Fast Cut-and-Choose Based Protocols for Malicious and Covert Adversaries[*]

Yehuda Lindell

Dept. of Computer Science
Bar-Ilan University, Israel
lindell@biu.ac.il

September 1, 2014

Abstract

In the setting of secure two-party computation, two parties wish to securely compute a joint function of their private inputs, while revealing only the output. One of the primary techniques for achieving efficient secure two-party computation is that of Yao's garbled circuits (FOCS 1986). In the semi-honest model, where just one garbled circuit is constructed and evaluated, Yao's protocol has proven itself to be very efficient. However, a malicious adversary who constructs the garbled circuit may construct a garbling of a different circuit computing a different function, and this cannot be detected (due to the garbling). In order to solve this problem, many circuits are sent and some of them are opened to check that they are correct while the others are evaluated. This methodology, called *cut-and-choose*, introduces significant overhead, both in computation and in communication, and is mainly due to the number of circuits that must be used in order to prevent cheating.

In this paper, we present a cut-and-choose protocol for secure computation based on garbled circuits, with security in the presence of malicious adversaries, that vastly improves on all previous protocols of this type. Concretely, for a cheating probability of at most 2^{-40}, the best previous works send between 125 and 128 circuits. In contrast, in our protocol 40 circuits alone suffice (with some additional overhead). Asymptotically, we achieve a cheating probability of 2^{-s} where s is the number of garbled circuits, in contrast to the previous best of $2^{-0.32s}$. We achieve this by introducing a new cut-and-choose methodology with the property that in order to cheat, *all* of the evaluated circuits must be incorrect, and not just the *majority* as in previous works. The security of our protocol relies on the Decisional Diffie-Hellman assumption.

Keywords: two-party computation, Yao's protocol, cut-and-choose, concrete efficiency

1 Introduction

Background. Protocols for secure two-party computation enable a pair of parties P_1 and P_2 with private inputs x and y, respectively, to compute a function f of their inputs while preserving a number of security properties. The most central of these properties are *privacy* (meaning that the parties learn the output $f(x,y)$ but nothing else), *correctness* (meaning that the output received is indeed $f(x,y)$ and not something else), and *independence of inputs* (meaning that neither party can choose its input as a function of the other party's input). The standard way of formalizing these security properties is to compare the output of a real protocol execution to an "ideal execution" in which the parties send their inputs to an incorruptible trusted party who computes the output for the parties. Informally speaking, a protocol is then secure if no real adversary attacking the real protocol can do more harm than an ideal adversary (or simulator) who interacts in the ideal model [12, 10, 26, 2, 4, 11]. An important parameter when considering this problem relates to the power of the adversary. Three important models are the *semi-honest model* (where the adversary follows the protocol specification exactly but tries to learn more than it should by inspecting the protocol transcript), the *malicious model* (where the adversary can follow any arbitrary polynomial-time strategy), and the *covert model* (where the adversary may behave maliciously but is guaranteed to be caught with probability ϵ if it does [1]).

Efficient secure computation and Yao's garbled circuits. The problem of efficient secure computation has recently gained much interest. There are now a wide variety of protocols, achieving great efficiency in a variety of settings. These include protocols that require exponentiations for every gate in the circuit [30, 18] (these can be reasonable for small circuits but not large ones with tens or hundreds of thousands of gates), protocols that use the "cut and choose" technique on garbled circuits [22, 24, 31, 27], and more [29, 16, 17, 6, 20, 3, 28, 7]. The recent protocols of [28, 7] have very fast online running time. However, for the case of Boolean circuits and when counting the entire running time (and not just the online time), the method of cut-and-choose on garbled circuits is still the most efficient way of achieving security in the presence of covert and malicious adversaries.

Protocols for cut-and-choose on garbled circuits [22, 24, 31, 27] all work in the following way. Party P_1 constructs a large number of garbled circuits and sends them to party P_2. Party P_2 then chooses a subset of the circuits which are opened and checked. If all of these circuits are correct, then the remaining circuits are evaluated as in Yao's protocol [32], and P_2 takes the *majority* output value as the output. The cut-and-choose approach forces P_1 to garble the *correct* circuit, since otherwise it will be caught cheating. However, it is important to note that even if all of the opened circuits are correct, it is not guaranteed that all of the unopened circuits are correct. This is due to the fact that if there are only a small number of incorrect circuits, then with reasonable probability these may not be chosen to be opened. For this reason, it is critical that P_2 outputs the majority output, since the probability that a majority of unopened circuits are incorrect when all opened circuits are correct is exponentially small in the number of circuits. We stress that it is not possible for P_2 to abort in case it receives different outputs in different circuits, even though in such a case it knows that P_1 cheated, because this opens the door to the following attack. A malicious P_1 can construct a single incorrect circuit that computes the following function: if the first bit of P_2's input equals 0 then output random garbage; else compute the correct function. Now, if this circuit is not opened (which happens with probability 1/2) and if the first bit of P_2's input equals 0,

1

then P_2 will receive a different output in this circuit and in the others. In contrast, if the first bit of P_2's input equals 1 then it always receives the same output in all circuits. Thus, if the protocol instructs P_2 to abort if it receives different outputs, then P_1 will learn the first bit of P_2's input (based on whether or not P_2 aborts). By having P_2 take the majority value as output, P_1 can only cheat if the majority of the unopened circuits are incorrect, while all the opened ones are correct. In [24] it was shown that when s circuits are sent and half of them are opened, the probability that P_1 can cheat is at most $2^{-0.311s}$. Thus, concretely, in order to obtain an error probability of 2^{-40}, it is necessary to set $s = 128$ and so use 128 circuits, which means that the approximate cost of achieving security in the presence of malicious adversaries is 128 times the cost of achieving security in the presence of semi-honest adversaries. In [31], it was shown that by opening and checking 60% of the circuits instead of 50%, then the error becomes $2^{-0.32s}$ which means that it suffices to send 125 circuits in order to obtain a concrete error of 2^{-40}. When the circuits are large, this means a little less bandwidth. However, as we describe in Appendix C, this also results in more work and so is not necessarily better and may depend on the exact system setup. In [31] it was shown that checking 60% of the circuits is optimal for the cut-and-choose method as used until that time (i.e., checking a percentage of the circuits and taking the majority result from the rest). In this paper, we present an alternative cut-and-choose method that achieves far better parameters.

Our results. In this paper, we present a novel twist on the cut-and-choose strategy used in [22, 24, 31, 27] that enables us to achieve an error of just 2^{-s} with s circuits (and some small additional overhead). Concretely, this means that just 40 circuits are needed for error 2^{-40}. Our protocol is therefore much more efficient than previous protocols (there is some small additional overhead but this is greatly outweighed by the savings in the garbled circuits themselves unless the circuit being computed is small). We stress that the bottleneck in protocols of this type is the computation and communication of the s garbled circuits. This has been demonstrated in implementations. In [9], the cost of the circuit communication and computation for secure AES computation is approximately 80% of the work. Likewise in [19, Table 7] regarding secure AES computation, the bandwidth due to the circuits was 83% of all bandwidth and the time was over 50% of the time. On large circuits, as in the edit distance, this is even more significant with the circuit generation and evaluation taking 99.999% of the time [19, Table 9]. Thus, the reduction of this portion of the computation to a third of the cost is of great significance.

We present a high-level outline of our new technique in Section 2. For now, we remark that the cut-and-choose technique on Yao's garbled circuits introduces a number of challenges. For example, since the parties evaluate numerous circuits, it is necessary to enforce that the parties use the same input in all circuit computations. In addition, a selective input attack whereby P_1 provides correct garbled inputs only for a subset of the possible inputs of P_2 must be prevented (since otherwise P_2 will abort if its input is not in the subset because it cannot compute any circuit in this case, and thus P_1 will learn something about P_2's input based on whether or not it aborts). There are a number of different solutions to these problems that have been presented in [22, 24, 31, 27, 9]. The full protocol that we present here is based on the protocol of [24]. However, these solutions are rather "modular" (although this is meant in an informal sense), and can also be applied to our new technique; this is discussed at the end of Section 2. Understanding which technique is best will require implementation since they introduce tradeoffs that are not easily comparable. We leave this for future work, and focus on the main point of this work which is that it is possible to achieve error 2^{-s} with just s circuits. In Section 3.1 we present an exact efficiency count of our protocol, and compare it to [24].

Covert adversaries. Although not always explicitly proven, the known protocols for cut-and-choose on garbled circuits achieve covert security where the deterrent probability ϵ that the adversary is caught cheating equals 1 minus the statistical error of the protocol. That is, the protocol of [24] yields covert security of $\epsilon = 1 - 2^{-0.311s}$ (actually, a little better), and the protocol of [31] yields covert security with $\epsilon = 1 - 2^{-0.32s}$. Our protocol achieves covert security with deterrent $\epsilon = 1 - 2^{-s+1}$ (i.e., the error is 2^{-s+1}) which is far more efficient than all previous work. Specifically, in order to obtain $\epsilon = 0.99$, the number of circuits needed in [24] is 24. In contrast, with our protocol, it suffices to use 8 circuits. Furthermore, with just 11 circuits, we achieve $\epsilon = 0.999$, which is a very good deterrent. We note that s circuits do not suffice for obtaining $\epsilon = 1 - 2^{-s}$ deterrent, since the case that no circuits are computed (i.e., all are chosen to be checked) needs to be taken into account for small s.

Concurrent work. In independent, concurrent work, Huang et al. [14] presented a different method for reducing the number of garbled circuits. In their protocol, both parties generate s circuits and send them to each other. Thus, the overall number of circuits sent in their protocol is double that of ours. However, in a single-execution setting, the latency is likely to be similar.

2 The New Technique and Protocol Outline

The idea behind our new cut-and-choose strategy is to design a protocol with the property that the party who constructs the circuits (P_1) can cheat if and only if *all* of the opened circuits are correct and *all* of the evaluated circuits are incorrect. Recall that in previous protocols, if the circuit evaluator (P_2) aborts if the evaluated circuits don't all give the same output, then this can reveal information about P_2's input to P_1. This results in an absurd situation: P_2 knows that P_1 is cheating but cannot do anything about it. In our protocol, we run an additional *small* secure computation after the cut-and-choose phase so that if P_2 catches P_1 cheating (namely, if P_2 receives inconsistent outputs) then in the second secure computation it learns P_1's full input x. This enables P_2 to locally compute the correct output $f(x, y)$ once again. Thus, it is no longer necessary for P_2 to take the majority output. Details follow.

Phase 1 – first cut-and-choose:

- Parties P_1 (with input x) and P_2 (with input y) essentially run a protocol based on cut-and-choose of garbled circuits, that is secure for malicious adversaries (like [24] or [31]). P_1 constructs just s circuits (for error 2^{-s}) and the strategy for choosing check or evaluation circuits is such that each circuit is *independently* chosen as a check or evaluation circuit with probability $1/2$ (unlike all previous protocols where a fixed number of circuits are checked).

- If all of the circuits successfully evaluated by P_2 give the same output z, then P_2 locally stores z. Otherwise, P_2 stores a "proof" that it received two inconsistent output values in two different circuits. Such a proof could be having a garbled value associated with 0 on an output wire in one circuit, and a garbled value associated with 1 on the same output wire in a different circuit. (This is a proof since if P_2 obtains a single consistent output then the garbled values it receives on an output wire in different circuits are all associated with the same bit.)

Phase 2 – secure evaluation of cheating: P_1 and P_2 run a protocol that is secure for malicious adversaries with error 2^{-s} (e.g., they use the protocol of [24, 31] with approximately $3s$ circuits),

in order to compute the following:

- P_1 inputs the same input x as in the computation of phase 1 (and *proves* this).

- P_2 inputs random values if it received a single output z in phase 1, and inputs the proof of inconsistent output values otherwise.

- If P_2's input is a valid proof of inconsistent output values, then P_2 receives P_1's input x; otherwise, it receives nothing.

If this secure computation terminates with abort, then the parties abort.

Output determination: If P_2 received a single output z in phase 1 then it outputs z and halts. Otherwise, if it received inconsistent outputs then it received x in phase 2. P_2 locally computes $z = f(x, y)$ and outputs it. We stress that P_2 does not provide any indication as to whether z was received from phase 1 or locally computed.

Security.   The argument for the security of the protocol is as follows. Consider first the case that P_1 is corrupted and so may not construct the garbled circuits correctly. If all of the check circuits are correct and all of the evaluation circuits are incorrect, then P_2 may receive the same incorrect output in phase 1 and will therefore output it. However, this can only happen if each incorrect circuit is an evaluation circuit and each correct circuit is a check circuit. Since each circuit is an evaluation or check circuit with probability exactly $1/2$ this happens with probability exactly 2^{-s}. Next, if all of the evaluation circuits (that yield valid output) are correct, then the correct output will be obtained by P_2. This leaves the case that there are two different evaluation circuits that give two different outputs. However, in such a case, P_2 will obtain the required "proof of cheating" and so will learn x in the 2nd phase, thereby enabling it to still output the correct value. Since P_1 cannot determine which case yielded output for P_2, this can be easily simulated.

Next consider the case that P_2 is corrupted. In this case, the only way that P_2 can cheat is if it can provide output in the second phase that enables it to receive x. However, since P_1 constructs the circuits correctly, P_2 will not obtain inconsistent outputs and so will not be able to provide such a "proof". (We remark that the number of circuits s sent is used for the case that P_1 is corrupted; for the case that P_2 is corrupted a single circuit would actually suffice. Thus, there is no need to justify the use of fewer circuits than in previous protocols for this corruption case.)

Implementing phase 2.   The main challenge in designing the protocol is phase 2. As we have hinted, we will use the knowledge of two different garbled values for a single output wire as a "proof" that P_2 received inconsistent outputs. However, it is also necessary to make sure that P_1 uses the same input in phase 1 and in phase 2; otherwise it could use x or x', respectively, and then learn whether P_2 received output via phase 1 or 2. The important observation is that all known protocols already have a mechanism for ensuring that P_1 uses the same input in all computed circuits. This same mechanism can be used for the circuits in phase 1 and 2, since it does not depend on the circuits being computed being the same.

Another issue that arises is the efficiency of the computation in phase 2. In order to make the circuit for phase 2 small, it is necessary to construct all of the output wires in all the circuits of phase 1 so that they have the same garbled values on the output wires. This in turn makes it necessary to open and check the circuits only after phase 2 (since opening a circuit to check it

reveals both garbled values on an output wire which means that this knowledge can no longer be a proof that P_1 cheated). Thus, the structure of the actual protocol is more complex than previous protocols; however, this relates only to its description and not efficiency.

We remark that we use the method of [24] in order to prove the consistency of P_1's input in the different circuits and between phase 1 and phase 2. However, we believe that the methods used in [31, 27], for example, would also work, but have not proven this.

3   The Protocol

Preliminaries – modified batch single-choice cut-and-choose OT.   The cut-and-choose OT primitive was introduced in [24]. Intuitively, a cut-and-choose OT is a series of 1-out-of-2 oblivious transfers with the special property that in some of the transfers the receiver obtains a single value (as in regular oblivious transfer), while in the others the receiver obtains both values. In the context of cut-and-choose on Yao's garbled circuits, the functionality is used for the receiver to obtain all garbled input values in the circuits that it wishes to open and check, and to obtain only the garbled input values associated with its input on the circuits to be evaluated.

In [24], the functionality defined is such that the receiver obtains both values in exactly half of the transfers; this is because in [24] exactly half of the circuits are opened. In this work, we modify the functionality so that the receiver can choose at its own will in which transfers it receives just one value and in which it receives both. We do this since we want P_2 to check each circuit with probability exactly $1/2$, independently of all other circuits. This yields an error of 2^{-s} instead of $\binom{s}{s/2}^{-1}$, which is smaller (this is especially significant in the setting of covert adversaries; see Section 5).

This modification introduces a problem since at a later stage in the protocol the receiver needs to prove to the sender for which transfers it received both values and for which it received only one. If it is known that the receiver obtains both values in exactly half of the transfers, or for any other known number, then the receiver can just send both values in these transfers (assuming that they are otherwise unknown, as is the case in the Yao circuit use of the functionality), and the sender knows that the receiver did *not* obtain both values in all others; this is what is done in [24]. However, here the receiver can obtain both values in an unknown number of transfers, as it desires. We therefore need to introduce a mechanism enabling the receiver to prove to the sender in which transfers it did not receive both values, in a way that it cannot cheat. We solve this by having the sender input s random "check" values, and having the receiver obtain such a value in every transfer for which it receives a single value only. Thus, at a later time, the receiver can send the appropriate check values, and this constitutes a proof that it did not receive both values in these transfers. See Figure 3.1, and Section 6 for the construction and details.

FIGURE 3.1 (Modified Batch Single-Choice Cut-and-Choose OT Functionality $\mathcal{F}_{\mathrm{ccot}}$)

- **Inputs:**

  - S inputs ℓ *vectors* of pairs $\vec{x}_i$ of length s, for $i = 1, \ldots, \ell$. (Every vector consists of s pairs; i.e., $\vec{x}_i = \langle (x_0^{i,1}, x_1^{i,1}), (x_0^{i,2}, x_1^{i,2}), \ldots, (x_0^{i,s}, x_1^{i,s}) \rangle$. There are ℓ such vectors.) In addition, S inputs s "check values" $\chi_1, \ldots, \chi_s$. All values are in $\{0,1\}^n$.

  - R inputs $\sigma_1, \ldots, \sigma_\ell \in \{0,1\}$ and a set of indices $\mathcal{J} \subseteq [s]$.

- **Output:** The sender receives no output. The receiver obtains the following:

  - For every $i = 1, \ldots, \ell$ and for every $j \in \mathcal{J}$, the receiver R obtains the jth pair in vector $\vec{x}_i$. (I.e., for every $i = 1, \ldots, \ell$ and every $j \in \mathcal{J}$, R obtains $(x_0^{i,j}, x_1^{i,j})$.)

  - For every $i = 1, \ldots, \ell$, the receiver R obtains the σ_i value in every pair of the vector $\vec{x}_i$. (I.e., for every $i = 1, \ldots, \ell$, R obtains $\langle x_{\sigma_i}^{i,1}, x_{\sigma_i}^{i,2}, \ldots, x_{\sigma_i}^{i,s} \rangle$.)

  - For every $k \notin \mathcal{J}$, the receiver R obtains χ_k.

Encoded translation tables. We modify the output translation tables typically used in Yao's garbled circuits as follows. Let k_i^0, k_i^1 be the garbled values on wire i, which is an output wire, and let H be a collision-resistant hash function. Then, the encoded output translation table for this wire is simply $\left[H(k_i^0), H(k_i^1) \right]$. We require that $k_i^0 \neq k_i^1$ and if this doesn't hold (which will be evident since then $H(k_i^0) = H(k_i^1)$), P_2 will automatically abort.

Observe that given a garbled value k, it is possible to determine whether k is the 0 or 1 key (or possibly neither) by just computing $H(k)$ and seeing if it equals the first or second value in the pair, or neither. However, given the encoded translation table, it is not feasible to find the actual garbled values, since this is equivalent to inverting the one-way function. This is needed in our protocol, as we will see below. We remark that both k_i^0, k_i^1 are revealed by the end of the protocol, and only need to remain secret until Step 7 has concluded (see the protocol below). Thus, they can be relatively short values (in practice, 80 bits should be well long enough).

We require that H be collision resistant, since otherwise P_1 who constructs the circuit is not bounded by the values there. This can cause an issue since there may now be multiple values that are consistent with the encoded translation tables, and different values may be used in the actual garbled circuits constructed in phase 1, and in the circuits used to prevent cheating in phase 2 (this would cause a situation where P_2 will not learn P_1's input x even if it obtained inconsistent output values in phase 1).

PROTOCOL 3.2 (Computing $f(x, y)$)

Inputs: P_1 has input $x \in \{0, 1\}^\ell$ and P_2 has input $y \in \{0, 1\}^\ell$.

Auxiliary input: a statistical security parameter s, the description of a circuit C such that $C(x, y) = f(x, y)$, and $(\mathbb{G}, q, g)$ where $\mathbb{G}$ is a cyclic group with generator g and prime order q, and q is of length n. In addition, they hold a hash function H that is a suitable randomness extractor; see [8].

Specified output: Party P_2 receives $f(x, y)$ and party P_1 receives no output (this suffices for the general case where the parties receive different outputs). Denote the length of the output of $f(x, y)$ by m.

The protocol:

1. INPUT KEY CHOICE AND CIRCUIT PREPARATION:

    (a) P_1 chooses random values $a_1^0, a_1^1, \ldots, a_\ell^0, a_\ell^1; r_1, \ldots, r_s \in_R \mathbb{Z}_q$ and $b_1^0, b_1^1, \ldots, b_m^0, b_m^1 \in_R \{0, 1\}^n$.

    (b) Let $w_1, \ldots, w_\ell$ be the input wires corresponding to P_1's input in C, and denote by $w_{i,j}$ the instance of wire w_i in the jth garbled circuit, and by $k_{i,j}^b$ the key associated with bit b on wire $w_{i,j}$. Then, P_1 sets the keys for its input wires to:
    $$k_{i,j}^0 = H(g^{a_i^0 \cdot r_j}) \quad \text{and} \quad k_{i,j}^1 = H(g^{a_i^1 \cdot r_j})$$

    (c) Let $w_1', \ldots, w_m'$ be the output wires in C. Then, the keys for wire w_i' in *all* garbled circuits are b_i^0 and b_i^1 (we stress that unlike all other wires in the circuit, the same values are used for the output wires in all circuits).

    (d) P_1 constructs s independent copies of a garbled circuit of C, denoted $GC_1, \ldots, GC_s$, using random keys except for wires $w_1, \ldots, w_\ell$ (P_1's input wires) and $w_1', \ldots, w_m'$ (the output wires) for which the keys are as above.

2. OBLIVIOUS TRANSFERS: P_1 and P_2 run a modified batch single-choice cut-and-choose oblivious transfer, with parameters ℓ (the number of parallel executions) and s (the number of pairs in each execution):

    (a) P_1 defines vectors $\vec{z}_1, \ldots, \vec{z}_\ell$ so that $\vec{z}_i$ contains the s pairs of random symmetric keys associated with P_2's ith input bit y_i in all garbled circuits $GC_1, \ldots, GC_s$. P_1 also chooses random values $\chi_1, \ldots, \chi_s \in_R \{0, 1\}^n$. P_1 inputs these vectors and the $\chi_1, \ldots, \chi_s$ values.

    (b) P_2 chooses a random subset $\mathcal{J} \subset [s]$ where every $j \in \mathcal{J}$ with probability exactly $1/2$, under the constraint that $\mathcal{J} \neq [s]$. P_2 inputs the set $\mathcal{J}$ and bits $\sigma_1, \ldots, \sigma_\ell \in \{0, 1\}$, where $\sigma_i = y_i$ for every i.

    (c) P_2 receives all the keys associated with its input wires in all circuits GC_j for $j \in \mathcal{J}$, and receives the keys associated with its input y on its input wires in all other circuits.

    (d) P_2 receives χ_j for every $j \notin \mathcal{J}$.

3. SEND CIRCUITS AND COMMITMENTS: P_1 sends P_2 the garbled circuits (i.e., the garbled gates). In addition, P_1 sends P_2 the "seed" for the randomness extractor H, and the following displayed values (which constitute a "commitment" to the garbled values associated with P_1's input wires):
    $$\left\{ (i, 0, g^{a_i^0}), (i, 1, g^{a_i^1}) \right\}_{i=1}^\ell \quad \text{and} \quad \left\{ (j, g^{r_j}) \right\}_{j=1}^s$$
    In addition, P_1 sends P_2 the encoded output translation tables, as follows:
    $$\left[\left(H(b_1^0), H(b_1^1) \right), \ldots, \left(H(b_m^0), H(b_m^1) \right) \right].$$
    If $H(b_i^0) = H(b_i^1)$ for any $1 \leq i \leq m$, then P_2 aborts.

4. SEND CUT-AND-CHOOSE CHALLENGE: P_2 sends P_1 the set $\mathcal{J}$ along with the values χ_j for every $j \notin \mathcal{J}$. If the values received by P_1 are incorrect, it outputs $\perp$ and aborts. Circuits GC_j for $j \in \mathcal{J}$ are called check-circuits, and for $j \notin \mathcal{J}$ are called evaluation-circuits.

5. P_1 SENDS ITS GARBLED INPUT VALUES IN THE EVALUATION-CIRCUITS: P_1 sends the keys associated with its inputs in the evaluation circuits: For every $j \notin \mathcal{J}$ and every wire $i = 1, \ldots, \ell$, party P_1 sends the value $k_{i,j}' = g^{a_i^{x_i} \cdot r_j}$; P_2 sets $k_{i,j} = H(k_{i,j}')$.

7

PROTOCOL 3.2 – continued

6. CIRCUIT EVALUATION: P_2 uses the keys associated with P_1's input obtained in Step 5 and the keys associated with its own input obtained in Step 2c to evaluate the circuits GC_j for every $j \notin \mathcal{J}$.

   If P_2 receives only one valid output value per output wire (i.e., one of b_i^0, b_i^1, verified against the encoded output translation tables) and it does not abort in the next step, then this will be its output. If P_2 receives two valid outputs on one output wire (i.e., both b_i^0 and b_i^1 for output wire w_i') then it uses these in the next step. If there exists an output wire for which P_2 did not receive a valid value in *any* evaluation circuit (neither b_i^0 nor b_i^1), then P_2 will abort in Step 8b below (and in the next step will use garbage input).

7. RUN SECURE COMPUTATION TO DETECT CHEATING:

   (a) P_1 defines a circuit with the values $b_1^0, b_1^1, \ldots, b_m^0, b_m^1$ hardcoded. The circuit computes the following function:

      i. P_1's input is a string $x \in \{0,1\}^\ell$, and it has no output.
      ii. P_2's input is a pair of values b_0, b_1.
      iii. If there exists a value i $(1 \le i \le m)$ such that $b_0 = b_i^0$ *and* $b_1 = b_i^1$, then P_2's output is x; otherwise it receives no output.

   (b) P_1 and P_2 run the secure computation protocol (Protocol 4.1) of [24] on the circuit defined in the previous step, with some differences, as follows:

      i. P_1 inputs its input x; if P_2 received b_i^0, b_1^i for some $1 \le i \le m$, then it inputs the pair b_i^0, b_i^1; otherwise it inputs garbage.
      ii. P_1 and P_2 use the statistical security parameter $s' = 3s$ (in order to obtain a maximum cheating probability of 2^{-s}).
      iii. The values $a_1^0, a_1^1, \ldots, a_\ell^0, a_\ell^1$ chosen by P_1 in Step 1a of Protocol 4.1 in [24] (for defining the garbled values on the wires associated with P_1's input) are the same $a_1^0, a_1^1, \ldots, a_\ell^0, a_\ell^1$ as chosen in Step 1a above in this protocol. (We stress that all other values are chosen randomly and independently.)
      iv. In Step 6 of Protocol 4.1 in [24], P_2 checks that the opened circuits are correctly constructed. Note that the values $b_1^0, b_1^1, \ldots, b_m^0, b_m^1$ are hardcoded in the definition of the circuit in Step 7a. P_2 adds an *additional check* that the hardcoded values $b_1^0, b_1^1, \ldots, b_m^0, b_m^1$ are consistent with the values sent in Step 3 above. Specifically, P_2 computes $H(b_i^0), H(b_i^i)$ for every i and compares the result with the values in the encoded translation tables.
      v. P_1 and P_2 do *not* run the proof of consistency of P_1's input values (Step 7b of Protocol 4.1 in [24]); this is carried out in Step 9 below.

      If this computation results in an abort, then both parties halt at this point and output $\perp$.

8. CHECK CIRCUITS FOR COMPUTING $f(x,y)$:

   (a) SEND ALL INPUT GARBLED VALUES IN CHECK-CIRCUITS: For every *check-circuit GC_j*, party P_1 sends the value r_j to P_2, and P_2 checks that these are consistent with the pairs $\{(j, g^{r_j})\}_{j \in \mathcal{J}}$ received in Step 3. If not, P_2 aborts outputting $\perp$.

   (b) CORRECTNESS OF CHECK CIRCUITS: For every $j \in \mathcal{J}$, P_2 uses the $g^{a_i^0}, g^{a_i^1}$ values it received in Step 3, and the r_j values it received in Step 8a, to compute the values $k_{i,j}^0 = H(g^{a_i^0 \cdot r_j}), k_{i,j}^1 = H(g^{a_i^1 \cdot r_j})$ associated with P_1's input in GC_j. In addition it sets the garbled values associated with its own input in GC_j to be as obtained in the cut-and-choose OT. Given all the garbled values for all input wires in GC_j, party P_2 decrypts the circuit and verifies that it is a garbled version of C, using the encoded translation tables for the output values. If there exists a circuit for which this does not hold, then P_2 aborts and outputs $\perp$. In addition, if there exists an output wire for which P_2 did not receive a valid value in *any* evaluation circuit in Step 6, then P_2 aborts and outputs $\perp$.

PROTOCOL 3.2 – continued

9. VERIFY CONSISTENCY OF P_1'S INPUT: Let $\hat{\mathcal{J}}$ be the set of check circuits in the execution of Protocol 4.1 of [24] run in Step 7, and let $\hat{r}_1, \ldots, \hat{r}_{3s}$ be the values chosen in Step 1a of Protocol 4.1 of [24] (analogous to the values $r_1, \ldots, r_s$ chosen in Step 1a in this protocol above). For $1 \leq i \leq \ell$ and $1 \leq \hat{j} \leq 3s$, let $\hat{k}_{i,\hat{j}} = g^{a_i^{x_i} \cdot \hat{r}_{\hat{j}}}$ be the value used for deriving the key associated with the ith bit of P_1's input in the $\hat{j}$th circuit in the protocol execution of Step 7 (analogous to the value $k'_{i,j}$ in Step 5 above).

   For every input wire $i = 1, \ldots, \ell$, party P_1 proves (using a zero-knowledge proof of knowledge) that there exists a $\sigma_i \in \{0,1\}$ such that for every $j \notin \mathcal{J}$ and every $\hat{j} \notin \hat{\mathcal{J}}$, $k'_{i,j} = g^{a_i^{\sigma_i} \cdot r_j}$ AND $\hat{k}_{i,\hat{j}} = g^{a_i^{\sigma_i} \cdot \hat{r}_{\hat{j}}}$. Concretely P_1 proves that:

   $$\left(\forall j \notin \mathcal{J} : (g, g^{a_i^0}, g^{r_j}, k'_{i,j}) \in DH \text{ and } \forall \hat{j} \notin \hat{\mathcal{J}} : (g, g^{a_i^0}, g^{\hat{r}_{\hat{j}}}, \hat{k}_{i,\hat{j}}) \in DH \right)$$
   $$\text{OR} \quad \left(\forall j \notin \mathcal{J} : (g, g^{a_i^1}, g^{r_j}, k'_{i,j}) \in DH \text{ and } \forall \hat{j} \notin \hat{\mathcal{J}} : (g, g^{a_i^1}, g^{\hat{r}_{\hat{j}}}, \hat{k}_{i,\hat{j}}) \in DH \right)$$

   where DH is the language of all Diffie-Hellman tuples.[1]
   If any of the proofs fail, then P_2 aborts and outputs $\perp$.

10. OUTPUT EVALUATION: If P_2 received no inconsistent outputs from the evaluation circuits GC_i ($i \notin \mathcal{J}$), then it decodes the outputs it received using the encoded translation tables, and outputs the string received. If P_2 received inconsistent output, then let x be the output that P_2 received from the second computation in Step 7. Then, P_2 computes $f(x, y)$ and outputs it.

The circuit for step 7. A naive circuit for computing the function in Step 7 can be quite large. Specifically, to compare two n bit strings requires $2n$ XORs followed by $2n$ ORs; if the output is 0 then the strings are equal. This has to be repeated m times, once for every i, and then the results have to be ORed. The result of all of the ORs is then ORed with x. Thus, there are $2mn + m + 1$ non-XOR gates. Assuming n is of size 80 (e.g., which suffices for the output values) and m is of size 128, this requires $20,481$ non-XOR gates, which is very large. An alternative is therefore to compute the following garbled circuit:

1. For every $i = 1, \ldots, m$,

   (a) Compare $b_0 \| b_1$ to $b_i^0 \| b_i^1$ (where '$\|$' denotes concatenation) by XORing bit-by-bit, and take the NOT of each bit. This is done as in a regular garbled circuit; by combining the NOT together with the XOR this has the same cost as a single XOR gate.

   (b) Compute the $2n$-wise AND of the bits from above. Instead of using $2n - 1$ Boolean AND gates, this can be achieved by encrypting the 1-key on the output wire under all n keys (together with redundancy so that the circuit evaluator can know if it received the correct value). Furthermore, this encryption can be a "one-time pad" and thus is just the XOR of all of the 1-keys on the input wires together with the 1-key on the output wire. The 0-key for the output can be given in the clear, since it provides no additional information, but is not needed so can just not be given (note that P_2 knows exactly which case it is in). Note that the result of this operation is 1 if and only if $b_0 \| b_1 = b_i^0 \| b_i^1$ and so P_2 had both keys on the ith output wire.

---

[1] This can be efficiently proven by combining all the ANDs into two tuples, and then proving that either the first is a Diffie-Hellman tuple, or the second is a Diffie-Hellman tuple. A description of how such batching can be carried out is described in [24].

2. Compute the OR of the m bits resulting from the above loop. Instead of using $m-1$ Boolean OR gates, this can be achieved by simply setting the 1-key on all of the output wires from the n-wise ANDs above to be the 1-key on the output wire of the OR. This ensures that as soon as the 1-key is received from an n-wise AND, the 1-key is received from the OR, as required. (This reveals for which i the result of the n-wise AND was 1. However, this is fine in this specific case since P_2 knows exactly where equality should be obtained in any case.)

3. Compute the AND of the output from the previous step with all of the input bits of P_1. This requires ℓ Boolean AND gates.

4. The output wires include the output of the OR (so that P_2 can know if it received x or nothing), together with the output of all of the ANDs with the input bits of P_1.

The original and optimized circuits are depicted in Appendices A.1 and A.2. The number of non-XOR operations required to securely compute this circuit is just ℓ binary AND gates. Assuming $\ell = 128$ (e.g., as in the secure AES example), we have that there are only 128 non-XOR gates. When using 128 circuits as in our instantiation of Step 7 via [24], this comes to 16,384 garbled gates *overall*, which is significant but not too large. We stress that the size of this circuit is independent of the size of the circuit for the function f to be computed. Thus, this becomes less significant as the circuit becomes larger. On the other hand, for very small circuits or when the input size is large relative to the overall circuit size, our approach will not be competitive. To be exact, assume a garbled circuit approach that requires $3s$ circuits. If $3s|C| < s|C| + 3s \cdot \ell$ then our protocol will be slower (since the cost of our protocol is $s|C|$ for the main computation plus $3s\ell$ for the circuit of Step 7, in contrast to $3s|C|$ for the other protocol). This implies that our protocol will be faster as long as $|C| > \frac{3\ell}{2}$. Concretely, if $\ell = 128$ and $s = 40$, it follows that our protocol will be faster as long as $|C| > 192$. Thus, our protocol is much faster, except for the case of very small circuits.

Additional optimizations.[2] Observe that although the above circuit is very small, P_2's input size is $2n$ and this is quite large. Since the input size has a significant effect on the cost of the protocol (especially when using cut-and-choose oblivious transfer), it would be desirable to reduce this. This can be achieved by first having P_2 input $b_0 \oplus b_1$ instead of $b_0 \| b_1$, reducing the input length to n (this is sound since if P_2 does not have both keys on any output wire then it cannot know their XOR). Furthermore, in order to obtain a cheating probability of 2^{-40} it suffices for the circuit to check only the first 40 bits of $b_0 \oplus b_1$. (Note that b_0^i and b_1^i have to be longer since $H(b_i^0), H(b_i^1)$ are published; nevertheless, only 40 bits need to be included in the circuit. When using this optimization, the length of b_0^i, b_1^i can be 128 bits and not 80, which is preferable.) Finally, by choosing all of the b_i^0, b_i^1 values so that they have the same fixed XOR (i.e., for some Δ it holds that for all i, $b_i^0 \oplus b_i^1 = \Delta$, as in the free XOR technique), the size of the circuit is further reduced. In this case, an *additional check* has to be introduced, which is to check that there exists a single Δ such that $b_i^0 \oplus b_i^1 = \Delta$ for all $i = 1, \ldots, m$ (in the check phase, the output translation tables are already checked, so the additional step here is to just check that the XOR is all fine). This optimization requires that we assume that the hash function is correlation robust [15] with such correlated inputs (i.e., these look random); clearly this holds in the random oracle model. This significantly reduces the bandwidth; see the figure in Appendix A.3.

---

[2]We thank Benny Pinkas for pointing these optimizations out to us.

3.1 A Detailed Efficiency Count and Comparison

In this section we provide an exact efficiency count of our protocol. This will enable an exact comparison of our protocol to previous and future works, as long as they also provide an exact efficiency count. We stress that implementations and empirical comparisons are also necessary to obtain a full picture, due to additional factors like memory, parallelism and so on.

We count exponentiations, symmetric encryptions and bandwidth. We let n denote the length of a symmetric encryption, and an arbitrary string of length of the security parameter (e.g., χ_j). The cost of Step 2 (cut-and-choose oblivious transfer) is taken from the exact count provided in Section 6.4. The counts below are for the overall number of operations for *both* parties P_1 and P_2.

Step	Fixed-base exponent.	Regular exponent.	Symmetric Encryptions	Group elems sent	Symmetric comm
1	$2s\ell$	0	$8s\lvert C\rvert$		
2	$9s\ell$	$1.5s\ell$		$5s\ell$	
3	$\ell + s$	0		$2\ell + s$	$4ns\lvert C\rvert$
4					$\frac{s}{2}\cdot n$
5					nm
6			$\frac{s}{2}\cdot 2\lvert C\rvert$		
7	$9s\ell + 5040s$	$480s$	$39s\ell$	$21s\ell$	$12sn\ell$
8	$s/2 + s\ell$		$\frac{s}{2}\cdot 8\lvert C\rvert$		$\frac{s}{2}\cdot n$
9		$2s\ell + 18\ell$		10	$2s\ell n$
TOTAL	$\mathbf{21s\ell + 5040s}$	$\mathbf{3.5s\ell + 18\ell}$ $\mathbf{+480s}$	$\mathbf{13s\lvert C\rvert +}$ $\mathbf{39s\ell}$	$\mathbf{26s\ell}$	$\mathbf{4ns\lvert C\rvert + 14s\ell n}$

The number of symmetric encryptions is counted as follows: each circuit requires $8\lvert C\rvert$ symmetric encryptions to construct (counting $\lvert C\rvert$ as the number of non-XOR gates when using free XOR), $48\lvert$ symmetric encryption to check, and $2\lvert C\rvert$ encryptions to evaluate (using standard double encryption, each entry in a garbled gate requires 2 encryptions). Since approximately half of the circuits are checked and half are evaluated, the garbling, checking and evaluation of the main garbled circuit accounts for approximately $s\cdot 8\lvert C\rvert + \frac{s}{2}\cdot 8\lvert C\rvert + \frac{s}{2}\cdot 2\lvert C\rvert = 13s\lvert C\rvert$ symmetric encryptions. The garbled circuit used in Step 7 has ℓ non-XOR gates and so the same analysis applies on this size. However, the number of circuits sent in this step is $3s$ and thus we obtain an additional $3\times 13\cdot s\cdot \ell = 39s\ell$.

The bandwidth count for Step 7 is computed based on the counts provided in [24], using $3s$ circuits. The cost of the exponentiations is based on the fact that in [24], if P_1 has input of length ℓ_1 and P_2 has input of length ℓ_2, and s' circuits are used, then there are $3.5s'\ell_1 + 10.5s'\ell_2$ fixed-base exponentiations and $s'\ell_2$ regular exponentiations. However, $0.5s'\ell_1$ of the fixed-base exponentiations are for the proof of consistency and these are counted in Step 9 instead. Now, in Step 7, P_1's input length is ℓ (it is the same x as for the entire protocol) and P_2's input is comprised of two garbled values for the output wires. Since these must remain secret for only a short amount of time, it is possible to take 80-bit values only and so P_2's input length is 160 bits (this is irrespective of P_2's input length to the function f). Taking $s' = 3s$ and plugging these lengths this into the above, we obtain the count appearing in the table.

The proof of consistency of P_1's input is carried out ℓ times (once for each bit of P_1's input) and over $s + 3s = 4s$ circuits (since there are s circuits for the main computation of C, plus another $3s$ circuits for the computation in Step 7). By the count in [24], this proof therefore costs

$\frac{4s\ell}{2} + 18\ell$ exponentiations, and bandwidth of 10 group elements and another $8s\ell$ short strings (this can therefore be counted as $2s\ell n$.

A comparison to [24].   In order to to get a concrete understanding of the efficiency improvement, we will compare the cost to [24] for the AES circuit of size 6,800 gates [33], and input and output sizes of 128. Now, as we have mentioned, the overall cost of the protocol of [24] is $3.5s'\ell_1 + 10.5s\ell_2$ fixed-base exponentiations, $s'\ell_2$ regular exponentiations and $13s'|C|$ symmetric encryptions. In this case, $\ell_1 = \ell_2 = 128$, $s' = 125$ ($s' = 125$ was shown to suffice for 2^{-40} security in [19]), and so we have that the cost is 224,000 fixed-base exponentiations, 32,000 regular exponentiations, and $1625|C| =$11,050,000 symmetric encryptions. In contrast, taking $\ell = 128$ and $s = 40$ we obtain here 309,120 fixed-base exponentiations, $37,120$ regular exponentiations, and 3,749,600 symmetric encryptions. In addition, the bandwidth of [24] is approximately $112,000$ group elements and 3,400,000 symmetric ciphertexts. At the minimal cost of 220 bits per group element (e.g., using point compression) and 128 bits per ciphertext, we have that this would come to approximately 449,640,000 bits, or close to half a gigabyte (in practice, it would be significantly larger due to communication overheads). In contrast, the bandwidth of our protocol for this circuit would be 133,120 group elements and 1,159,680 ciphertexts. With the same parameters as above, this would be approximately 177,725,440 bits, which is under 40% of the cost of [24]. This is very significant since bandwidth is turning out to be the bottleneck in many cases.

Protocol	Fixed-base exp.	Regular exp.	Symmetric encryptions	Bandwidth
[24]	224,000	16,000	11,050,000	449,640,000
Here	309,120	37,120	3,749,600	177,725,440

Figure 1: Comparison of protocols for secure computation of AES

We stress that in larger circuits the difference would be even more striking.

4   Proof of Security

We prove security via the standard definition of secure two-party computation following the real/ideal paradigm, and using modular sequential composition; see [4, 11, 13] for details.

Theorem 4.1 *Assume that the Decisional Diffie-Hellman assumption holds in* $(\mathbb{G}, g, q)$, *that* H *is a collision-resistant function, and that the garbling of the circuit is secure as in [24]. Then, Protocol 3.2 securely computes* f *in the presence of malicious adversaries (with error* $2^{-s} + \mu(n)$ *where* $\mu(\cdot)$ *is some negligible function).*[3]

Proof: We prove Theorem 4.1 in a hybrid model where a trusted party is used to compute the modified batch single-choice cut-and-choose oblivious transfer functionality and the zero-knowledge proof of knowledge of Step 9 (the fact that any zero-knowledge proof of knowledge fulfills the zero-knowledge proof of knowledge functionality was shown in [13]). We separately prove the case that P_1 is corrupted and the case that P_2 is corrupted.

---

[3]The error is actually $\frac{1}{2^s - 1}$ since J must be chosen to not equal all of $[s]$. Thus, there are "only" $2^s - 1$ possible cuts. We ignore this minor detail.

P_1 is corrupted. Intuitively, P_1 can only cheat by constructing some of the circuits in an incorrect way. However, in order for this to work, all of the check circuits must be valid, and all of the evaluation circuits that give output must give the *same* output. Otherwise, P_2 will abort (if there is an invalid check circuit) or will obtain x and output the correct value $f(x, y)$ (if two different outputs are obtained). Now, once P_1 sends the circuits, and the cut-and-choose oblivious transfer has been completed, the question of whether a circuit is valid or not, or can be computed or not, is fully determined. (It is true that P_1 can send incorrect values for its input wires but this will be detected in the zero knowledge phase and so will cause an abort except with negligible probability.) Now, if there exists even just one valid circuit that is an evaluation circuit then P_2 will always output the correct value; either because all evaluation circuits that can be computed are valid or because it will obtain two different outputs and so will learn x in Step 7. Thus, in order to cheat, every valid circuit must be a check circuit and every invalid circuit must be an evaluation circuit. Since every circuit is chosen to be a check or evaluation circuit with probability exactly $1/2$, it follows that this probability of cheating is exactly 2^{-s}. We now proceed to the formal proof.

Let $\mathcal{A}$ be an adversary controlling P_1 in an execution of Protocol 3.2 where a trusted party is used to compute the modified batch single-choice cut-and-choose OT functionality $\mathcal{F}_{\mathrm{ccot}}$ and the zero-knowledge proof of knowledge of Step 9. We construct a simulator $\mathcal{S}$ who runs in the ideal model with a trusted party computing f. $\mathcal{S}$ runs $\mathcal{A}$ internally and simulates the honest P_2 for $\mathcal{A}$ as well as the trusted party computing the oblivious transfer and zero-knowledge proof of knowledge functionalities. In addition, $\mathcal{S}$ interacts *externally* with the trusted party computing f. $\mathcal{S}$ works as follows:

1. $\mathcal{S}$ interacts with $\mathcal{A}$ and plays the honest P_2 for the entire protocol execution with input $y = 0^\ell$ (in the $\mathcal{F}_{\mathrm{ccot}}$ and zero-knowledge hybrid models).

2. Let $x = \sigma_1, \ldots, \sigma_\ell$ be the witness used by P_1 in the zero-knowledge proof of knowledge of Step 9 (i.e., x is the concatenation of the $\sigma_1, \ldots, \sigma_\ell$ values where σ_i is the appropriate part of the witness used in the ith proof). $\mathcal{S}$ obtains this witness directly from $\mathcal{A}$ in the zero-knowledge proof of knowledge hybrid model (since $\mathcal{A}$ hands it to the trusted party computing the functionality). If the witness is not valid, then P_2 would abort and this is dealt with in the next step.

3. If P_2 would abort in the execution, then $\mathcal{S}$ sends $\bot$ to the ideal functionality computing f. Otherwise, it sends x.

First, observe that P_2 uses its input only in $\mathcal{F}_{\mathrm{ccot}}$. Therefore, the view of $\mathcal{A}$ in the simulation is *identical* to its view in a real execution, and so $\mathcal{S}$'s output distribution in an ideal execution is identical to $\mathcal{A}$'s output distribution in a real execution of Protocol 3.2. However, in order to prove security, we need to prove that the *joint distribution* consisting of $\mathcal{A}$ and P_2's output after a real protocol execution is indistinguishable from the *joint distribution* of $\mathcal{S}$ and P_2's output after an ideal execution. Now, if P_2 were to abort in a real execution, then $\mathcal{S}$ sends $\bot$ to the trusted party and thus the joint distributions will be the same. However, the probability that $\mathcal{S}$ sends $\bot$ to the trusted party may *not* be the same as the probability that P_2 aborts in a real execution since this may actually depend on P_2's input (recall that $\mathcal{S}$ uses input 0^ℓ and not the same y that P_2 uses). In addition to this difference, we must also show that when P_2 does not abort in a real execution then its output is $f(x, y)$, relative to the same x as sent by $\mathcal{S}$, except with probability that equals $2^{-s} + \mu(n)$, where $\mu(\cdot)$ is some negligible function.

We begin by defining the notion of a bad and good circuit. For a garbled circuit GC_j we define the circuit input keys as follows:

1. *Circuit input keys associated with P_1's input:* Let $(i, 0, g^{a_i^0}), (i, 1, g^{a_i^1}), (j, g^{r_j})$ be the values sent by P_1 to P_2 in Step 3 of the protocol. Then, the circuit input keys associated with P_1's input in GC_j are the keys $H(g^{a_1^0 \cdot r_j}), H(g^{a_1^1 \cdot r_j}), \ldots, H(g^{a_\ell^0 \cdot r_j}), H(g^{a_\ell^1 \cdot r_j})$.

2. *Circuit input keys associated with P_2's input:* Let $(z_0^{1,j}, z_1^{1,j}), \ldots, (z_0^{\ell,j}, z_1^{\ell,j})$ be the set of symmetric keys associated with the jth circuit that were input by P_1 to the cut-and-choose oblivious transfer of Step 2. (These keys are the jth pair in each vector $\vec{z}_1, \ldots, \vec{z}_\ell$.) These values are called the circuit input keys associated with P_2's input in GC_j.

Then, a garbled circuit is bad if the circuit keys associated with both P_1's and P_2's input do *not* open it to the correct circuit C. We stress that after Step 3 of the protocol, each circuit is either "bad" or "not bad", and this is fully determined. (The simulator cannot know at this point if a circuit is bad or not, but it is nevertheless fully determined.) This holds because P_1 has already sent the $\{g^{a_i^0}, g^{a_i^1}, g^{r_j}\}$ values, the garbled circuits and the seed to the randomness extractor in this step; note that once the seed to the randomness extractor is fixed, the symmetric keys derived from the Diffie-Hellman values are fully determined. In addition, the keys associated with P_2's input are already fixed since these are fully determined in Step 2. A garbled circuit is good if it is not bad.

Clearly, if a check-circuit is bad, then P_2 aborts. In addition, we claim that if a *single* evaluation circuit is good and P_2 does not abort, then the distribution over P_2's output is the same in a real and simulated execution (except with probability $2^{-s} + \mu(n)$). In order to see this, observe that if there exists an evaluation circuit that is good, then P_2 obtains $f(x, y)$ when evaluating that circuit. Any bad evaluation circuit that does not give correct output (with respect to the encoded output tables) is ignored. In addition, if there exists another bad evaluation circuit that yields a different output to $f(x, y)$, then P_2 receives x in the computation in Step 7. This holds because if P_2 receives two different outputs then there exists a wire for which it receives both outputs; equivalently, there exists an encoded output pair for which P_2 receives both preimages. Since the computation in Step 7 is a secure protocol (as proven in [24]), it follows that P_2 either aborts or receives the correct output x (except with probability $2^{-s} + \mu(n)$). Note that the change in the computation with respect to the zero-knowledge proof of consistency of P_1's input is inconsequential with respect to the proof of security of [24]. Thus the security obtained is except with probability $2^{-s} + \mu(n)$; the fact that 2^{-s} is achieved is by using $3s$ circuits (to be exact, one would need to the exact computation based on the error given in [24] but this is a very close approximation). In this case where P_2 learns x in this step, it also outputs the correct $f(x, y)$. Thus, we conclude that if there exists a single evaluation circuit that is good and P_2 does not abort, then P_2 outputs $f(x, y)$. The last remaining case is the case where there exists an output wire for which P_2 did not receive a valid value in *any* evaluation circuit. In this case, P_2 aborts. However, this case means that *all* evaluation circuits are bad. Now, if an additional check circuit is also bad, then P_2 aborts due to a circuit check failing. In contrast, if all of the check circuits are good, then P_2 will not abort due to the check failing but rather only due to it not receiving a valid value on the output wire. This may be different from the simulation, since the question of P_2 aborting or not in this case may depend on the input.[4]

---

[4]Specifically, P_1 may construct all of the circuits correctly, but may use one incorrect key – say, the 1-key – for one of the input wires associated with P_2's input in a subset of the circuits in the oblivious transfer. In this case, if P_2's associated input is 1 and all of the evaluation circuits are in the subset, then P_2 will abort. In contrast, the simulator uses input $y = 0^\ell$ and so will not abort due to this error.

This implies that the simulation can deviate from the real output distribution if and only if *every* check circuit is good and *every* evaluation circuit is bad. Since the association of a circuit as being bad or good is fixed before $\mathcal{J}$ is revealed (and it is information theoretically hidden from $\mathcal{A}$ in the $\mathcal{F}_{\text{ccot}}$-hybrid model), it follows that a bad circuit is an evaluation circuit with probability exactly $1/2$ and a good circuit is a check circuit with probability exactly $1/2$ (and these probabilities between circuits are independent). There are s circuits and so the probability that the above holds is exactly 2^{-s}, as required.

We observe that when P_1 is corrupted, $\mathcal{A}$ has zero probability of cheating in the oblivious transfers and zero probability of cheating in the proof of consistency (this holds in the hybrid models; in the real model there is a negligible probability of cheating). Thus, it can only cheat if the circuits are in the bad case described above, or in the secure computation of Step 7. This completes the proof of this corruption case.

We remark that it is implicit in the above analysis that $\mathcal{A}$ cannot detect whether P_2 outputs $f(x, y)$ due to all evaluation circuits giving the same output or due to it obtaining x in Step 7. This is because the simulator does not need to be able to distinguish between these cases, and indeed this can depend on the real input used by P_2. Likewise, $\mathcal{A}$ cannot detect whether P_2 aborts due to all evaluation circuits not giving a valid output, or whether it aborts due to the fact that it detected a bad circuit in the "correctness of check circuits" step.

P_2 is corrupted. The intuition behind the security in this corruption case is simple, and is identical to [24] and other cut-and-choose protocols based on garbled circuits. Specifically, every circuit is either an evaluation circuit, in which case P_2 can only learn the output, or a check circuit in which case P_2 learns nothing. The security here relies on the DDH assumption and the fact that the garbled values sent by P_1 and P_2 for the input wires reveal nothing about P_1's input. In addition, it relies on the fact that P_2 cannot learn two different outputs by evaluating the evaluation circuits, or in general cannot learn two different values on an output wire (which would be a "proof that P_1 cheated" even though it is honest). This is needed since otherwise it could input these into the secure computation of Step 7 and learn P_1's input x. Nevertheless, this follows directly from the security of [24] which in the case of a corrupted P_2 holds for any number of circuits, from the security of the computation in Step 7, and from the fact that it is hard to invert the one-way function H used to compute the encoded output tables.

Formally, the simulator $\mathcal{S}$ works as follows:

1. $\mathcal{S}$ obtains P_2's input to the oblivious transfers (where $\mathcal{A}$ controls P_2) and uses this to define the set $\mathcal{J}$ and the input y.

2. $\mathcal{S}$ sends y to the trusted party computing f and receives back $z = f(x, y)$.

3. For every $j \in \mathcal{J}$, $\mathcal{S}$ constructs a proper garbled circuit. For every $j \notin \mathcal{J}$, $\mathcal{S}$ constructs a garbled circuit with exactly the same structure as the correct one, but that outputs the fixed string z received from the trusted party, irrespective of the input. As in the protocol, all the circuits have the same garbled values on the output wires, and they have the same structure for P_1's inputs.

4. $\mathcal{S}$ continues to run the protocol with $\mathcal{A}$, playing the honest P_1 but using the circuits constructed in the previous step.

5. At the conclusion of the protocol execution with $\mathcal{A}$, simulator $\mathcal{S}$ outputs whatever $\mathcal{A}$ outputs and halts.

This completes the simulation. The formal proof that the simulation is indistinguishable from a real execution follows from the indistinguishability of "fake" garbled circuits from real ones. The formal reduction is almost identical to [24] and so is omitted in this abstract. We remark only that the output distributions between the simulated and real executions is negligible in n, and independent of s; see [24] for details. This completes the proof. ∎

5 Variants – Universal Composability and Covert Adversaries

Universal composability [5].   As in [24], by instantiating the cut-and-choose oblivious transfer and the zero-knowledge proofs with variants that are universally composable, the result is that Protocol 3.2 is universally composable.

Covert adversaries [1] (see Appendix B).   Observe that in the case that P_2 is corrupted, the protocol is fully secure irrespective of the value of s used. In contrast, when P_1 is corrupted, then the cheating probability is $2^{-s} + \mu(n)$. However, this cheating probability is *independent* of the input used by the P_2 (as shown in the proof of Theorem 4.1). Thus, Protocol 3.2 is suitable for the model of security in the presence of *covert adversaries*. Intuitively, since the adversary can cheat with probability only 2^{-s} and otherwise it is caught cheating, the protocol achieves covert security with deterrent $\epsilon = 1 - 2^{-s}$. However, on closer inspection, this is incorrect. Specifically, as we have discussed above, if P_2 catches P_1 cheating due to the fact that two different circuits yield two different outputs, then it is not allowed to reveal this fact to P_1. Thus, P_2 cannot declare that P_1 is a cheat in this case, as is required in the model of covert adversaries. However, if P_2 detects even a single bad circuit in the check phase, then it can declare that P_1 is cheating, and this happens with probability at least $1/2$ (even if only a single circuit is bad). We can use this to show that for every s, Protocol 3.2 securely computes f in the presence of covert adversaries **with deterrent $\epsilon = 1 - 2^{-s+1}$**. In actuality, we need to make a slight change to Protocol 3.2, as follows:

- In Step 3 of Protocol 3.2 (SEND CIRCUITS AND COMMITMENTS), P_1 also proves that it knows the discrete log r_j in every pair (j, g^{r_j}) sent, using a zero-knowledge proof of knowledge.

- In Step 8b of Protocol 3.2 (CORRECTNESS OF CHECK CIRCUITS), if there exists a circuit for which the checks do not hold, then P_2 outputs corrupted$_1$ and halts (instead of just aborting and outputting $\perp$).

The additional zero-knowledge proof of knowledge enables the simulator $\mathcal{S}$ to completely open all circuits by the end of Step 3 of the protocol (by obtaining all of the values on P_2's input wires via the oblivious transfers and all of the values on P_1's inputs wires using r_j values). We remark also this zero-knowledge proof of knowledge requires only a small constant number of exponentiations for every j (note that only a small number of circuits are used here also since this is the setting of covert adversaries).

In addition, we stress that the honest P_2 *only declares the P_1 is* corrupted *(and attempted cheating) if one of the check circuits is found to be invalid.* In particular, P_2 does not declare that P_1 is corrupted if all checked circuits are correct, but two different evaluated circuits yield different outputs. In this latter case, like in the setting of malicious adversaries, P_2 obtains P_1's input x and uses it to locally compute the correct output without revealing to P_1 that anything happened.

We now proceed to security. We refer the reader to the definition of security for covert adversaries in Appendix B and [1]; we use the strong explicit definition here. The simulator $\mathcal{S}$ for the case that P_1 is corrupted works as follows:

1. $\mathcal{S}$ works like the simulator in the case of malicious adversaries and determines the set $\mathcal{G}$ of good (or valid) circuits and the set $\mathcal{B}$ of bad (or invalid) circuits. It determines this by using all of P_1's inputs to the oblivious transfers and r_j values extracted from the zero-knowledge proof of knowledge described above. If none of circuits are bad, then $\mathcal{S}$ simulates exactly as in the case of the malicious setting when no circuits are invalid. Otherwise, there is at least one bad circuit and $\mathcal{S}$ proceeds to the next step.

2. With probability $\frac{1}{2}$, the simulator $\mathcal{S}$ sends cheat_1 to the trusted party.

   - If $\mathcal{S}$ receives back $\mathsf{undetected}$ then it receives the input y of the honest party P_2 and can determine its output as it wishes. Thus, it proceeds in the simulation by choosing the set of circuits to opened and checked to be exactly the set of good circuits $\mathcal{G}$. Otherwise, it runs the code of the honest P_2 with input y with $\mathcal{A}$, and then sets the output of the honest P_2 in the ideal execution to be whatever the simulated P_2 receives (including the possibility that P_2 outputs $\mathsf{corrupted}_1$).

   - If $\mathcal{S}$ receives back $\mathsf{corrupted}_1$ then $\mathcal{S}$ chooses the set of circuits to be opened and checked in the simulation of the protocol at random under the constraint that at least one circuit from $\mathcal{B}$ is opened and checked.

3. Else, (i.e., in the case that $\mathcal{S}$ does *not* send cheat_1 to the trusted party):

   - With probability $1 - 2^{-|\mathcal{B}|+1} + 2^{-s+1}$, $\mathcal{S}$ chooses the set of circuits to be opened and checked in the simulation of the protocol at random under the constraint that at least one circuit from $\mathcal{B}$ is opened and checked. In addition, $\mathcal{S}$ sends $\mathsf{corrupted}_1$ as P_1's input to the trusted party. (This "strange" probability is explained below.)

   - Otherwise, $\mathcal{S}$ chooses the set of circuits to be opened and checked in the simulation of the protocol at random under the constraint that *no* circuits from $\mathcal{B}$ are opened and checked.

4. Apart from the above $\mathcal{S}$ behaves exactly as in the case of simulation of malicious adversaries.

We now analyze the above. In order to understand the probability distribution, observe that if $\mathcal{B}$ circuits are invalid then in a real protocol execution the following occurs:

- With probability 2^{-s} the adversary $\mathcal{A}$ succeeds in cheating (this is the case that the set of evaluated circuits is exactly $\mathcal{B}$).

- With probability $1 - 2^{-|B|}$ party P_2 catches $\mathcal{A}$ cheating and outputs $\mathsf{corrupted}_1$ (this holds since P_2 catches $\mathcal{A}$ cheating unless every circuit in $\mathcal{B}$ is not opened, which happens with probability $2^{-|\mathcal{B}|}$).

- Otherwise, with probability $2^{-|\mathcal{B}|} - 2^{-s}$, party P_2 outputs the correct result as if no cheating took place.

Now, in the simulation, we have the following probabilities:

- With probability $\frac{1}{2}$, $\mathcal{S}$ sends cheat$_1$ to the trusted party who returns undetected with probability 2^{-s+1}. Thus, successful cheating occurs with probability $\frac{1}{2} \cdot 2^{-s+1} = 2^{-s}$, exactly as in a real execution.

- With probability $1 - 2^{-s+1}$, the trusted party returns corrupted$_1$ after $\mathcal{S}$ sends cheat$_1$. Thus, P_2 outputs corrupted$_1$ due to this case with probability $\frac{1}{2} \cdot (1 - 2^{-s+1}) = \frac{1}{2} - 2^{-s}$. However, in addition, P_2 outputs corrupted$_1$ with probability $1 - 2^{-|\mathcal{B}|+1} + 2^{-s+1}$ in the case that $\mathcal{S}$ does not send cheat$_1$. Since this case occurs with probability $\frac{1}{2}$ as well, this contributes probability $\frac{1}{2} \cdot (1 - 2^{-|\mathcal{B}|+1} + 2^{-s+1}) = \frac{1}{2} - 2^{-|\mathcal{B}|} + 2^{-s}$. Summing these together, we have that P_2 outputs corrupted$_1$ in the simulation with probability $\frac{1}{2} - 2^{-s} + \frac{1}{2} - 2^{-|\mathcal{B}|} + 2^{-s} = 1 - 2^{-|\mathcal{B}|}$, exactly as in a real execution.

- Finally, in all other cases (with probability $2^{-|\mathcal{B}|} - 2^{-s}$), P_2 outputs the correct output as if no cheating took place, exactly as in a real execution.

As discussed in the introduction, this yields a huge efficiency improvement over previous results, especially for small values of s. For example, 100 circuits are needed to obtain $\epsilon = 0.99$ in [1], 24 circuits are needed to obtain $\epsilon = 0.99$ in [24], and here *8 circuits alone* suffice to obtain $\epsilon = 0.99$. Observe that when covert security is desired, the number of circuits sent in Step 7 needs to match the level of covert security. For example, in order to obtain $\epsilon = 0.99$, 8 circuits are used in our main protocol and 24 circuits are used in Step 7.

We remark that our protocol would be a little simpler if P_2 always asked to open exactly half the circuits (especially in the cut-and-choose oblivious transfer). In this case, the error would be $\binom{s}{s/2}^{-1}$ instead of 2^{-s}. In order to achieve an error of 2^{-40} this would require 44 circuits which is a 10% increase in complexity, and reason enough to use our strategy of opening each circuit independently with probability $1/2$. However, when considering covert security, the difference is huge. For example, with $s = 8$ we have that $\binom{8}{4}^{-1} = 1/70$ whereas $2^{-8} = 1/256$. This is a very big difference.

6 Constructing Modified Cut-and-Choose Oblivious Transfer

6.1 Preliminaries – The $RAND$ Function

Let $(\mathbb{G}, g, q)$ be such that $\mathbb{G}$ is a group of order q, with generator g. We define the function $RAND(w, x, y, z) = (u, v)$, where $u = (w)^t \cdot (y)^{t'}$ and $v = (x)^t \cdot (z)^{t'}$, and the values $t, t' \leftarrow \mathbb{Z}_q$ are random. $RAND$ has the following properties:

Claim 6.1 *Let $(\mathbb{G}, g, q)$ be as above and let $w, x, y, z \in \mathbb{G}$. Furthermore, let (w, x, y, z) be a Diffie-Hellman tuple, and let $a \in \mathbb{Z}_q$ be such that $x = w^a$ and $z = y^a$. Then, for $(u, v) \leftarrow RAND(w, x, y, z)$ it holds that $u^a = v$.*

Proof: By the definition of $RAND$, $u = w^t \cdot y^{t'}$ and $v = x^t \cdot z^{t'}$. Since $x = w^a$ and $z = y^a$ we can write that

$$v = x^t \cdot z^{t'} = (w^a)^t \cdot (y^a)^{t'} = (w^t \cdot y^{t'})^a = u^a,$$

as required. ∎

Claim 6.2 *Let* $(\mathbb{G}, g, q)$ *be as above and let* $w, x, y, z \in \mathbb{G}$. *If* (w, x, y, z) *is not a Diffie-Hellman tuple, then the distributions* $(w, x, y, z, RAND(w, x, y, z))$ *and* $(w, x, y, z, g^\alpha, g^\beta)$, *where* $\alpha, \beta \leftarrow \mathbb{Z}_q$ *are random, are* equivalent.

Proof: Let $a, b \in \mathbb{Z}_q$ such that $x = w^a$ and $z = y^b$ and $a \neq b$ (such a pair a, b exists since (w, x, y, z) is not a Diffie-Hellman tuple). Let $\alpha, \beta \in \mathbb{Z}_q$; we show that $\Pr[RAND(w, x, y, z) = (g^\alpha, g^\beta)] = \frac{1}{q^2}$, where the probability is taken over t, t' used to compute $RAND$. Let $\gamma \in \mathbb{Z}_q$ be such that $y = w^\gamma$. Then, we have that $(w, x, y, z) = (w, w^a, w^\gamma, w^{\gamma \cdot b})$, and so $u = w^t \cdot w^{\gamma \cdot t'}$ and $v = w^{a \cdot t} \cdot w^{\gamma \cdot b \cdot t'}$. Thus, $(u, v) = (g^\alpha, g^\beta)$ if and only if $t + \gamma \cdot t' = \alpha$ and $a \cdot t + \gamma \cdot b \cdot t' = \beta$. These equations have a single solution if and only if the matrix

$$\begin{pmatrix} 1 & \gamma \\ a & \gamma \cdot b \end{pmatrix}$$

is invertible, which holds here since the determinant is $\gamma \cdot \beta - \gamma \cdot a = \gamma \cdot (b - a) \neq 0$, where the inequality holds since $a \neq b$. Thus, there is a single pair t, t' such that $(u, v) = (g^\alpha, g^\beta)$ implying that the probability is $\frac{1}{q^2}$, as required. ∎

6.2 The Protocol for Modified Cut-and-Choose OT

Our protocol is based on the batch single-choice cut-and-choose OT in [24], with appropriate changes. The most significant difference here is that unlike [24] the number of transfers in which the receiver obtains both strings is not fixed a priori, but is rather determined singlehandedly by the receiver. Thus, there is no proof that half of the pairs (h_0^i, h_1^i) are such that (g_0, g_1, h_0^i, h_1^i) are Diffie-Hellman tuples (i.e., a Diffie-Hellman tuple is of the form (g_0, g_1, h_0, h_1) where there exists a value w such that $h_0 = (g_0)^w$ and $h_1 = (g_1)^w$). In addition, there are additional steps used for R to receive strings χ_j for every $j \notin \mathcal{J}$. This uses a similar mechanism to the transfer of the other strings. Specifically, R is able to obtain χ_j for any j for which the tuple $(g_0, g_1, h_0^j, h_1^j/g_1)$ is a Diffie-Hellman tuple, which is exactly for all $j \notin \mathcal{J}$. We remark that in Step 3, R first "rerandomizes" the values h_0^j and h_1^j/g_1, obtaining $\tilde{h}_0^j, \tilde{h}_1^j$. This seems to be unnecessary, since the protocol works when the sender computes $RAND$ on the initial values h_0^j and h_1^j/g_1. However, in the proof of security, the simulator needs to be able to cheat and obtain all values χ_j. This is a problem since the simulator also needs to obtain all pairs from a cheating sender and so must choose h_0^j, h_1^j such that *all* (g_0, g_1, h_0^j, h_1^j) *are* Diffie-Hellman tuples. Since the simulator cannot make all of these tuples be Diffie-Hellman tuples at the same time as making all $(g_0, g_1, h_0^j, h_1^j/g_1)$ Diffie-Hellman tuples, the rerandomization is introduced in order to enable the simulator to cheat. Specifically, the simulator will choose all h_0^j, h_1^j such that (g_0, g_1, h_0^j, h_1^j) is a Diffie-Hellman tuple, and will then choose all $\tilde{h}_0^j, \tilde{h}_1^j$ so that all $(g_0, g_1, \tilde{h}_0^j, \tilde{h}_1^j)$ are also Diffie-Hellman tuples. This enables it to obtain all values. (Of course, the simulator will have to cheat in the zero-knowledge proof that $(h_0^j, h_1^j/g_1, \tilde{h}_0^j, \tilde{h}_1^j)$ since will this *not* be the case in the aforementioned way that the simulator chooses the values. Nevertheless, this is indistinguishable by the Decisional Diffie-Hellman assumption.

The protocol appears below.

PROTOCOL 6.3 (Modified Batch Single-Choice Cut-and-Choose Oblivious Transfer)

Sender's Input: The sender's input is ℓ vectors of pairs $\vec{x}_1, \ldots, \vec{x}_\ell$. We denote $\vec{x}_i = \langle (x_0^{i,1}, x_1^{i,1}), (x_0^{i,2}, x_1^{i,2}), \ldots, (x_0^{i,s}, x_1^{i,s}) \rangle$. In addition, the sender inputs $\chi_1, \ldots, \chi_s$. All of the sender's input values are of length exactly n.

Receiver's Input: The receiver's input is comprised of ℓ bits $\sigma_1, \ldots, \sigma_\ell$ and a set $\mathcal{J} \subset [s]$.

Auxiliary input: Both parties hold a security parameter 1^n and $(\mathbb{G}, q, g_0)$, where $\mathbb{G}$ is an efficient representation of a group of order q with a generator g_0, and q is of length n. In addition, they hold a hash function H that is a suitable randomness extractor from random elements of $\mathbb{G}$ to uniform strings of length n; see [8].

THE PROTOCOL:

1. **Setup phase:**
   (a) R chooses a random $y \leftarrow \mathbb{Z}_q$, sets $g_1 = (g_0)^y$, and sends g_1 to S.
   (b) R proves to S in zero-knowledge that it knows the discrete log of g_1, relative to g_0 (i.e., formally, the relation for the zero-knowledge is $\{((g_0, g_1), a) \mid g_1 = (g_0)^a\}$.
   (c) For every $j = 1, \ldots, s$, R chooses a random $\alpha_j \leftarrow \mathbb{Z}_q$ and computes $h_0^j = (g_0)^{\alpha_j}$. For every $j \in \mathcal{J}$, R computes $h_1^j = (g_1)^{\alpha_j}$, and for every $j \notin \mathcal{J}$, R computes $h_1^j = (g_1)^{\alpha_j + 1}$.
   (d) R sends $(h_0^1, h_1^1, \ldots, h_0^s, h_1^s)$ to S

2. **Transfer phase (repeated in parallel for every $i = 1, \ldots, \ell$):**
   (a) The receiver chooses a random value $r_i \leftarrow \mathbb{Z}_q$ and computes $\tilde{g}_i = (g_{\sigma_i})^{r_i}$ and $\tilde{h}_{i,1} = (h_{\sigma_i}^1)^{r_i}, \ldots, \tilde{h}_{i,s} = (h_{\sigma_1}^s)^{r_i}$. It sends $(\tilde{g}_i, \tilde{h}_{i,1}, \ldots, \tilde{h}_{i,s})$ to the sender.
   (b) The receiver proves in zero-knowledge that either all of $\{(g_0, \tilde{g}_i, h_0^j, \tilde{h}_{i,j})\}_{j=1}^s$ are Diffie-Hellman tuples, or all of $\{(g_1, \tilde{g}_i, h_1^j, \tilde{h}_{i,j})\}_{j=1}^s$ are Diffie-Hellman tuples.
   (c) The sender operates in the following way:
      - For every $j = 1, \ldots, s$, S sets $(u_0^{i,j}, v_0^{i,j}) = RAND(g_0, \tilde{g}_i, h_0^j, \tilde{h}_{i,j})$, and $(u_1^{i,j}, v_1^{i,j}) = RAND(g_1, \tilde{g}_i, h_1^j, \tilde{h}_{i,j})$.
      - For every $j = 1, \ldots, s$, S sends the receiver the values $(u_0^{i,j}, w_0^{i,j})$ where $w_0^{i,j} = H(v_0^{i,j}) \oplus x_0^{i,j}$, and $(u_1^{i,j}, w_1^{i,j})$ where $w_1^{i,j} = H(v_1^{i,j}) \oplus x_1^{i,j}$.

3. **Transfer of χ_j values:**
   (a) For every $j = 1, \ldots, s$, R chooses a random $\rho_j \leftarrow \mathbb{Z}_q$ and computes $\tilde{h}_0^j = (h_0^j)^{\rho_j}$. For every $j \notin \mathcal{J}$, R computes $\tilde{h}_1^j = (h_1^j/g_1)^{\rho_j}$, and for every $j \in \mathcal{J}$, R computes $\tilde{h}_1^j = (h_1^j)^{\rho_j}$. R sends $(\tilde{h}_0^1, \tilde{h}_1^1, \ldots, \tilde{h}_0^s, \tilde{h}_1^s)$ to S.
   (b) R proves in zero-knowledge that all $\{(g_0, g_1, \tilde{h}_0^j, \tilde{h}_1^j)\}_{j=1}^s$ are Diffie-Hellman tuples.
   (c) For every $j = 1, \ldots, s$, S sets $(u_j, v_j) = RAND(\tilde{h}_0^j, h_1^j/g_1, \tilde{h}_0^j, \tilde{h}_1^j)$.
   (d) For every $j = 1, \ldots, s$, S sends R the values (u_j, w_j) where $w_j = H(v_j) \oplus \chi_j$.

4. **Output:** S outputs nothing. R's output is as follows:
   (a) For every $i = 1, \ldots, \ell$ and for every $j = 1, \ldots, s$, R computes $x_{\sigma_i}^{i,j} = w_{\sigma_i}^{i,j} \oplus H((u_{\sigma_i}^{i,j})^{r_i})$.
   (b) For every $i = 1, \ldots, \ell$ and every $j \in \mathcal{J}$, R computes $x_{1-\sigma_i}^{i,j} = w_{1-\sigma_i}^{i,j} \oplus H((u_{1-\sigma_i}^{i,j})^{r_i \cdot z})$, where $z = y^{-1} \bmod q$ if $\sigma = 0$, and $z = y$ if $\sigma = 1$.
   (c) R outputs $(x_{\sigma_1}^1, \ldots, x_{\sigma_1}^s), \ldots, (x_{\sigma_\ell}^1, \ldots, x_{\sigma_\ell}^s)$ and $\left\{ \left\langle (x_0^{1,j}, x_1^{1,j}), \ldots, (x_0^{\ell,j}, x_1^{\ell,j}) \right\rangle \right\}_{j \in \mathcal{J}}$.
   (d) For every $j \notin \mathcal{J}$, R outputs $\chi_j = w_j \oplus H((u_j)^{\rho_j})$.

Correctness. Since it is not immediate from the protocol description, before proceeding to prove security we show that the output of the protocol is correct. First, we show that R receives the values χ_j for every $j \notin \mathcal{J}$. For every j the receiver obtains (u_j, w_j) where $w_j = H(v_j) \oplus \chi_j$ and $(u_j, v_j) = RAND(h_0^j, h_1^j, \tilde{h}_0^j, \tilde{h}_1^j)$. For every $j \notin \mathcal{J}$ we have that $\tilde{h}_0^j = (h_0^j)^{\rho_j}$ and $\tilde{h}_1^j = (h_1^j/g_1)^{\rho_j}$. Thus $(h_0^j, h_1^j/g_1, \tilde{h}_0^j, \tilde{h}_1^j)$ is a Diffie-Hellman tuple. By Claim 6.1 this implies that $v_j = (u_j)^{\rho_j}$, for every $j \notin \mathcal{J}$. Thus, $w_j \oplus H((u_j)^{\rho_j}) = \chi_j$, as required. Observe also that all $(g_0, g_1, \tilde{h}_0^j, \tilde{h}_1^j)$ are Diffie-Hellman tuples, because if $j \in \mathcal{J}$ then $\tilde{h}_0^j = (h_0^j)^{\rho_j} = (g_0)^{\alpha_j \cdot \rho_j}$ and $\tilde{h}_1^j = (h_1^j)^{\rho_j} = (g_1)^{\alpha_j \cdot \rho_j}$, and if $j \in \mathcal{J}$ then $\tilde{h}_0^j = (h_0^j)^{\rho_j} = (g_0)^{\alpha_j \cdot \rho_j}$ and $\tilde{h}_1^j = (h_1^j/g_1)^{\rho_j} = (g_1)^{\alpha_j \cdot \rho_j}$. Thus, the zero-knowledge proof of Step 3b is accepted.

Likewise, for every $i = 1, \ldots, \ell$ and every $j = 1, \ldots, s$, by the way R chooses its values we have that $(g_\sigma, \tilde{g}_i, h_\sigma^j, \tilde{h}_{i,j})$ is a Diffie-Hellman tuple, where $\tilde{g}_i = (g_\sigma)^{r_i}$ and thus by Claim 6.1 it holds that $(u_{\sigma_i}^{i,j})^{r_i} = v_{\sigma_i}$.

Finally, for every $j \in \mathcal{J}$, the values are chosen by P_2 so that $h_0^j = (g_0)^{\alpha_j}$ and $h_1^j = (g_1)^{\alpha_j}$. Recall also that $g_1 = (g_0)^y$, and thus $h_1^j = (g_1)^{\alpha_j} = (g_0)^{y \cdot \alpha_j} = (h_0^j)^y$. We now show that for all $j \in \mathcal{J}$, both $(g_0, \tilde{g}_i, h_0^j, \tilde{h}_{i,j})$ and $(g_1, \tilde{g}_i, h_1^j, \tilde{h}_{i,j})$ are Diffie-Hellman tuples:

1. If $\sigma_i = 0$ then $(g_0, \tilde{g}_i, h_0^j, \tilde{h}_{i,j})$ is a Diffie-Hellman tuple as shown above. Regarding $(g_1, \tilde{g}_i, h_1^j, \tilde{h}_{i,j})$, observe that $\tilde{g}_i = (g_0)^{r_i}$ and $\tilde{h}_{i,j} = (h_0^j)^{r_i}$. Since $g_1 = (g_0)^y$ and $h_1^j = (h_0^j)^y$, we have that $\tilde{g}_i = (g_1)^{r_i/y}$ and $\tilde{h}_{i,j} = (h_0^j)^{r_i} = (h_1^j)^{r_i/y}$. By Claim 6.1, it follows that $v_1^{i,j} = (u_1^{i,j})^{r_i/y}$.

2. If $\sigma_i = 1$ then $(g_1, \tilde{g}_i, h_1^j, \tilde{h}_{i,j})$ is a Diffie-Hellman tuple as shown above. Regarding $(g_0, \tilde{g}_i, h_0^j, \tilde{h}_{i,j})$, observe that $\tilde{g}_i = (g_1)^{r_i}$ and $\tilde{h}_{i,j} = (h_1^j)^{r_i}$. Since $g_1 = (g_0)^y$ and $h_1^j = (h_0^j)^y$, we have that $\tilde{g}_i = (g_0)^{y \cdot r_i}$ and $\tilde{h}_{i,j} = (h_1^j)^{r_i} = (h_0^j)^{y \cdot r_i}$. By Claim 6.1, it follows that $v_1^{i,j} = (u_1^{i,j})^{y \cdot r_i}$.

6.3 The Proof of Security of Protocol 6.3

Theorem 6.4 *If the Decisional Diffie-Hellman problem is hard in $(\mathbb{G}, g, q)$, then Protocol 6.3 securely computes the $\mathcal{F}_{\mathrm{ccot}}$ functionality. If the zero-knowledge proofs of Step 1b of the setup phase and Step 2b of the transfer phase are universally composable, then Protocol 6.3 UC-computes the $\mathcal{F}_{\mathrm{ccot}}$ functionality.*

Proof: We prove security in a hybrid model where the zero-knowledge proofs and proofs of knowledge (ZKPOK) are computed by ideal functionalities (where the prover sends (x, w) and the functionality sends 1 to the verifier if and only if (x, w) is in the relation); the fact that any zero-knowledge proof of knowledge securely computes this functionality has been formally proven in [**?**, Section 6.5.3].

Case 1 – corrupted receiver: Let $\mathcal{A}$ be an adversary that controls the receiver R. We construct a simulator $\mathcal{S}$ that invokes $\mathcal{A}$ on its input and works as follows:

1. $\mathcal{S}$ receives g_1 and receives the input $((g_0, g_1), y)$ that $\mathcal{A}$ sends to the ideal ZKPOK functionality. If $g_1 \neq (g_0)^y$, then $\mathcal{S}$ simulates S aborting, and outputs whatever $\mathcal{A}$ outputs. Else, $\mathcal{S}$ proceeds to the next step.

2. $\mathcal{S}$ receives $(h_0^1, h_1^1, \ldots, h_0^s, h_1^s)$ from $\mathcal{A}$. For every $j = 1, \ldots, s$, $\mathcal{S}$ checks if $h_1^j = (h_0^j)^y$ then it sets $j \in \mathcal{J}$ and otherwise it sets $j \notin \mathcal{J}$.

3. For every $i = 1, \ldots, \ell$, $\mathcal{S}$ receives $(\tilde{g}_i, \tilde{h}_{i,1}, \ldots, \tilde{h}_{i,s})$ from $\mathcal{A}$.

4. For every $i = 1, \ldots, \ell$, $\mathcal{S}$ receives the witness r_i that $\mathcal{A}$ sends to ZKPOK $(\tilde{g}_i, \tilde{h}_{i,1}, \ldots, \tilde{h}_{i,s})$. If it does not hold that $[\tilde{g}_i = (g_0)^{r_i}$ and every $\tilde{h}_{i,j} = (h_0^j)^{r_i}]$ or $[\tilde{g}_i = (g_1)^{r_i}$ and every $\tilde{h}_{i,j} = (h_1^j)^{r_i}]$, then $\mathcal{S}$ simulates S aborting, and outputs whatever $\mathcal{A}$ outputs. Otherwise, for every i, let σ_i be the bit for which it holds.

5. $\mathcal{S}$ sends $\mathcal{J}$ and $\sigma_1, \ldots, \sigma_s$ to the trusted party:

   (a) For every $i = 1, \ldots, \ell$ and every $j \in \mathcal{J}$, $\mathcal{S}$ receives back a pair $(x_0^{i,j}, x_1^{i,j})$.

   (b) For every $i = 1, \ldots, \ell$, $\mathcal{S}$ receives back the vector $\vec{x}_i = \langle x_{\sigma_i}^{i,1}, \ldots, x_{\sigma_i}^{i,s} \rangle$.

   (c) For every $k \notin \mathcal{J}$, $\mathcal{S}$ receives back χ_k.

6. $\mathcal{S}$ simulates the transfer phase by handing the following messages to $\mathcal{A}$:

   (a) For every $i = 1, \ldots, \ell$ and every $j \in \mathcal{J}$, $\mathcal{S}$ computes $(u_0^{i,j}, w_0^{i,j})$ and $(u_1^{i,j}, w_1^{i,j})$ exactly like the honest sender (it can do this because it knows both $x_0^{i,j}$ and $x_1^{i,j}$).

   (b) For every $i = 1, \ldots, \ell$ and every $k \notin \mathcal{J}$, $\mathcal{S}$ computes $(u_{\sigma_i}^{i,k}, w_{\sigma_i}^{i,k})$ like the honest sender using $x_{\sigma_i}^{i,k}$, and sets $(u_{1-\sigma_i}^{i,k}, w_{1-\sigma_i}^{i,k})$ to be a pair of a random element of $\mathbb{G}$ and a random string of length n.

7. $\mathcal{S}$ simulates the transfer of χ_j values, as follows:

   (a) $\mathcal{S}$ receives the $\tilde{h}_0^j, \tilde{h}_1^j$ values from $\mathcal{A}$ and the witness it sends to ZKPOK functionality. If the witness is not correct, it simulates S aborting, and outputs whatever $\mathcal{A}$ outputs. Otherwise it proceeds to the next step.

   (b) For every $j \notin \mathcal{J}$, $\mathcal{S}$ computes (u_j, w_j) as the honest sender does (it can do this since it knows χ_k.

   (c) For every $j \in \mathcal{J}$, $\mathcal{S}$ sets (u_j, w_j) to be a pair of a random element of $\mathbb{G}$ and a random string of length n.

8. $\mathcal{S}$ outputs whatever $\mathcal{A}$ outputs.

We claim that the output of the ideal execution with $\mathcal{S}$ is *identical* to the output of a real execution with $\mathcal{A}$ and an honest sender. This is due to the fact that the only difference is with respect to the way the tuples $(u_{1-\sigma_i}^{i,k}, w_{1-\sigma_i}^{i,k})$ for $k \notin \mathcal{J}$ are formed, and the way that the tuples (u_j, w_j) $j \in \mathcal{J}$ are formed (namely, these are generated as random elements by the simulator).

Beginning first with the (u_j, w_j) tuples for $j \in \mathcal{J}$, observe that by the way that $\mathcal{S}$ defines $\mathcal{J}$ it holds that $h_1^j = (h_0^j)^y$ for all $j \in \mathcal{J}$. This implies that (g_0, g_1, h_0^j, h_1^j) is a Diffie-Hellman tuple, and $(g_0, g_1, h_0^j, (h_1^j/g_1)$ is *not* a Diffie-Hellman tuple. Now, in the real protocol, S generates (u_j, w_j) by first computing $RAND(g_0, g_1, h_0^j, (h_1^j/g_1))$. Thus, by Claim 6.2, the values (u_j, v_j) are random in $\mathbb{G}$ and so (u_j, w_j) are distributed identically to the way they are generated by $\mathcal{S}$ (i.e., u_j is uniform in $\mathbb{G}$ and w_j is uniform in $\{0,1\}^n$).

Regarding the tuples $(u_{1-\sigma_i}^{i,k}, w_{1-\sigma_i}^{i,k})$ for $k \notin \mathcal{J}$, the same argument holds. Specifically, for $k \notin \mathcal{J}$ we know that $h_1^k \neq (h_0^k)^y$. Thus, for a given tuple $(\tilde{g}_i, \tilde{h}_{i,1}, \ldots, \tilde{h}_{i,s})$ it *cannot* hold that both $[\tilde{g}_i = (g_0)^{r_i}$ and every $\tilde{h}_{i,j} = (h_0^j)^{r_i}]$ and $[\tilde{g}_i = (g_1)^{r_i}$ and every $\tilde{h}_{i,j} = (h_1^j)^{r_i}]$. Since the simulator

set σ_i to be such that $[\tilde{g}_i = (g_{\sigma_i})^{r_i}$ and every $\tilde{h}_{i,j} = (h_{\sigma_i}^j)^{r_i}]$ holds, we have that $[\tilde{g}_i = (g_{1-\sigma_i})^{r_i}$ and every $\tilde{h}_{i,j} = (h_{1-\sigma_i}^j)^{r_i}]$ does *not* hold. This in turn applies that $(g_{1-\sigma_i}, \tilde{g}_i, h_{1-\sigma_i}, \tilde{h}_{i,j})$ is not a Diffie-Hellman tuple, and so the distribution generated by $\mathcal{S}$ by choosing the values $(u_{1-\sigma_i}^{i,k}, w_{1-\sigma_i}^{i,k})$ uniformly is identical, as above.

Case 2 – corrupted sender: We now proceed to the case that $\mathcal{A}$ controls the sender S. We construct a simulator $\mathcal{S}$ as follows:

1. $\mathcal{S}$ computes $g_1 = (g_0)^y$ for a random $y \leftarrow \mathbb{Z}_q$, and hands g_1 internally to $\mathcal{A}$.

2. $\mathcal{S}$ simulates the ideal zero-knowledge proof of knowledge from R to S by handing 1 ("success") to $\mathcal{A}$ as if it came from the ZKPOK ideal functionality.

3. $\mathcal{S}$ plays the same strategy as an honest R with $\mathcal{J} = [s]$ and $\sigma_1 = \sigma_2 = \cdots = \sigma_\ell = 0$. That is:

   (a) For every $j = 1, \ldots, s$, $\mathcal{S}$ chooses a random $\alpha_j \leftarrow \mathbb{Z}_q$ and sets $h_0^j = (g_0)^{\alpha_j}$ and $h_1^j = (g_1)^{\alpha_j}$ (i.e., setting $\mathcal{J} = [s]$). $\mathcal{S}$ hands $(h_0^1, h_1^1, \ldots, h_0^s, h_1^s)$ to $\mathcal{A}$.

   (b) For every $i = 1, \ldots, \ell$, $\mathcal{S}$ computes $\tilde{g}_i = (g_0)^{r_i}$ and $\tilde{h}_{i,1} = (h_0^1)^{r_i}, \ldots, \tilde{h}_{i,s} = (h_0^s)^{r_i}$ and sends the vector $(\tilde{g}_i, \tilde{h}_{i,1}, \ldots, \tilde{h}_{i,s})$ to $\mathcal{A}$.

4. Upon receiving back pairs $(u_0^{i,j}, w_0^{i,j})$ and $(u_1^{i,j}, w_1^{i,j})$ for every $i = 1, \ldots, \ell$ and every $j = 1, \ldots, s$, simulator $\mathcal{S}$ computes $x_0^{i,j} = w_0^{i,j} \oplus H((u_0^{i,j})^{r_i}$ and $x_1^{i,j} = w_1^{i,j} \oplus H((u_1^{i,j})^{r_i/y}$, just like an honest R (for when $j \in \mathcal{J}$).

5. In the simulation of the transfer of χ_j values, $\mathcal{S}$ behaves differently, as follows:

   (a) For every $j = 1, \ldots, s$ it computes $\tilde{h}_0^j = (h_0^j)^{\rho_j}$ and $\tilde{h}_1^j = (h_1^j/g_1)^{\rho_j}$ (as if $j \notin \mathcal{J}$, even though in actuality $j \in \mathcal{J}$ for all j).

   (b) $\mathcal{S}$ simulates the zero-knowledge proof that all the tuples $\{(g_0, g_1, \tilde{h}_0^j, \tilde{h}_1^j)\}_{j=1}^s$ are Diffie-Hellman tuples (even though none of them are) by handing 1 ("success") to $\mathcal{A}$ as if it came from the ZKPOK ideal functionality.

   (c) Upon receiving all the pairs $\{(u_j, w_j)\}_{j=1}^s$ from $\mathcal{A}$, simulator $\mathcal{S}$ computes $\chi_j = w_j \oplus H((u_j)^{\rho_j}$.

6. $\mathcal{S}$ sends all of the vectors pairs $\vec{x}_i = \langle (x_0^{i,1}, x_1^{i,1}), \ldots, (x_0^{i,s}, x_1^{i,s}) \rangle$ for $i = 1, \ldots, \ell$ to the trusted party, as well as all of the values $\chi_1, \ldots, \chi_s$.

7. $\mathcal{S}$ outputs whatever $\mathcal{A}$ outputs, and halts.

There are two main observations regarding the simulation. First, by the way the simulator chooses g_0 and g_1, all of the $(g_0, \tilde{g}_i, h_0^j, \tilde{h}_{i,j})$ and $(g_1, \tilde{g}_i, h_1^j, \tilde{h}_{i,j})$ tuples are Diffie-Hellman tuples. Thus, $\mathcal{S}$ learns all of the pairs $(x_0^{i,j}, x_1^{i,j})$ and these are consistent with the values that the honest receiver would receive in a real execution. Likewise, $\mathcal{S}$ learns all of the values $\chi_1, \ldots, \chi_s$ and these are consistent with the χ_j values that R would receive in a real execution for values of $j \notin \mathcal{J}$. Second, by the Decisional Diffie-Hellman assumption, the output of a simulated execution with $\mathcal{S}$ in the ideal model is indistinguishable from the output of a real execution between $\mathcal{A}$ and an honest receiver. This is due to the fact that the only differences between the real and ideal executions are:

1. The simulator chooses the values $(h_0^1, h_1^1, \ldots, h_0^s, h_1^s)$ as if $\mathcal{J} = [s]$ (and the honest receiver uses the subset $\mathcal{J}$ in its input)

2. The simulator generates the vectors $(\tilde{g}_i, \tilde{h}_{i,1}, \ldots, \tilde{h}_{i,s})$ as if all the receiver inputs are $\sigma_1 = \cdots = \sigma_\ell = 0$ (whereas the honest receiver uses its real input $\sigma_1, \ldots, \sigma_\ell$, and

3. The simulator cheats when it generates the pair $(\tilde{h}_0^j, \tilde{h}_1^j)$ making it so that all tuples $(h_0^j, h_1^j/g_1, \tilde{h}_0^j, \tilde{h}_1^j)$ are Diffie-Hellman tuples even though none of the tuples $(g_0, g_1, \tilde{h}_0^j, \tilde{h}_1^j)$ are.

Nevertheless, these differences are all indistinguishable by the DDH assumption. This is demonstrated formally as in [24]. ∎

6.4 Concrete Efficiency – Cut-and-Choose OT

We count the number of exponentiations, and the bandwidth. We also distinguish between fixed-base exponentiations and regular exponentiations (in most groups, the cost of a fixed-base exponentiation is about a third of a full exponentiation [25, Sec. 14.6.3]). The cost of the zero-knowledge proof of knowledge of discrete log is 9 exponentiations and the exchange of 4 group elements [13], and the cost of the zero-knowledge proof in Step 2b of the transfer phase is $s + 16$ exponentiations overall, and the exchange of 10 group elements; see [24]. Finally, the cost of the s zero-knowledge proofs of knowledge of Step 3b is $12s$ exponentiations and the exchange of $5s$ group elements.

In the setup phase, not counting the zero-knowledge proof of knowledge, R carries out $2s + 1$ fixed-base exponentiations and sends $2s + 1$ group elements. In the transfer phase (excluding the zero-knowledge), R computes $(s + 1) \cdot \ell$ fixed-base exponentiations in Step 2a (these bases are used ℓ times and so for not very small values of ℓ the fixed-base optimization will be significant), and S computes $RAND$ $2s\ell$ times at the cost of $8s\ell$ fixed-base exponentiations. In addition, R sends $(s + 1)\ell$ group elements and S sends $4s\ell$ group elements. In the transfer of χ_j values phase (excluding the zero-knowledge), R computes $2s$ regular exponentiations, and S computes $RAND$ s times at the cost of $4s$ regular exponentiations. In addition, R sends $2s$ group elements and S sends $2s$ group elements. Finally, in the output phase, R computes $(s + |\mathcal{J}|)\ell + s - |\mathcal{J}| \approx \frac{3s}{2} \cdot \ell - \frac{s}{2}$ regular exponentiations. Overall we have:

| Operation	Exact Cost	Approximate Cost
Regular exponentiations	$1.5s\ell + 18.5s + 25$	$1.5s\ell$
Fixed-base exponentiations	$9s\ell + \ell + 2s + 1$	$9s\ell$
Bandwidth (group elements)	$5s\ell + \ell + 11s + 15$	$5s\ell$

Acknowledgements

References

[1] Y. Aumann and Y. Lindell. Security Against Covert Adversaries: Efficient Protocols for Realistic Adversaries. In the *Journal of Cryptology*, 23(2):281–343, 2010.

[2] D. Beaver. Foundations of Secure Interactive Computing. In *CRYPTO'91*, Springer-Verlag (LNCS 576), pages 377–391, 1991.

[3] R. Bendlin, I. Damgård, C. Orlandi and S. Zakarias. Semi-homomorphic Encryption and Multiparty Computation. In *EUROCRYPT 2011*, Springer (LNCS 6632), pages 169–188, 2011.

[4] R. Canetti. Security and Composition of Multiparty Cryptographic Protocols. *Journal of Cryptology*, 13(1):143–202, 2000.

[5] R. Canetti. Universally Composable Security: A New Paradigm for Cryptographic Protocols. In *42nd FOCS*, pages 136–145, 2001. Full version available at `http://eprint.iacr.org/2000/067`.

[6] I. Damgård and C. Orlandi. Multiparty Computation for Dishonest Majority: From Passive to Active Security at Low Cost. In *CRYPTO 2010*, Springer (LNCS 6223), pages 558–576, 2010.

[7] I. Damgård, V. Pastro, N.P. Smart and S. Zakarias. Multiparty Computation from Somewhat Homomorphic Encryption. In *CRYPTO 2012*, Springer (LNCS 7417), pages 643–662, 2012.

[8] Y. Dodis, R. Gennaro, J. Hastad, H. Krawczyk and T. Rabin. Randomness Extraction and Key Derivation Using the CBC, Cascade and HMAC Modes. In *CRYPTO 2004*, Springer (LNCS 3152), pages 494–510, 2004.

[9] T.K. Frederiksen and J.B. Nielsen. Fast and Maliciously Secure Two-Party Computation Using the GPU. *Cryptology ePrint Archive: Report 2013/046*, 2013.

[10] S. Goldwasser and L. Levin. Fair Computation of General Functions in Presence of Immoral Majority. In *CRYPTO'90,* Springer-Verlag (LNCS 537), pages 77–93, 1990.

[11] O. Goldreich. *Foundations of Cryptography: Volume 2 – Basic Applications.* Cambridge University Press, 2004.

[12] O. Goldreich, S. Micali and A. Wigderson. How to Play any Mental Game – A Completeness Theorem for Protocols with Honest Majority. In *19th STOC,* pages 218–229, 1987.

[13] C. Hazay and Y. Lindell. *Efficient Secure Two-Party Protocols: Techniques and Constructions.* Springer, November 2010.

[14] Y. Huang, J. Katz and D. Evans. Efficient Secure Two-Party Computation Using Symmetric Cut-and-Choose. In *CRYPTO 2013* vol. 2, Springer (LNCS 8043), pages 1–17, 2013.

[15] Y. Ishai, J. Kilian, K. Nissim and E. Petrank. Extending Oblivious Transfer Efficiently. In *CRYPTO 2003*, Springer (LNCS 2729), pages 145–161, 2003.

[16] Y. Ishai, M. Prabhakaran and A. Sahai. Founding Cryptography on Oblivious Transfer – Efficiently. In *CRYPTO 2008*, Springer (LNCS 5157), pages 572–591, 2008.

[17] Y. Ishai, M. Prabhakaran and A. Sahai. Secure Arithmetic Computation with No Honest Majority. In *TCC 2009*, Springer (LNCS 5444), pages 294–314, 2009.

[18] S. Jarecki and V. Shmatikov. Efficient Two-Party Secure Computation on Committed Inputs. In *EUROCRYPT 2007*, Springer (LNCS 4515), pages 97–114, 2007.

[19] B. Kreuter, A.Shelat, and C. Shen. Billion-Gate Secure Computation with Malicious Adversaries. In the 21*st USENIX Security Symposium*, 2012.

[20] Y. Lindell, E. Oxman and B. Pinkas. The IPS Compiler: Optimizations, Variants and Concrete Efficiency. In *CRYPTO 2011*, Springer (LNCS 6841), pages 259–276, 2011.

[21] Y. Lindell and B. Pinkas. A Proof of Yao's Protocol for Secure Two-Party Computation. In the *Journal of Cryptology*, 22(2):161–188, 2009.

[22] Y. Lindell and B. Pinkas. An Efficient Protocol for Secure Two-Party Computation in the Presence of Malicious Adversaries. To appear in the *Journal of Cryptology*. (Extended abstract appeared in *EUROCRYPT 2007*, Springer (LNCS 4515), pages 52–78, 2007.)

[23] Y. Lindell and B. Pinkas. Secure Two-Party Computation via Cut-and-Choose Oblivious Transfer. In *Journal of Cryptology*, 25(4):680722, 2012. (Extended abstract appeared in *TCC 2011*, Springer (LNCS 6597), pages 329–346, 2011.)

[24] Y. Lindell and B. Pinkas. Secure Two-Party Computation via Cut-and-Choose Oblivious Transfer. In *TCC 2011*, Springer (LNCS 6597), pages 329–346, 2011

[25] A.J. Menezes, P.C. van Oorschot and S.A. Vanstone. *Handbook of Applied Cryptography*, CRC Press, 2001.

[26] S. Micali and P. Rogaway. Secure Computation. Unpublished manuscript, 1992. Preliminary version in *CRYPTO'91*, Springer-Verlag (LNCS 576), pages 392–404, 1991.

[27] P. Mohassel and B. Riva. Garbled Circuits Checking Garbled Circuits: More Efficient and Secure Two-Party Computation. In *CRYPTO 2013*, Springer (LNCS 8043), pages 36–53, 2013.

[28] J.B. Nielsen, P.S. Nordholt, C. Orlandi and S.Sheshank Burra. A New Approach to Practical Active-Secure Two-Party Computation. In *CRYPTO 2012*, Springer (LNCS 7417), pages 681–700, 2012.

[29] J.B. Nielsen and C. Orlandi. LEGO for Two-Party Secure Computation. In *TCC 2009*, Springer (LNCS 5444), pages 368–386, 2009.

[30] B. Schoenmakers and P. Tuyls. Practical Two-Party Computation Based on the Conditional Gate. In *ASIACRYPT 2004*, Springer (LNCS 3329), pages 119–136, 2004.

[31] A. Shelat, C.H. Shen. Two-Output Secure Computation with Malicious Adversaries. In *EUROCRYPT 2011*, Springer (LNCS 6632), pages 386–405, 2011.

[32] A. Yao. How to Generate and Exchange Secrets. In 27*th FOCS*, pages 162–167, 1986. See [21] for details.

[33] Bristol Cryptography Group. Circuits of Basic Functions Suitable For MPC and FHE. http://www.cs.bris.ac.uk/Research/CryptographySecurity/MPC/.

A The Circuit for Step 7 of Protocol 3.2

A.1 The Original Circuit

A.2   The Optimized Circuit

Free-XOR garbled values here

Independent garbled values on wires from here on

Compare $b_1{}^0||b_1{}^1$ to $b_0||b_1$ bit by bit

XOR — NOT

Compute XOR of all 1-input values with 1-output value

Compare $b_2{}^0||b_2{}^1$ to $b_0||b_1$ bit by bit

XOR — NOT

Compute XOR of all 1-input values with 1-output value

Use the same k_0, k_1 value in all of these wires and thus in all of the ANDs with P_1's input $x_1,...,x_l$ (success for a single $b_i{}^0||b_i{}^1$ gives k_1 which is enough to learn $x_1,...,x_l$

Compare $b_m{}^0||b_m{}^1$ to $b_0||b_1$ bit by bit

XOR — NOT

Compute XOR of all 1-input values with 1-output value

x_1

x_2

x_l

AND

AND

AND

A.3 The Final Optimized Circuit
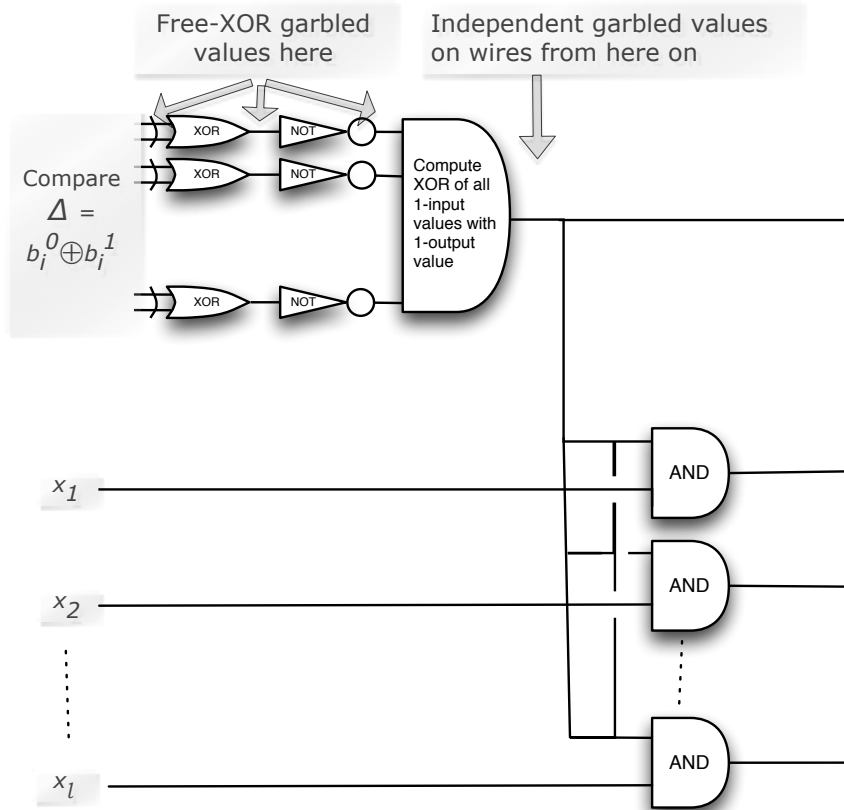
The actual computation of the above circuit can be carried out simply and efficiently, as follows. Let k_0, k_1 be random garbled values, and assign them to the wire that is output from the "Compare Δ" part of the circuit (i.e., the output wire of the gate that computes the XOR of all 1-input values). Then, assuming Δ is of length s, choose s random values $k_1^1, \ldots, k_1^s$ under the constraint that $\oplus_{i=1}^{s} k_1^i = k_1$. Next, in the cut-and-choose oblivious transfer for P_2's inputs, P_1 defines the ith input pair to be (x_0^i, x_1^i) where $x_{\Delta_i}^i = k_1^i$ and $x_{1-\Delta_i}^i$ is random (we denote $\Delta = \Delta_1, \ldots, \Delta_s$). In addition, to the above, k_0 is sent in the clear.

Observe that if P_2 knows Δ and inputs $\Delta_1, \ldots, \Delta_s$ to the oblivious transfer, then it will receive for output $k_1^1, \ldots, k_1^s$ and can compute $k_1 = \oplus_{i=1}^{s} k_1^i$. However, if even one bit input to P_2 is not correct, then k_1 will be completely hidden.

In addition, recall that in the check circuits P_2 receives both inputs to each oblivious transfer. Thus, given Δ (which is sent in the clear *after* the cut-and-choose OT terminates), P_2 can compute $k_1 = \oplus_{i=1}^{s} x_{\Delta_i}^i$. It can then use k_0 sent in the clear, together with this computed k_1 as input to the procedure that checks if the (standard) garbled circuit with the AND gates is correct.

B Definition of Security for Covert Adversaries

For the sake of completeness, and since it is less well known, we provide the definition of security for covert adversaries here.

The high-level idea behind the definition. The definition of security for covert adversaries is based on the standard ideal/real simulation paradigm for malicious adversaries, and provides the guarantee that if the adversary cheats, then it will be caught by the honest parties (with some probability). In order to understand what we mean by this, we have to explain what we mean by "cheating". Loosely speaking, we say that an adversary successfully cheats if it manages to do something that is impossible in the ideal model. Stated differently, successful cheating is behavior that cannot be simulated in the ideal model. Thus, for example, an adversary who learns more about the honest parties' inputs than what is revealed by the output has cheated. In contrast, an adversary who uses pseudorandom coins instead of random coins (where random coins are what are specified in the protocol) has not cheated.

   We begin by informally describing the guarantee provided by this notion. Let $0 < \epsilon \le 1$ be a value (called the *deterrence factor*). Then, any attempt to cheat by a real adversary $\mathcal{A}$ is detected by the honest parties with probability at least ϵ. Thus, provided that ϵ is sufficiently large, an adversary that wishes not to be caught cheating will refrain from *attempting* to cheat, lest it be caught doing so. Clearly, the higher the value of ϵ, the greater the probability adversarial behavior is caught and thus the greater the *deterrent* to cheat. This notion is therefore called security in the presence of covert adversaries with ϵ-deterrent. Note that the security guarantee does not preclude successful cheating. Indeed, if the adversary decides to cheat it may gain access to the other parties' private information or bias the result of the computation. The only guarantee is that if it attempts to cheat, then there is a fair chance that it will be caught doing so. This is in contrast to standard definitions, where absolute privacy and security are guaranteed for the given type of adversary. We remark that by setting $\epsilon = 1$, the definition can be used to capture a requirement that cheating parties are always caught.

The actual definition. We begin by presenting the modified ideal model. In this model, we add new instructions that the adversary can send to the trusted party. Recall that in the standard ideal model, the adversary can send a special abort_i message to the trusted party, in which case the honest party receives abort_i as output. In the ideal model for covert adversaries, the adversary can send the following additional special instructions:

- *Special input* $\mathsf{corrupted}_i$: If the ideal-model adversary sends $\mathsf{corrupted}_i$ instead of an input, the trusted party sends $\mathsf{corrupted}_i$ to the honest party and halts. This enables the simulation of behavior by a real adversary that always results in detected cheating. (It is not essential to have this special input, but it sometimes makes proving security easier.)

- *Special input* cheat_i: If the ideal-model adversary sends cheat_i instead of an input, the trusted party tosses coins and with probability ϵ determines that this "cheat strategy" by P_i was detected, and with probability $1 - \epsilon$ determines that it was not detected. If it was detected, the trusted party sends $\mathsf{corrupted}_i$ to the honest party. If it was not detected, the trusted party hands the adversary the honest party's input and gives the ideal-model adversary the ability to set the output of the honest party to whatever value it wishes. Thus, a cheat_i

input is used to model a protocol execution in which the real-model adversary decides to cheat. However, as required, this cheating is guaranteed to be detected with probability at least ϵ. Note that if the cheat attempt is not detected then the adversary is given "full cheat capability", including the ability to determine the honest party's output.

The idea behind the new ideal model is that given the above instructions, the adversary in the ideal model can choose to cheat, with the caveat that its cheating is guaranteed to be detected with probability at least ϵ. We stress that since the capability to cheat is given through an "input" that is provided to the trusted party, the adversary's decision to cheat must be made before the adversary learns anything (and thus independently of the honest party's input and the output).

We are now ready to present the modified ideal model. Let $\epsilon : \mathbb{N} \to [0, 1]$ be a function. Then, the ideal execution for a function $f : \{0,1\}^* \times \{0,1\}^* \to \{0,1\}^* \times \{0,1\}^*$ with parameter ϵ proceeds as follows:

Inputs: Let x denote the input of party P_1, and let y denote the input of party P_2. The adversary $\mathcal{A}$ also has an auxiliary input z.

Send inputs to trusted party: The honest party P_j sends its received input to the trusted party. The corrupted party P_i, controlled by $\mathcal{A}$, may send a special input signalling abort or attempted cheating (by replacing the input with a special abort_i, $\mathsf{corrupted}_i$ or cheat_i message), send its received input, or send some other input of the same length to the trusted party. This decision is made by $\mathcal{A}$ and may depend on the input value of P_i and the auxiliary input z. Denote the pair of inputs sent to the trusted party by (x', y').

Abort options: If a corrupted party sends abort_i to the trusted party as its input, then the trusted party sends abort_i to the honest party and halts. If a corrupted party sends $\mathsf{corrupted}_i$ to the trusted party as its input, then the trusted party sends $\mathsf{corrupted}_i$ to the honest party and halts.

Attempted cheat option: If a corrupted party sends cheat_i to the trusted party as its input, then the trusted party works as follows:

1. With probability ϵ, the trusted party sends $\mathsf{corrupted}_i$ to the adversary and the honest party.

2. With probability $1 - \epsilon$, the trusted party sends $\mathsf{undetected}$ to the adversary along with the honest party's input. Following this, the adversary sends the trusted party an output value τ of its choice for the honest party. The trusted party then sends τ to P_j as its output (where P_j is the honest party).

If the adversary sent cheat_i, then the ideal execution ends at this point. Otherwise, the ideal execution continues below.

Trusted party sends output to adversary: At this point the trusted party computes $f_1(x', y')$ and $f_2(x', y')$ and sends $f_i(x', y')$ to P_i (i.e., it sends the corrupted party its output).

Adversary instructs trusted party to continue or halt: After receiving its output, the adversary sends either $\mathsf{continue}$ or abort_i to the trusted party. If the trusted party receives $\mathsf{continue}$ then it sends $f_j(x', y')$ to the honest party P_j. Otherwise, if it receives abort_i, it sends abort_i to the honest party P_j.

Outputs: The honest party always outputs the output value it obtained from the trusted party. The corrupted party outputs nothing. The adversary $\mathcal{A}$ outputs any arbitrary (probabilistic polynomial-time computable) function of the initial inputs of the corrupted party, the auxiliary input z, and the value $f_i(x', y')$ obtained from the trusted party.

The output of the honest party and the adversary in an execution of the above ideal model is denoted by $\text{IDEAL}^\epsilon_{f,\mathcal{S}(z),i}(x,y,n)$.

Notice that there are two types of "cheating" here. The first is the classic abort and is used to model "early aborting" due to the impossibility of achieving fairness in general when there is no honest majority. The other type of cheating in this ideal model is more serious for two reasons: first, the ramifications of the cheating are greater (the adversary may learn the honest party's input and may be able to determine its output), and second, the cheating is only guaranteed to be detected with probability ϵ. Nevertheless, if ϵ is high enough, this may serve as a deterrent. We stress that in the ideal model the adversary must decide whether to cheat obliviously of the honest party's input and before it receives any output (and so it cannot use the output to help it decide whether or not it is "worthwhile" cheating). We have the following definition.

Definition B.1 (security – strong explicit cheat formulation [1]): *Let f be a two-party functionality, let π be a protocol, and let $\epsilon : \mathbb{N} \rightarrow [0,1]$ be a function. Protocol π is said to* securely compute f in the presence of covert adversaries with ϵ-deterrent *if for every non-uniform probabilistic polynomial-time adversary $\mathcal{A}$ for the real model, there exists a non-uniform probabilistic polynomial-time adversary $\mathcal{S}$ for the ideal model such that for every $i \in \{1,2\}$:*

$$\left\{ \text{IDEAL}^\epsilon_{f,\mathcal{S}(z),i}(x,y,n) \right\}_{x,y,z,n} \stackrel{\text{c}}{\equiv} \left\{ \text{REAL}_{\pi,\mathcal{A}(z),i}(x,y,n) \right\}_{x,y,z,n}$$

where $x,y,z \in \{0,1\}^*$ under the constraint that $|x| = |y|$, and $n \in \mathbb{N}$.

C   A Discussion on the Cut-and-Choose Parameters of [31]

In [31] it was shown that opening and checking 60% of the circuits is preferable to opening and checking 50% of the circuits, since it provides a lower error of $2^{-0.32s}$ instead of $2^{-0.311s}$. Although this difference may be small, it can still be significant: already for an error rate of 2^{-40} this enables the use of 125 circuits instead of 128, and so constitutes a saving. As we will show now, although this results in a saving in bandwidth, it also increases the work.

In order to see this, note that in every circuit that is checked, P_2 computes 4 symmetric encryptions per gate, in contrast to a single symmetric encryption in circuits that are computed (P_1 pays 4 symmetric encryptions for all circuits in any case). Thus, if half of the s circuits are checked, the overall number of symmetric encryptions by the parties is $4|C| \cdot s + 4|C| \cdot \frac{s}{2} + |C| \cdot \frac{s}{2} = 6.5|C|s$. In contrast, if 60% of the circuits are checked then the overall number of symmetric encryptions is $4|C| \cdot s + 4|C| \cdot \frac{3}{5} \cdot s + |C| \cdot \frac{2}{5} \cdot s = 6.8|C|s$. To make this comparison concrete, for an error level of 2^{-40} one can either open half and use 128 circuits requiring 832 symmetric encryptions per gate, or one can open 60% and use 125 circuits requiring 850 symmetric encryptions per gate.

Thus, although opening 60% means saving on bandwidth, it also involves about a 2% increase in work in the form of symmetric encryptions. The preferred method may depend on the actual network setup and machines being used for the computation. In most cases, the bandwidth seems to be more significant; however, there may be settings where computation is the bottleneck.