly constructed case. In this situation we must now argue that both Bob and $\mathcal{S}$ can find out which of the garbled circuits and its corresponding $\mathbb{H}^{\text{In}} \oplus \mathsf{a}$ augmentation is correct. In particular, we have two cases:

- Either the semantic output of the garbled circuit is the same in both cases and there is no issue for Bob since the correct output is known in the real world and Bob can output this. In the simulation $\mathcal{S}$ can easily find out which

circuit is correct since it can extract the semantic value of Alice's input keys as described above in the proof of 2.a and 2.b. The input for a correct circuit whose associated $\mathbb{H}^{\mathsf{In}} \oplus \mathsf{a}$ augmentation is correct $\mathcal{S}$ inputs to $\mathcal{F}_{\mathsf{SFE}}$.

- If the semantic output of the garbled circuits are different the simulator can do the same as in the case above or it can do the same as a real-world Bob would; use the *Reconstruction* phase to extract x'^j for all $j \in \phi$ committed to by Alice and evaluate $\mathbb{H}^{\mathsf{In}} \oplus \mathsf{a}$ of the disputed circuits in plain and find out which input is used in a correctly constructed circuit and the use this input to evaluate the function f in plain.

Proof of 2.d The fact that at least one good circuit exists or both $\mathcal{S}$ and Bob will terminate the protocol follows directly from Lemma 4.

Proof of 2.e The fact that we cannot have several correctly constructed augmented garbled circuits evaluate to different outputs stems from the fact that if they evaluate differently then Alice's input must be different. However, if that is the case then $\mathbb{H}^{\mathsf{In}} \oplus \mathsf{a}$, except with probability 2^{-s}, will output different values and thus both Bob and $\mathcal{S}$ will terminate the protocol by Lemma 1 and the discussion of case 2.c.

Proof of Case 2.f We finally argue that even with a malicious Alice, if x' is the semantic value of the input that Alice commits to in the *Commitment* phase, associated with a correctly constructed augmented garbled circuit, then $f(x', y)$ will be the output of Bob if he does not abort.

To see this we only need to argue that in the real world, and despite a corrupted Alice, Bob outputs $f(x', y)$ where x' is the input that Alice commits to in the *Commitment* Step. To show this we consider the following lemma:

Lemma 5. *Assume we have c evaluation circuits and w output wires where at most half are bad, then the probability that we cannot find the Δ for each circuit, and thus might fail the reconstruction, is at most*

$$1 + (w/2)(2^c - 1) + \sum_{i=2}^{w/2} \left(\binom{w/2}{i} \cdot c \cdot 2^{i(c-1)} \right) \ .$$

Proof. See Appendix G. $\qquad\square$

According to Lemma 5 we can bound the probability that a corrupted Alice passes the cut-and-choose test of polynomials and that we cannot restore all the Δ values for the evaluable circuits. When we use the following lemma we get that the amount of wires needed in order to make sure that we can find Δ_j for all $j \in \phi$ except with probability at most 2^{-s}:

Lemma 6. *Assuming the polynomials on at least half of the wires are good then we achieve failure probability at most 2^{-s} for $s \geq 1$ if we have at least $w = \lceil 4.82s + 4.82 \rceil \geq -\frac{s}{\frac{\log(3)}{2} - 1} - \frac{1}{\frac{\log(3)}{2} - 1}$ output wires.*

Proof. See Appendix G. $\qquad\square$

Next, to bound the amount of polynomials needed in the cut-and-choose challenge remember that we need at least half of the $\lceil 4.82s + 4.82 \rceil$ evaluation polynomials to be good. Thus there is a total of $\binom{s_2}{s_2 - (\lceil 4.82s + 4.82 \rceil)}$ possible ways to choose the check polynomials. If more than half of the evaluation polynomials are "bad" then there is a total of $\binom{s_2 - \left(\frac{\lceil 4.82s + 4.82 \rceil}{2} + 1 \right)}{s_2 - (\lceil 4.82s + 4.82 \rceil)}$ "bad" cut-and-choose challenges. Since there is only a set of $s_2 - \left(\frac{\lceil 4.82s + 4.82 \rceil}{2} + 1 \right)$ good polynomials to choose from and the check size remains $s_2 - (\lceil 4.82s + 4.82 \rceil)$. Thus the probability of a "bad" cut-and-choose challenge is

$$\frac{\binom{s_2 - \left(\frac{\lceil 4.82s + 4.82 \rceil}{2} + 1 \right)}{s_2 - (\lceil 4.82s + 4.82 \rceil)}}{\binom{s_2}{s_2 - (\lceil 4.82s + 4.82 \rceil)}}$$

We must now find a value for s_2 such that this expression is less than 2^{-s}. According to the following lemma this is the case if we set $s_2 \geq 6s + 7$.

Lemma 7.

$$\frac{\binom{6s+7-\left(\frac{\lceil 4.82s+4.82 \rceil}{2}+1\right)}{6s+7-(\lceil 4.82s+4.82 \rceil)}}{\binom{6s+7}{6s+7-(\lceil 4.82s+4.82 \rceil)}} \leq 2^{-s}$$

for $s > 1$.

Proof. See Appendix G. □

Next see that at least half of the polynomials used to augmented the output wires of $\mathbb{H}^{\mathsf{Out}}$ are good (the event in Lemma 6) then we have exactly the case described in the *Reconstruction* phase except if some of the links of $\mathbb{H}^{\mathsf{Out}}$ are wrong. However, since Alice sends the links before the cut-and-choose phase then Lemma 4 tells us that there will be at least one circuit where the links are correctly constructed. In turn Bob will always be able to recover the input x' used by $\mathcal{A}$.

Whether a garbled circuit is good or bad is determined before partitioning them into check and evaluation circuits in *Cut-and-Choose*. Thus when using a replication factor of ℓ (which is the replication factor actually used in the protocol), the random partitioning into check and evaluation circuits will ensure that there will be no good evaluation circuit with probability at most $2^{-s} \leq 2^{-\left(\ell - \frac{1}{2}\log \ell + \log\left(2\frac{\sqrt{2\pi}}{e^2}\right)\right)}$ (by Lemma 4).

This completes the proof for the case where Alice is corrupt.

E.2 Corrupt Bob

Lemma 8. *Let π be the protocol in Fig. 14. Assume that the garbling scheme $\mathcal{G} = (\mathsf{CYao}, \mathsf{CEval})$ used in π is IND-CIA secure according to Def. 2 and the link scheme (G, F) used is secure according to Def. 3 then π securely realizes $\mathcal{F}_{\mathrm{SFE}}$ in Fig. 7 in the ($\mathcal{F}_{\mathrm{BOTCC}}$, $\mathcal{F}_{\mathrm{CT}}$, $\mathcal{F}_{\mathrm{COM}}$)-hybrid model with abort in the presence of static, malicious adversary corrupting Bob.*

The intuition for this part of the proof is quite similar to most other two-party protocols based on garbled circuits and cut-and-choose [LP07, LP11]. Plainly speaking a malicious Bob cannot cheat since all the input he gives to the protocol is his plain input bits (to $\mathcal{F}_{\mathrm{COM}}$) and some random choices, which are used only to protect himself against a potentially malicious Alice.

Still, from a simulation point of view, we must argue that the simulated $\mathcal{B}$ outputs the same as $\mathcal{B}$ in the real (hybrid) execution, including the literal output of the protocol $f(x, y') = z$. This is not as obvious as in the case of a corrupted Alice and is achieved by $\mathcal{S}$ learning the cut-and-choose challenge directly from $\mathcal{F}_{\mathrm{CT}}$ and then constructing the evaluation garbled circuits "maliciously" such that they always output $f(x, y') = z$ no matter what the semantic value of the input is, in particular it will still give the correct output even if the x' used in the garbled circuit is the all-0 string.

The formal proof proceeds as follows: Given an adversary $\mathcal{B}$ that controls Bob, we construct a simulator $\mathcal{S}$ in a way analogous to the case where Alice is corrupted. Formally the simulator $\mathcal{S}$ works as follows:

- $\mathcal{S}$ starts simulating internally the real world execution of the protocol consisting of Alice, $\mathcal{B}$ and the ideal functionalities $\mathcal{F}_{\mathrm{BOTCC}}^{\bar{n},\ell,\kappa}$, $\mathcal{F}_{\mathrm{CT}}$ and $\mathcal{F}_{\mathrm{COM}}$. If Alice aborts at any time during the simulation, then $\mathcal{S}$ sends abort to $\mathcal{F}_{\mathrm{SFE}}$.
- The simulator then executes the *Setup* and *Polynomial Setup* phases as an honest Alice would.
- When reaching the *Oblivious Transfer* $\mathcal{S}$ first constructs the garbled circuits as an honest Alice would. It then extracts $\mathcal{B}$'s input $\bar{y}'$ from his call to $\mathcal{F}_{\mathrm{BOTCC}}^{\bar{n},\ell,\kappa}$. $\mathcal{S}$ then computes $M^{\mathsf{Sec}}\bar{y}' = y'$ and sends y' to $\mathcal{F}_{\mathrm{SFE}}$, which returns $z = f(x, y')$. Simulating $\mathcal{F}_{\mathrm{BOTCC}}^{\bar{n},\ell,\kappa}$, $\mathcal{S}$ returns the the keys $k_{\bar{m}+i,j}^{\bar{y}_i'}$ for $i \in [\bar{n}]$ and $j \in [\ell]$ to $\mathcal{B}$.
- In the *Commitment* phase $\mathcal{S}$ sends the garbled circuits it just constructed to $\mathcal{B}$. Then, $\mathcal{S}$ simulates $\mathcal{F}_{\mathrm{CT}}$ to select the partition of check and verification circuits. For all the check circuits it constructs the verifiable commitments of the output keys correctly. For the evaluation circuits it first sets Alice's input to $x' = 0^m \| a$ with $a \in_R \{0, 1\}^s$ and evaluates $z' = f(0^m, y')$ in plain. It then constructs the verifiable commitments to each bit i correctly if

$z_i' = z_i$, otherwise it switches the meaning of the 0- and 1-key, i.e. it does $\mathcal{F}_{\text{COM}}\left(\text{VC}, \text{ID}_{i,j}^{\text{out0}}, k_{\bar{m}+\bar{n}+|C|+i,j}^0 \oplus \Delta_j\right)$ and $\mathcal{F}_{\text{COM}}\left(\text{VC}, \text{ID}_{i,j}^{\text{out1}}, k_{\bar{m}+\bar{n}+|C|+i,j}^0\right)$. Thus when $\mathcal{B}$ evaluates the evaluation circuits and learns the semantic meaning of the output keys they will always be $z = f(x, y')$. The rest of the commitments in the *Commitment* phase $\mathcal{S}$ constructs as an honest Alice would.

- In the *Augmentation* phase $\mathcal{S}$ receives the specification for $\mathbb{H}^{\text{In}}$ and $\mathbb{H}^{\text{Out}}$ from Bob. For the check circuits $\mathcal{S}$ then constructs the verifiable commitments to the output keys of $\mathbb{H}^{\text{In}} \oplus a$ along with the links, as an honest Alice would. For the evaluation circuits $\mathcal{S}$ decides on a uniformly random string of s bits and hardcodes the $\mathbb{H}^{\text{In}} \oplus a$ part of the evaluation circuits to compute this. That is, it constructs the verifiable commitments to the output keys of $\mathbb{H}^{\text{In}} \oplus a$ such that they hit the randomly chosen string when evaluated on $0^m \| a$. For $\mathbb{H}^{\text{Out}}$ $\mathcal{S}$ uses the output z received from $\mathcal{F}_{\text{SFE}}$ and constructs the verifiable commitments as an honest Alice would.
- $\mathcal{S}$ now proceeds with the *Cut-and-Choose*, *Evaluation* and *Reconstruction* phases as an honest Alice would.

To prove security we must show the following:

1. We must argue that Alice outputs `abort` with the same probability in the real and ideal world. In particular that the probability with which she aborts does not depend on whether or not her input is $x\|a$ or $0^m\|a'$ with $a, a' \in_R \{0,1\}^s$.
2. If Alice does not abort then we must argue that, except with negligible probability in s, Alice outputs the same in the real and ideal world.
3. We need to argue that the output of $\mathcal{B}$ is the same in the real and ideal world. In particular that $\mathcal{B}$ does not learn anything that depends on Alice's input x except $f(x, y') = z$.

Together, the steps avove ensure that the *joint* distributions of $\mathcal{B}$'s output and Alice's output in the ideal and in the real world are indistinguishable.

Proof of Case 1 Notice that an honest Alice following the protocol will not abort during the protocol.

Proof of Case 2 This case is quite easy to argue about since Alice does not receive any "real" output. What she receives is perhaps a message saying `abort`. However, since Alice is honest this will happen with the same probability in both the real and ideal world.

Proof of Case 3 This is the hardest case and we must first argue that $\mathcal{S}$ can convince $\mathcal{B}$ to learn the true value $z = f(x, y')$ with the same probability in the real and ideal world. To do this we first observe that in the real protocol even a corrupt Bob is committed to a specific input $\bar{y}'$ at the time he sends his input to $\mathcal{F}_{\text{BOTCC}}^{\bar{n},\ell,\kappa}$. Since $\mathcal{S}$ fully controls the simulation of $\mathcal{F}_{\text{BOTCC}}^{\bar{n},\ell,\kappa}$, $\mathcal{S}$ can therefore easily extract $\bar{y}'$ from Bob's communication with $\mathcal{F}_{\text{BOTCC}}^{\bar{n},\ell,\kappa}$, compute $y' = M^{\text{Sec}}\bar{y}'$ and input y' to $\mathcal{F}_{\text{SFE}}$ (as done in the simulation above). Thus, in the real world, $\mathcal{B}$ learns $f(x, y') = z$ from $\mathcal{F}_{\text{SFE}}$. This is the same as in the ideal world since $\mathcal{S}$ hardcodes the evaluation garbled circuits to be exactly $f(x, y') = z$.

Furthermore, $\mathcal{B}$ cannot learn anything about the semantic values of Alice's input keys (or any other wire keys in the circuit) for the evaluation circuits because of the IND-CIA security of the garbling scheme. Notice that IND-CIA is needed as we use a simulated input for Alice in the simulation.

More formally we reduce a distinguishability advantage of the view of the execution of the protocol in either the real or ideal (simulated) world to an advantage in winning the IND-CIA game. Specifically we construct a distinguisher $\mathcal{D}$ in the IND-CIA game which makes use of $\mathcal{B}$ and an environment that tries to distinguish executions of the protocol in the real or ideal (simulated) world. We will construct $\mathcal{D}$, such that depending on the value of the challenge bit c in IND-CIA, $\mathcal{D}$ will produce a transcript of the protocol that is indistinguishable from the real, respectively $\mathcal{S}$'s simulation of the ideal world.

We have $\mathcal{D}$ make the environment execute the real, respectively simulated ideal protocol, up to the *Oblivious Transfer* phase when $\mathcal{B}$ inputs $\bar{y}'$. At this point $\mathcal{D}$ sets the bit vector $x^0 = \bar{x}\|\bar{y}'$ and $x^1 = 0^m\|a\|\bar{y}'$ where a is a uniformly random vector of s bits. It then initializes $\ell/2$ executions of the IND-CIA game by asking for $\ell/2$ independent garbled versions of the circuit C and at each instance it queries the game (x^0, x^1).

Next, $\mathcal{D}$ embeds these circuits in the protocol as evaluation circuits. This means that $\mathcal{B}$ will receive garbled circuits constructed as an honest Alice would. The keys for the evaluation circuits he receives are constructed by the IND-CIA game. These will again be made as in the real world if $c = 0$ and as in the ideal, simulated world, if $c = 1$.

Next we let $\mathcal{D}$ evaluate the garbled circuits received by the IND-CIA games using the input keys it also got from the IND-CIA games. Then, it constructs verifiable commitments on the output keys so that they hit their corresponding bits in $z = f(x, y)$. That is, if the output key on the i'th wire in the j'th evaluation circuit is $k_{i,j}$ and $z_i = 0$ it does $\mathcal{F}_{\mathrm{COM}}\left(\mathrm{vc}, \mathsf{ID}_{i,j}^{\mathsf{out0}}, k_{i,j}\right)$, otherwise it does $\mathcal{F}_{\mathrm{COM}}\left(\mathrm{vc}, \mathsf{ID}_{i,j}^{\mathsf{out1}}, k_{i,j}\right)$. It then simulates the other verifiable commitment by calling $\mathcal{F}_{\mathrm{COM}}\left(\mathrm{vc}, \mathsf{ID}_{i,j}^{\mathsf{out1}}, r_{i,j}\right)$ if $z_i = 0$ and $\mathcal{F}_{\mathrm{COM}}\left(\mathrm{vc}, \mathsf{ID}_{i,j}^{\mathsf{out0}}, r_{i,j}\right)$ otherwise, where $r_{i,j}$ is a uniformly random sampled κ bit string. In a similar manner it simulates the order commitments by just committing to uniformly random sampled strings.

Since we are in the hybrid model for commitments this goes unnoticed until the commitment functionality is queried on one of the keys $\mathcal{D}$ did not learn, so if $\mathcal{B}$ could distinguish between the random commitments and the true commitments, then it is because it can query the functionality for commitment on the other key it did not learn. But then we have both keys for one of the wires and can compute the value Δ for that wire, which allows $\mathcal{D}$ to win the IND-CIA game. Specifically, for all queries x to the ideal functionality for verifiable commitment done by $\mathcal{B}$, if $x \in \{0,1\}^\kappa$ $\mathcal{D}$ will query $\mathcal{O}_\Delta(x \oplus k)$ in each IND-CIA game for all output keys k which it learned from the keys and circuits returned by the game. If instead $x \in \{0,1\}^{2\kappa}$ then $\mathcal{D}$ will query $\mathcal{O}_\Delta(x' \oplus x'')$ in each IND-CIA game, where x' is the first κ bits of x and x'' the last κ bits of x.

Another subtlety in the embedding done in the above reduction is that $\mathcal{D}$ needs to give the links $L_{i,j} = G\left(P_{\chi[i]}(j), k_{i,j}'^0\right)$ to $\mathcal{B}$. This is only possible for the wires where $\mathcal{D}$ has learned the 0 key. However, for all the wires where $\mathcal{D}$ learns only the 1 key, it holds that if it did not give the links to $\mathcal{B}$, then it would be hard for $\mathcal{B}$ to compute any point on the corresponding polynomial or the 0 key: The points on the polynomials have high min-entropy in the view of $\mathcal{B}$, as he knowns only t points, and computing a 0 key would be the same as breaking the security of the garbling scheme. Hence, if we consider the distribution $(x_1, y_1, \ldots, x_\ell, y_\ell, z)$ generated by running the protocol and letting z be the view of $\mathcal{B}$ with the links removed which correspond to wires for which he learned the 1 key and if $\mathcal{D}$ let (x_i, y_i) be the corresponding $\left(P_{\chi[i]}(j), k_{i,j}'^0\right)$, then this distribution is hard. Hence $\mathcal{D}$ can replace the links $L_{i,j} = G\left(P_{\chi[i]}(j), k_{i,j}'^0\right)$ by $G(a, b)$ for uniformly random a and b. Note that it only needs to do this step twice. First $\mathcal{D}$ replaces all the involved links by random links, then it can do the hybrids where we circuit by circuit replace input a with input 0 in the circuits, and then we can go back and replace the random links by correct links.

The arguments above imply that if it is possible for the environment to distinguish between a real world execution and a simulated execution, then this translates directly into an advantage for $\mathcal{D}$ in the IND-CIA game.

Finally, the reason a PPT $\mathcal{B}$ cannot learn anything from the output of $\mathbb{H}^{\mathsf{In}} \oplus \mathsf{a}$, which will be different in the real and ideal world, is because the output is encrypted under a one-time pad, i.e., the random bits a chosen by $\mathcal{S}$ are xor'ed with $M^{\mathsf{In}}x$. So both in the real and ideal world this value will be uniformly distributed.

Hence, the above series of hybrids take us from the real execution to the simulation.

$\square$

A Note On the Statistical Security Parameter. In the proof above we use that each of the bad events that a cheating Alice can try to cause can only happen with probability at 2^{-s}. These events are:

1. A majority of inconsistent polynomials remains after cut-and-choose of polynomials.
   – This can only happen if she incorrectly constructs at least $\lceil 2.41s + 3.41 \rceil$ of the polynomials.
2. *Every* check circuit is good and *every* evaluation circuit is bad.
   – This only happens if she constructs half of the circuits correctly and the other half incorrectly.
3. There is a collision of $\mathbb{H}^{\mathsf{In}} \oplus \mathsf{a}$.
   – This can only happen if she gives at least two different inputs to the protocol.[14]
4. Alice is successful in a selective failure attack.
   – This can only happen if she gives garbage for some inputs to the $\mathcal{F}_{\mathrm{BOTCC}}$.

---

[14] There can seemingly be a collision if Alice maliciously constructs the verifiable commitments to the output keys of $\mathbb{H}^{\mathsf{In}} \oplus \mathsf{a}$ for one of the circuits where she has given divergent input. However, this is not actually a collision of $\mathbb{H}^{\mathsf{In}} \oplus \mathsf{a}$ but an incorrect circuit.

5. There is a collision of $\mathbb{H}^{\text{Out}}$.
   - This can only happen if she constructs at least one circuit maliciously such that it computes something else than $f(x, y)$.

However, nothing is preventing Alice from trying to cause each of these events. Still, it should be noted that if Alice is unsuccessful in her cheating attempt of event 1, 2 or 3 then Bob will notice and can safely abort the protocol. Thus trying to cause more than just a single of the events 1, 2 or 3 will only increase the probability of Alice being caught and Bob terminating the protocol without the compromising security. Now notice that list of bad events are in the order of which Bob will discover if they are unsuccessful (and in turn for the three first events, abort). Thus, to increase the probability of a successful cheating attempt Alice should try to make event 4 and 5 occur. This is so since if there is a collision of $\mathbb{H}^{\text{Out}}$ Bob cannot terminate the protocol without possibly leaking some information on his private input and since a selective failure attack does not necessarily make Bob abort the protocol. She cannot combine an attack for event 4 or 5 with attack 1, 2 or 3 since if one of these are unsuccessful then Bob will terminate before event 4 and 5 could happen.

Now since event 4 and 5 can happen independently with probability 2^{-s} we get from the union bound that if Alice tries to make both occur she will be successful in having at least one of them occur with probability at most $2^{-s} + 2^{-s} = 2^{-s+1}$. This means that in order to achieve a total failure probability of at most 2^{-s} we must instead of s use $s' = s + 1$ in the part of the protocol concerned with $\mathbb{H}^{\text{Out}}$ and the selective failure attack prevention. Thus M^{Sec} must be a $n \times \max(4n, 8(s+1))$ matrix and $\mathbb{H}^{\text{Out}}$ must have $\lceil 4.82(s+1) + 4.82 \rceil$ output wires (in turn we must also construct $6(s+1) + 7$ polynomials to be certain that the success probability of guessing the cut-and-choose challenge of polynomials is still at most 2^{-s} unless we wish to reprove Lemma 7).

UC Security. We note the following about the simulator in the proof:

1. The simulator does not at any time rewind the adversary it simulates.
2. The simulator can extract the modified input of the malicious player in its simulation (and hence obtain the honest players output from $\mathcal{F}_{\text{SFE}}$).

Therefore, we immediately achieve UC security [Can01] if the underlying OT protocol used for the OTs and coin tossing protocol are UC-secure.

For example, assuming that the discrete log problem (DLP) is hard, the UC-secure oblivious transfer of [PVW08] can be used to realize the seed OTs.

Theorem 2 is stated in the hybrid world where the protocol depends on ideal functionalities for commitments, coin tossing and oblivious transfers with consistent choice. In order to achieve a protocol that can be implemented in practice we apply the UC composition theorem, as stated in the following corollary.

Corollary 1. *Let* (CYao, CEval) *be the concrete Yao circuit garbling scheme of Fig. 5 and let the links be implemented as discussed in Appendix B. Let π be the protocol in Fig. 14 to Fig. 16, but where the call to $\mathcal{F}_{\text{BOTCC}}$ is replaced by invocation of the protocol π_{BOT} in Fig. 11 (in which calls to $\mathcal{F}_{\text{OT}}$ are replaced by an execution of the OT protocol in [PVW08]) and where calls to $\mathcal{F}_{\text{COM}}$ are replaced with invocations of the protocol π_{COM} from Appendix C and where the calls to $\mathcal{F}_{\text{CT}}$ are replaced by invocations of the protocol π_{CT} from Appendix C and where calls to $H(\cdot)$ in the subprotocols above are done with calls to the random oracle. Then in the random oracle model π UC-securely realizes $\mathcal{F}_{\text{SFE}}$ with abort in the presence of static, malicious adversaries assuming DLP is hard and both parties have access to a common reference string.*

Proof. The (CYao, CEval) scheme in Fig. 5 is shown to be IND-CIA secure in the random oracle model in [PSSW09]. As proved in Appendix C the protocol π_{COM} UC-securely realizes $\mathcal{F}_{\text{COM}}$ in the random oracle model. Next, consider that $\mathcal{F}_{\text{BOTCC}}$ can be implemented by an invocation of π_{BOT} along with a random oracle and a protocol for $\mathcal{F}_{\text{OT}}$ as described in Appendix C, where $\mathcal{F}_{\text{OT}}$ can be realized as in [PVW08] using a common reference string and the assumption that DLP is hard. Then $\mathcal{F}_{\text{CT}}$ can be realized as described in Appendix C. Finally, all of the above realizations are UC-secure in the presence of static and malicious adversaries. Hence, by Theorem 2 and the UC-composition theorem Cor. 1 follows immediately from Theorem 2. $\qquad\square$

F The Replication Factor

Using Stirling's approximation formula we find a lower bound for the circuit replication factor ℓ required for a malicious Alice being able to guess the cut-and-choose challenge with probability at most 2^{-s}. We restate Lemma 4 here:

The probability of guessing exactly the elements in a set of size $\ell/2$ sampled randomly from a set of ℓ elements is less than 2^{-s} when $\log\left(2\frac{\sqrt{2\pi}}{e^2}\right) - \frac{1}{2}\log(\ell) + \ell \geq s$.

Proof (of Lemma 4). First notice that a cheating Alice can only succeed if she guesses exactly the set of check circuits in the cut-and-choose phase. For circuit replication factor ℓ there are exactly $\binom{\ell}{\ell/2}$ possible sets, of which one is randomly selected based on the coin tossing. We now bound $\binom{\ell}{\ell/2}$ using Stirling's approximation, assuming $\ell \geq 2$:

$$\binom{\ell}{\frac{\ell}{2}} = \frac{\ell!}{\frac{\ell}{2}!\left(\ell - \frac{\ell}{2}\right)!} = \frac{\ell!}{\left(\frac{\ell}{2}!\right)^2} \geq \frac{\sqrt{2\pi\ell}\left(\frac{\ell}{e}\right)^\ell}{\left(e\sqrt{\frac{\ell}{2}}\left(\frac{\frac{\ell}{2}}{e}\right)^{\frac{\ell}{2}}\right)^2} = \frac{\sqrt{2\pi\ell}\left(\frac{\ell}{e}\right)^\ell}{e^2\frac{\ell}{2}\left(\frac{\ell}{2e}\right)^\ell} = \frac{\sqrt{2\pi\ell}\ell^\ell e^{-\ell}}{e^2 2^{-1}\ell\ell^\ell 2^{-\ell}e^{-\ell}}$$

$$= \frac{2\sqrt{2\pi\ell}2^\ell}{e^2\ell} = \frac{2\sqrt{2\pi}}{e^2}\cdot\frac{\sqrt{\ell}2^\ell}{\ell} = \frac{2\sqrt{2\pi}}{e^2}\cdot\ell^{-\frac{1}{2}}2^\ell$$

We now want to find ℓ such that $\frac{2\sqrt{2\pi}}{e^2}\cdot\ell^{-\frac{1}{2}}2^\ell \geq 2^s$. This will mean that there will be at least 2^s possible choices for the cut-and-choose phase and thus Alice's probability of guessing the choice will be at most 2^{-s}. Setting $c = \frac{2\sqrt{2\pi}}{e^2}$, taking the logarithm, and isolating s we get:

$$\log\left(c\ell^{-\frac{1}{2}}2^\ell\right) \geq \log\left(2^s\right)$$

$$\log(c) - \frac{1}{2}\log(\ell) + \ell \geq s$$

$\square$

Approximating the constant we get:

$$\ell - \frac{1}{2}\log(\ell) - 0.5596 \geq s$$

G Polynomial Cut-and-Choose

We restate Lemma 5 here:

Assume that we have c evaluation circuits and w output wires and at most half of the wires have an associated polynomial that is bad (of degree greater than t), then the probability that we cannot find the Δ for each circuit, and thus might fail the reconstruction, is at most

$$\frac{1 + (w/2)(2^c - 1) + \sum_{i=2}^{w/2}\left(\binom{w/2}{i}\cdot c\cdot 2^{i(c-1)}\right)}{2^{cw}}.$$

Proof (of Lemma 5). In the following denote the bit value on the i'th wires in the j'th circuit by $b_{i,j}$. We now prove this lemma using a counting argument. First notice that the total amount of cases of c circuits with w wires are $2^{c\cdot w}$ since we can view it as a set of cw binary wires where each outcome is possible. We now count the number of "bad" configurations of these wires, i.e., configurations of the polynomial points of the wires in the circuits that does not make it possible to do interpolations such that we can learn the Δ used in each of the c circuits. First define a "very bad" wire to be a wire whose associated polynomial is not correctly constructed. Any very bad wire will be completely

useless, no matter how its associated bit is set. Furthermore, the position of a very bad wire does not matter as all wires are equally usable in the manner that they can be used to restore Δ_j for all $j \in [c]$. Next define an "AND0 wire" to be a wire where the bits from all c circuits AND'ed together gives 0, assuming it is not a very bad wire. Notice that for a given wire on c circuits there is only one configuration where the bits AND'ed together to 1, that is for the all 1-string.

Notice that in the view of finding Δ_j for all $j \in [c]$ the positions of very bad wires does not matter, as any wire that is not very bad can be used to do polynomial interpolation and thus find Δ values. Next see that because of the potential existence of very bad wires the worst cases will always be when we have $w/2$ very bad wires. So we can simply consider bad cases on the remaining $w/2$ wires. The reason being if we can restore the Δ values with only $w/2$ wires then we can always restore it with more wires, no matter the configuration of the bits on the extra wires, as we can simply choose to ignore them in the restoration process. So in the following we only count configurations of the $w/2$ wires which are not very bad. Now see that all bad configurations are captured in the following, possibly intersecting, sets:

1. All configurations when we do not have any AND0 wires.
2. All configurations when we only have a single AND0 wire.
3. For all $i \in [w/2]\backslash 1$, all configurations when we have i AND0 wires and $\exists j \in [c] \; : \; \vee_{k=1}^{i} b_{k,j} = 0$, that is, the cases where the OR of the bits on each of the i wires in any of the c circuits results in 0.

If the first case happens then we cannot do any polynomial interpolation and thus must be included. However, this case only occurs when the values on all the $w/2$ wires are 1's, i.e., only on one configuration.

For the second case assume the AND0 wire is wire i, then if this case occurs we will only be able to do interpolation on wire i, and in turn only be able to find Δ_j for each circuit j when $b_{i,j} = 1$ (which cannot occur for all circuits since wire i must have at least one 0-bit). Now see that a wire is an AND0 if and only if it is not the 1-string. Thus there are $2^c - 1$ possible ways a given wire can be of the AND0 type. Furthermore, notice that there are $w/2$ ways of selecting an AND0 wire. Finally see that when we only have a single AND0 wire, the values on the remaining wires must all be the 1-string. So we get a total of $(w/2)(2^c - 1)$ bad configurations in this case.

The third case is a bit harder to show and we consider it in an inductive manner. I.e., first consider the case of two AND0 wires, w.l.o.g. wire 1 and 2, and the event that $\exists j \in [c] \; : \; b_{1,j} \vee b_{2,j} = 0$. This means that for circuit j we will only be able to learn 0 keys on the two wires we a able to interpolate (the AND0 wires) and thus not be able to find Δ_j. Now to count configurations see that bits only OR to 0 if all of them are 0, so it must be the case that $b_{1,j} = b_{2,j} = 0$. Now since each AND0 wire has c bits, we see that the amount of remaining configurations of these two wires are at most[15] $2^{2(c-1)}$, since each wire constitute 2^{c-1} possible configurations (once the j'th bit has been set to 0). However, this occurs no matter which value j takes, i.e., it happens once for each of the c possible values j can take. That is at most $c2^{2(c-1)}$ cases. Finally see that we can select two AND0 wires in $\binom{w/2}{2}$ different ways, so we get $\binom{w/2}{2}c2^{2(c-1)}$ cases in total.

Now consider the general case where we have i AND0 wires. There must be no circuit where the bits of these i wires OR to 0. When they OR to 0 the remaining configuration of the wires are bad, i.e., at most $2^{i(c-1)}$ configurations. Again it is bad no matter for what circuit the wires OR to 0, so in total $c \cdot 2^{i(c-1)}$ cases. Also, these i AND0 wires can be chosen in $\binom{w/2}{i}$ ways. So in total $\binom{w/2}{i} \cdot c \cdot 2^{i(c-1)}$ bad configurations.

Finally we add all the cases together and get an upper bound on bad cases:

$$1 + (w/2)(2^c - 1) + \sum_{i=2}^{w/2}\left(\binom{w/2}{i} \cdot c \cdot 2^{i(c-1)}\right) \; .$$

Dividing this number with the total amount of possible cases (2^{cw}) yields the desired result.

$\square$

Notice that this analysis is not tight as our argumentation only considers a subset of all good cases. That is, there are several bad cases that get counted several times.

---

[15] We count some redundant configurations here.

We restate Lemma 6 here:

Assuming the polynomials on at least half of the wires are good then we achieve failure probability at most 2^{-s} for $s \geq 1$ if we have at least $w = \lceil 4.82s + 4.82 \rceil \geq -\frac{s}{\frac{\log(3)}{2}-1} - \frac{1}{\frac{\log(3)}{2}-1}$ output wires.

Proof. First observe that the optimal strategy for Alice, if she wishes to succeed in cheating, is to construct only two circuits that evaluate differently (assuming that the amount of circuits is less than the amount of output wires). This is so since the equation $\frac{1+(w/2)(2^c-1)+\sum_{i=2}^{w/2}\left(\binom{w/2}{i}\cdot c\cdot 2^{i(c-1)}\right)}{2^{c\cdot w/2}}$ is maximized for $c = 2$ under the constraint that $1 < c < w$, $c \in \mathbb{Z}$ and $w \geq 8$. Notice that we need $c > 1$ since otherwise there is no chance of failure as the output will be the same in all circuits. Obviously we also need $c \in \mathbb{Z}$. Finally notice that $c < w$ and $w \geq 8$ are constraints arriving as artifacts of the equation.

Under these constraints consider the probability of the worst case, i.e. that the output is discrepant on 2 circuits:

$$
\frac{1+(w/2)(2^2-1)+\sum_{i=2}^{w/2}\left(\binom{w/2}{i}\cdot 2\cdot 2^{i(2-1)}\right)}{2^{2\cdot w/2}} = \frac{1+3(w/2)+\sum_{i=2}^{w/2}\left(\binom{w/2}{i}\cdot 2\cdot 2^i\right)}{2^w}
$$

$$
= \frac{1+3(w/2)+2\cdot\sum_{i=2}^{w/2}\left(\binom{w/2}{i}\cdot 2^i\right)}{2^w}
$$

$$
= \frac{1+3(w/2)+2\cdot\left(\sum_{i=0}^{w/2}\left(\binom{w/2}{i}\cdot 2^i\right)-w-1\right)}{2^w}
$$

$$
= \frac{1+3(w/2)+2\cdot\left((1+2)^{w/2}-w-1\right)}{2^w}
$$

$$
= \frac{1+3(w/2)+2\cdot 3^{w/2}-2\cdot w-2}{2^w}
$$

$$
= \frac{2\cdot 3^{w/2}-(w/2)-1}{2^w}
$$

$$
= 2\cdot 2^{-w}\cdot 3^{w/2}-2^{-w}\cdot(w/2)-2^{-w}
$$

$$
= 2^{-w+1}\cdot 2^{\log(3)w/2}-2^{-w}((w/2)+1)
$$

$$
= 2^{-w+1+\log(3)w/2}-2^{-w}2^{\log((w/2)+1)}
$$

$$
= 2^{1+w\left(\frac{\log(3)}{2}-1\right)}-2^{-w+\log((w/2)+1)}
$$

$$
\leq 2^{1+w\left(\frac{\log(3)}{2}-1\right)}-2^{1-w}
$$

The last step follows when $w \geq 2$. So we finally need to show that $2^{1+w\left(\frac{\log(3)}{2}-1\right)}-2^{1-w} \leq 2^{-s}$. To do so we replace w by our hypothesis term, $w = -\frac{s}{\frac{\log(3)}{2}-1} - \frac{1}{\frac{\log(3)}{2}-1}$.

$$
2^{-s} \geq^? 2^{1+\left(-\frac{s}{\frac{\log(3)}{2}-1}-\frac{1}{\frac{\log(3)}{2}-1}\right)\left(\frac{\log(3)}{2}-1\right)}-2^{1-\left(-\frac{s}{\frac{\log(3)}{2}-1}-\frac{1}{\frac{\log(3)}{2}-1}\right)}
$$

$$
= 2^{1+(-s-1)}-2^{1+\frac{s}{\frac{\log(3)}{2}-1}+\frac{1}{\frac{\log(3)}{2}-1}}
$$

$$
= 2^{-s}-2^{1+\frac{s}{\frac{\log(3)}{2}-1}+\frac{1}{\frac{\log(3)}{2}-1}}
$$

Next observe that $2^{1+\frac{s}{\frac{\log(3)}{2}-1}+\frac{1}{\frac{\log(3)}{2}-1}} \approx 2^{-4.82s-3.82}$. Thus for $s \geq 1$ we have that

$$
2^{-s}-2^{1+\frac{s}{\frac{\log(3)}{2}-1}+\frac{1}{\frac{\log(3)}{2}-1}} < 2^{-s}.
$$

Finally, since the amount of evaluation circuits are less than s for $s \geq 1$ then $c < w \Rightarrow 2s < \lceil 4.82s + 4.82 \rceil$ is clearly true and we are done. $\square$

We restate Lemma 7 here:

$$\frac{\binom{6s+7-\left(\frac{\lceil 4.82s+4.82\rceil}{2}+1\right)}{6s+7-(\lceil 4.82s+4.82\rceil)}}{\binom{6s+7}{6s+7-(\lceil 4.82s+4.82\rceil)}} \le 2^{-s}$$

for $s > 1$.

Proof. Consider $2^{-s} / \frac{\binom{6s+7-\left(\frac{\lceil 4.82s+4.82\rceil}{2}+1\right)}{6s+7-(\lceil 4.82s+4.82\rceil)}}{\binom{6s+7}{6s+7-(\lceil 4.82s+4.82\rceil)}}$. If we can show that this is an increasing function for $s > 1$, then we are

done, as this means that 2^{-s} is growing faster than $\frac{\binom{6s+7-\left(\frac{\lceil 4.82s+4.82\rceil}{2}+1\right)}{6s+7-(\lceil 4.82s+4.82\rceil)}}{\binom{6s+7}{6s+7-(\lceil 4.82s+4.82\rceil)}}$ and in turn that $\frac{\binom{6s+7-\left(\frac{\lceil 4.82s+4.82\rceil}{2}+1\right)}{6s+7-(\lceil 4.82s+4.82\rceil)}}{\binom{6s+7}{6s+7-(\lceil 4.82s+4.82\rceil)}} <$

2^{-s}. We see this by taking the derivative of $2^{-s} / \frac{\binom{6s+7-\left(\frac{\lceil 4.82s+4.82\rceil}{2}+1\right)}{6s+7-(\lceil 4.82s+4.82\rceil)}}{\binom{6s+7}{6s+7-(\lceil 4.82s+4.82\rceil)}}$ and noticing this is always positive.

$\square$

H   Two-Output Functionality

Until now we have only considered functions where Bob receives the output and where Alice does not receive any output. This is an artifact of the Yao garbling approach, caused by its asymmetric nature, in which functions with only output to Bob is simpler. In the general case both Alice and Bob get private output $f_A(x, y) = z_A$ and $f_B(x, y) = z_B$, respectively. Since Alice can just input a one-time pad r (which is "free" when using the free-xor technique [KS08]), constructing two-output functions is essentially a matter of ensuring *authenticity* or *correctness*, meaning mechanisms letting Alice be sure that the result she receives from Bob is indeed the *correct* result $f_A(x, y)$.

In the basic Yao scheme designed for a semi-honest adversary, a common technique for authenticating Alice's result is to let Bob send the output keys obtained to Alice. But for a malicious adversary, when multiple circuits are being sent, sending keys to Alice would reveal to Alice from which circuit Bob choose the output, opening up for selective failure attacks.

Lindell and Pinkas [LP07] suggested embedding a one-time (information theoretic) MAC of the result in the function to be computed, i.e., where Alice as part of her input submits the secret MAC key and later uses it to verify the result. This increases Alice's input with $|z_A| + 2s$ bits and the circuit size with $O(s|z_A|)$ gates, that is, $O(s^2|z_A|)$ inexpensive cryptographic operations.

Shelat and Shen [sS11] instead uses digital signatures and witness-indistinguishable proofs, giving $O(s|z_A|)$ *heavy* cryptographic operations. Mohassel and Riva [MR13] brings this down to $O(s|z_A|)$ *inexpensive* operations.