

Authenticated Garbling and Communication-Efficient, Constant-Round, Secure Two-Party Computation

Jonathan Katz
University of Maryland
jkatz@cs.umd.edu

Samuel Ranellucci
University of Maryland
George Mason University
samuel@umd.edu

Xiao Wang
University of Maryland
wangxiao@cs.umd.edu

Abstract

We propose a simple and efficient framework for obtaining communication-efficient, constant-round protocols for (malicious) secure two-party computation. Our framework uses a preprocessing phase to generate authentication information for the two parties; in the online phase, this information is used to construct a *single* “authenticated” garbled circuit which is transmitted and evaluated. We also discuss various instantiations of our framework:

- The preprocessing phase can be instantiated efficiently using, e.g., TinyOT. Using this approach with our improvements to TinyOT, we obtain a protocol in which secure evaluation of an AES circuit (at 128-bit computational security and 40-bit statistical security) uses roughly 6 MB of communication in total. Most of the communication is circuit independent. A single execution of our protocol performs even better than the best previous work supporting circuit-independent preprocessing when amortized over 1024 executions.
- If the preprocessing phase is instantiated using the IPS compiler, we obtain a constant-round protocol whose communication complexity is asymptotically as small as a *semi-honest* garbled-circuit protocol in the OT-hybrid model.
- If the preprocessing phase is carried out by a trusted server, we obtain a constant-round protocol whose communication complexity is essentially the same as in the linear-round protocol of Mohassel et al. in the analogous setting.

1 Introduction

There have been incredible advances in the efficiency of protocols for (malicious) secure two-party computation (2PC) over the last decade. One popular approach for designing such protocols is to apply the “cut-and-choose” technique [LP07, SS11, LP11, HKE13, Lin13, Bra13, FJN14, AMPR14] to Yao’s garbled-circuit protocol [Yao86] for constant-round (semi-honest) secure two-party computation. For statistical security $2^{-\rho}$, the best protocols using this paradigm require the parties to generate, send, and evaluate ρ garbled circuits, which is optimal for the cut-and-choose approach. Recently, Wang et al. [WMK16] showed a protocol based on this technique that can securely evaluate an AES circuit (in the single-execution setting with no preprocessing) in only 65 ms.

The above approach incurs significant overhead when large circuits are evaluated precisely because ρ garbled circuits need to be transmitted. In order to mitigate this drawback, recent works have explored secure computation in an *amortized* setting where the same function is evaluated multiple times (on different inputs) [HKK⁺14, LR14, LR15, RR16]. When amortizing over τ

Protocol	Circuit-independent preprocessing	Circuit-dependent preprocessing	Online phase	Storage
Cut-and-choose	—	$O(\mathcal{C} \rho\kappa)$	$O(\mathcal{I} \rho\kappa)$	$O(\mathcal{C} \rho\kappa)$
Amortized	—	$O\left(\mathcal{C} \frac{\rho}{\log \tau}\kappa\right)$	$O\left(\mathcal{I} \frac{\rho}{\log \tau}\kappa\right)$	$O\left(\frac{ \mathcal{C} \rho\kappa}{\log \tau}\right)$
LEGO	$O\left(\frac{ \mathcal{C} \rho\kappa}{\log \tau + \log \mathcal{C} }\right)$	$O(\mathcal{C} \kappa)$	$O(\mathcal{I} \kappa)$	$O\left(\frac{ \mathcal{C} \rho\kappa}{\log \tau + \log \mathcal{C} }\right)$
SPDZ-BMR ¹ [LPSY15]	$O(\mathcal{C} \kappa^2)$	$O(\mathcal{C} \kappa)$	$O(\mathcal{I} \kappa)$	$O(\mathcal{C} \kappa)$
This work with TinyOT	$O\left(\frac{ \mathcal{C} \rho\kappa}{\log \tau + \log \mathcal{C} }\right)$	$O(\mathcal{C} \kappa)$	$ \mathcal{I} \kappa$	$O(\mathcal{C} \kappa)$
This work with IPS ²	$O(\mathcal{C} \kappa)$			

Table 1: **Communication complexity of constant-round 2PC protocols.** $|\mathcal{I}|$ represents the length of the inputs, and $|\mathcal{C}|$ denotes the circuit size. The statistical security parameter is ρ , and $\kappa \geq \rho$ denotes the computational security parameter. τ is the number of executions for protocols in the amortized setting. Storage is expressed as the amount of data to be stored after the preprocessing phase(s). Terms that are independent of the input/circuit size are ignored.

executions, only $O\left(\frac{\rho}{\log \tau}\right)$ garbled circuits are needed per execution. Rindal and Rosulek [RR16] recently reported 6.4 ms for evaluation of an AES circuit, amortized over 1024 executions.

Other techniques for secure two-party computation, with asymptotically better performance than cut-and-choose (without amortization), have also been investigated. The LEGO protocol and subsequent optimizations [NO09, FJN⁺13, FJNT15, HZ15, NST17] are based on a gate-level cut-and-choose approach that can be carried out during a preprocessing phase before the circuit to be evaluated is known. This class of protocols has good asymptotic performance (see Table 1) and small online time; however, the overall cost of the state-of-the-art LEGO implementation [NST17] is still slightly higher than the overall cost of the best protocol based on the cut-and-choose approach applied at the garbled-circuit level.

The Beaver-Micali-Rogaway compiler [BMR90] provides yet another approach to constructing constant-round protocols secure against malicious adversaries. This compiler uses an “outer” secure-computation protocol to generate a garbled circuit. Lindell et al. [LPSY15] apply a similar idea using SPDZ [DPSZ12] as the outer protocol.

Protocols running in a super-constant number of rounds have also been investigated. The TinyOT protocol [NNOB12] adds malicious security to the classical GMW protocol [GMW87] by adding an information-theoretic MAC to the shares held by the parties. TinyOT has smaller communication complexity than the LEGO family of protocols, but—just like the GMW protocol—has round complexity linear in the depth of the circuit being evaluated. (In contrast, all the results cited previously run in a constant number of rounds.) The IPS compiler [IPS08, LOP11] has asymptotic communication complexity (in the OT-hybrid model) proportional to the size of the underlying circuit being evaluated. It, too, has the disadvantage of requiring a number of rounds linear in the depth of the circuit. A more serious drawback is that the concrete complexity of the protocol is unclear, since it has not yet been implemented (and appears quite difficult to implement).

¹Based on [KOS16]; the complexity of circuit-independent preprocessing can be reduced to $O(|\mathcal{C}|\kappa)$ using somewhat homomorphic encryption [DPSZ12], but at the expense of requiring a number of public-key operations proportional to the circuit size.

²In the bit-OT-hybrid model.

In Table 1, we summarize the communication complexity of the various approaches for constructing constant-round 2PC protocols. Following [NST17], we divide execution of protocols into three phases:

- **Function-independent preprocessing.** During this phase, the parties need not know their inputs nor the function to be computed (beyond an upper bound on the number of gates).
- **Function-dependent preprocessing.** In this phase, the parties know what function they will compute, but need not know their inputs.

Often, the first two phases are combined and referred to simply as the *offline phase*.

- **Online phase.** In this phase, two parties evaluate the agreed-upon function on their respective inputs.

Our contribution. We propose a new approach for constant-round 2PC protocols with extremely low communication complexity. At a high level (further details are given in Section 3), our protocol relies on a function-independent preprocessing phase to realize an ideal functionality that we call \mathcal{F}_{Pre} . Following ideas of [NNOB12], we use this preprocessing phase to set up correlated information at the two parties that they can use during the online phase for information-theoretic authentication of different values. In contrast to [NNOB12], however, the parties in our protocol use this information in the online phase to distributively generate a *single* “authenticated” garbled circuit. (Conceptually similar ideas were used by Damgård and Ishai [DI05] in the context of multi-party computation with honest majority, and by Choi et al. [CKMZ14] for three-party computation with dishonest majority.) As in the semi-honest case, this garbled circuit can then be transmitted and evaluated using just one additional round of interaction.

Frederiksen et al. [FJNT15] introduces the notion of interactive garbling schemes and how to use it to construct a 2PC protocol. Our work also constructs garbled circuits interactively, but the underlying techniques and ideas are fundamentally different: in our protocol, garbled circuits are distributed to two parties and authenticated to the evaluator.

Regardless of how we realize \mathcal{F}_{Pre} , our protocol is extremely efficient in the function-dependent preprocessing phase and the online phase. Specifically, compared to a *semi-honest* garbled-circuit protocol, the cost of the function-dependent preprocessing phase of our protocol is only about 30% higher (assuming 128-bit computational security and 40-bit statistical security), and the cost of the online phase is essentially unchanged. The cost of the function-independent preprocessing phase—and thus the cost of the entire protocol—depends on precisely how we realize \mathcal{F}_{Pre} :

- If \mathcal{F}_{Pre} is instantiated using TinyOT, the asymptotic communication complexity of our protocol is as good as in protocols based on LEGO, but with two advantages. First, our protocol has better *concrete* communication complexity (see Section 8), especially in the online phase, and overall cost. Furthermore, the amount of storage needed by our protocol between the preprocessing phase and the online phase is (asymptotically) smaller. The latter is especially important when very large circuits are evaluated (see Table 1). We further improve the concrete efficiency by describing several improvements to TinyOT in Section 6.

Compared to the protocol of Lindell et al. [LPSY15], our protocol is asymptotically more efficient in the function-independent preprocessing phase; more importantly, the concrete efficiency of our protocol is much better since our work is compatible with free-XOR and we

do not suffer from any blowup in the size of the circuit being evaluated. In particular, Lindell et al. require five SPDZ-style multiplications per AND gate of the underlying circuit, while we only need one TinyOT-style AND computation per AND gate.

- When \mathcal{F}_{Pre} is instantiated using the IPS compiler, we obtain what is (to the best of our knowledge) the *first* constant-round protocol with communication complexity $O(|\mathcal{C}|\kappa)$ in the OT-hybrid model. Note that this (asymptotically) matches the communication complexity of *semi-honest* secure two-party computation based on garbled circuits.
- We can also realize \mathcal{F}_{Pre} using a (semi-)trusted server. In that case we obtain a constant-round protocol for server-aided 2PC with total communication $O(|\mathcal{C}|\kappa)$. Previous work in the same model [MOR16] achieves the same communication complexity but with number of rounds proportional to the circuit depth.

2 Notations and Preliminaries

We use κ to denote the computational security parameter (i.e., security should hold against attackers running in time $\approx 2^\kappa$), and ρ for the statistical security parameter (i.e., an adversary should succeed in cheating with probability at most $2^{-\rho}$). We use $=$ to denote equality and $:=$ to denote assignment.

A circuit is represented as a list of gates having the format $(\alpha, \beta, \gamma, T)$, where α and β denote the input-wire indices of the gate, γ denotes the output-wire index of the gate, and $T \in \{\oplus, \wedge\}$ denotes the type of the gate. We use \mathcal{I}_1 to denote the set of input-wire indices for P_A 's input, \mathcal{I}_2 to denote the set of input-wire indices for P_B 's input, \mathcal{W} to denote the set of output-wire indices of all AND gates, and \mathcal{O} to denote the set of output-wire indices of the circuit. We denote the parties running the secure-computation protocol by P_A and P_B .

2.1 Information-theoretic MACs

We use the information-theoretic message authentication codes (IT-MAC) of [NNOB12]. P_A holds a random *global key* $\Delta_A \in \{0, 1\}^\rho$. A bit b known by P_B can be authenticated by having P_A hold a random key $K[b]$ and having P_B hold the corresponding tag $M[b] := K[b] \oplus b\Delta_A$. Symmetrically, P_B holds an independent global key Δ_B ; a bit b known by P_A is authenticated by having P_B hold a random key $K[b]$ and having P_A hold the tag $M[b] := K[b] \oplus b\Delta_B$. We use $[b]_A$ to denote an authenticated bit known to P_A (i.e., $[b]_A$ means that P_A holds $(b, M[b])$ and P_B holds $K[b]$), and $[b]_B$ is defined symmetrically.

Observe that this MAC is XOR-homomorphic: given $[b]_A$ and $[c]_A$, the parties can (locally) compute $[b \oplus c]_A$ by having P_A compute $M[b \oplus c] = M[b] \oplus M[c]$ and P_B compute $K[b \oplus c] := (K[b] \oplus K[c])$.

It is possible to extend the above idea to XOR-shared values by having each party's share be authenticated. That is, say we have a value $\lambda := r \oplus s$, where P_A knows r and P_B knows s . Then by having P_A hold $(r, M[r], K[s])$ and P_B hold $(s, K[r], M[s])$, we end up with an authenticated secret-sharing of λ . It can be observed that this scheme is also XOR-homomorphic.

As described in the Introduction, we use a preprocessing phase that realizes a stateful ideal functionality \mathcal{F}_{Pre} . This functionality, described in Figure 1, is used to set up correlated values between the players along with their corresponding IT-MACs. The functionality chooses uniform global keys (once-and-for-all) for each party, with the malicious party being allowed to choose its global key. Then, when the parties request a random authenticated bit, the functionality generates

Functionality \mathcal{F}_{Pre}

- Upon receiving Δ_A from P_A and *init* from P_B , and assuming no values Δ_A, Δ_B are currently stored, choose uniform $\Delta_B \in \{0, 1\}^\rho$ and store Δ_A, Δ_B . Send Δ_B to P_B .
- Upon receiving (*random*, $r, M[r], K[s]$) from P_A and *random* from P_B , sample uniform $s \in \{0, 1\}$ and set $K[r] := M[r] \oplus r\Delta_B$ and $M[s] := K[s] \oplus s\Delta_A$. Send $(s, M[s], K[r])$ to P_B .
- Upon receiving (*AND*, $(r_1, M[r_1], K[s_1]), (r_2, M[r_2], K[s_2]), r_3, M[r_3], K[s_3]$) from P_A , and (*AND*, $(s_1, M[s_1], K[r_1]), (s_2, M[s_2], K[r_2])$) from P_B , verify that $M[r_i] = K[r_i] \oplus r_i\Delta_B$ and that $M[s_i] = K[s_i] \oplus s_i\Delta_A$ for $i \in \{1, 2\}$ and send *cheat* to P_B if not. Otherwise, set $s_3 := r_3 \oplus ((r_1 \oplus s_1) \wedge (r_2 \oplus s_2))$ and set $K[r_3] := M[r_3] \oplus r_3\Delta_B$ and $M[s_3] := K[s_3] \oplus s_3\Delta_A$. Send $(s_3, M[s_3], K[r_3])$ to P_B .

Figure 1: The preprocessing functionality, assuming P_A is corrupted. (It is defined symmetrically if P_B is corrupted. If neither party is corrupted, the functionality is adapted in the obvious way.)

an authenticated secret sharing of the random bit $r \oplus s$. (The malicious party may choose the “random values” it receives, but note that this does not reveal anything about $r \oplus s$ or the other party’s global key to the adversary.) Finally, the parties may also submit their authenticated shares for two bits; the functionality then computes a (fresh) authenticated share of the AND of those bits. We defer until Section 4.2 a discussion of how \mathcal{F}_{Pre} can be instantiated.

3 Protocol Intuition

We give a high-level overview of the core of our protocol in the \mathcal{F}_{Pre} -hybrid model. Our protocol is based on a garbled circuit that the parties compute in a distributed fashion, where the garbled circuit is “authenticated” in the sense that the circuit generator (P_A in our case) cannot change the logic of the circuit. We describe the intuition behind the garbled circuit we use in several steps.

We begin by reviewing standard garbled circuits. Each wire α of a circuit is associated with a random “mask” $\lambda_\alpha \in \{0, 1\}$ known to P_A . If the true value (i.e., the value when the circuit is evaluated on the parties’ inputs) of that wire is x , then the masked value observed by the circuit evaluator (namely, P_B) on that wire will be $\bar{x} = x \oplus \lambda_\alpha$. Each wire α is also associated with two labels $L_{\alpha,0}$ and $L_{\alpha,1} := L_{\alpha,0} \oplus \Delta$ known to P_A (here we are using the free-XOR technique[KS08]). If the masked bit on that wire is \bar{x} , then P_B learns $L_{\alpha,\bar{x}}$.

Let H be a hash function modeled as a random oracle. The garbled table for, e.g., an and-gate $(\alpha, \beta, \gamma, \wedge)$ is given by:

$\bar{x} = x \oplus \lambda_\alpha$	$\bar{y} = y \oplus \lambda_\beta$	Truth Table	Garbled Table
0	0	$\bar{z}_{00} = (\lambda_\alpha \wedge \lambda_\beta) \oplus \lambda_\gamma$	$H(L_{\alpha,0}, L_{\beta,0}, \gamma, 00) \oplus (\bar{z}_{00}, L_{\gamma,\bar{z}_{00}})$
0	1	$\bar{z}_{01} = (\lambda_\alpha \wedge \bar{\lambda}_\beta) \oplus \lambda_\gamma$	$H(L_{\alpha,0}, L_{\beta,1}, \gamma, 01) \oplus (\bar{z}_{01}, L_{\gamma,\bar{z}_{01}})$
1	0	$\bar{z}_{10} = (\bar{\lambda}_\alpha \wedge \lambda_\beta) \oplus \lambda_\gamma$	$H(L_{\alpha,1}, L_{\beta,0}, \gamma, 10) \oplus (\bar{z}_{10}, L_{\gamma,\bar{z}_{10}})$
1	1	$\bar{z}_{11} = (\bar{\lambda}_\alpha \wedge \bar{\lambda}_\beta) \oplus \lambda_\gamma$	$H(L_{\alpha,1}, L_{\beta,1}, \gamma, 11) \oplus (\bar{z}_{11}, L_{\gamma,\bar{z}_{11}})$

P_B , holding $(\bar{x}, L_{\alpha,\bar{x}})$ and $(\bar{y}, L_{\beta,\bar{y}})$, evaluates this garbled gate by picking the (\bar{x}, \bar{y}) -row and decrypting using the garbled labels it holds, thus obtaining $(\bar{z}, L_{\gamma,\bar{z}})$.

The standard garbled circuit just described ensures security against a malicious P_B , since (in an intuitive sense) P_B learns no information about the true values on any of the wires. Unfortunately, it provides no security against a malicious P_A who can potentially cheat by corrupting rows in the

various garbled tables. One particular attack a malicious P_A can carry out is a *selective-failure* attack. Say, for example, that a malicious P_A corrupts only the $(0,0)$ -row of the garbled table for the gate above, and assume P_B aborts if it detects an error during evaluation. If P_B aborts, then P_A learns that the masked values on the input wires of the gate above were $\bar{x} = \bar{y} = 0$, from which it learns that the true values on those wires were λ_α and λ_β .

The selective-failure attack just mentioned can be prevented if the masks are hidden from P_A . (In that case even if P_A learns the masked wire values as before, it learns nothing about the true wire values.) Since knowledge of the garbled table would leak information about the masks to P_A , the garbled table must be hidden from P_A as well. That is, we now want to set up a situation in which P_A and P_B hold *secret shares* of the garbled table, as follows:

$\bar{x} = x \oplus \lambda_\alpha$	$\bar{y} = y \oplus \lambda_\beta$	Truth Table	P_A 's share of Garbled Table	P_B 's share of Garbled Table
0	0	$\bar{z}_{00} = (\lambda_\alpha \wedge \lambda_\beta) \oplus \lambda_\gamma$	$H(L_{\alpha,0}, L_{\beta,0}, \gamma, 00) \oplus (r_{00}, R_{00} \oplus L_{\gamma, \bar{z}_{00}})$	$(s_{00} = \bar{z}_{00} \oplus r_{00}, R_{00})$
0	1	$\bar{z}_{01} = (\lambda_\alpha \wedge \bar{\lambda}_\beta) \oplus \lambda_\gamma$	$H(L_{\alpha,0}, L_{\beta,1}, \gamma, 01) \oplus (r_{01}, R_{01} \oplus L_{\gamma, \bar{z}_{01}})$	$(s_{01} = \bar{z}_{01} \oplus r_{01}, R_{01})$
1	0	$\bar{z}_{10} = (\bar{\lambda}_\alpha \wedge \lambda_\beta) \oplus \lambda_\gamma$	$H(L_{\alpha,1}, L_{\beta,0}, \gamma, 10) \oplus (r_{10}, R_{10} \oplus L_{\gamma, \bar{z}_{10}})$	$(s_{10} = \bar{z}_{10} \oplus r_{10}, R_{10})$
1	1	$\bar{z}_{11} = (\bar{\lambda}_\alpha \wedge \bar{\lambda}_\beta) \oplus \lambda_\gamma$	$H(L_{\alpha,1}, L_{\beta,1}, \gamma, 11) \oplus (r_{11}, R_{11} \oplus L_{\gamma, \bar{z}_{11}})$	$(s_{11} = \bar{z}_{11} \oplus r_{11}, R_{11})$

Once P_A sends its shares of all the garbled gates, P_B can evaluate the garbled circuit: Given $(\bar{x}, L_{\alpha, \bar{x}})$ and $(\bar{y}, L_{\beta, \bar{y}})$, it picks the appropriate row, decrypts P_A 's share of that row using the garbled labels it holds, and then XORs the result with its own shares of that same row to obtain $(\bar{z}, L_{\gamma, \bar{z}})$.

Informally, the above modification ensures *privacy* against a malicious P_A since (intuitively) the result of any changes P_A introduces will depend on the random masks but be *independent* of P_B 's inputs. However, P_A can still affect *correctness* by, e.g., flipping the masked value in one of the rows of a garbled gate. This can be addressed by adding an information-theoretic MAC on P_A 's share of the masked bit. That is, the shares of the garbled table now take the following form:

$\bar{x} = x \oplus \lambda_\alpha$	$\bar{y} = y \oplus \lambda_\beta$	P_A 's share of Garbled Table	P_B 's share of Garbled Table
0	0	$H(L_{\alpha,0}, L_{\beta,0}, \gamma, 00) \oplus (r_{00}, M[r_{00}], R_{00} \oplus L_{\gamma, \bar{z}_{00}})$	$(s_{00} = \bar{z}_{00} \oplus r_{00}, K[r_{00}], R_{00})$
0	1	$H(L_{\alpha,0}, L_{\beta,1}, \gamma, 01) \oplus (r_{01}, M[r_{01}], R_{01} \oplus L_{\gamma, \bar{z}_{01}})$	$(s_{01} = \bar{z}_{01} \oplus r_{01}, K[r_{01}], R_{01})$
1	0	$H(L_{\alpha,1}, L_{\beta,0}, \gamma, 10) \oplus (r_{10}, M[r_{10}], R_{10} \oplus L_{\gamma, \bar{z}_{10}})$	$(s_{10} = \bar{z}_{10} \oplus r_{10}, K[r_{10}], R_{10})$
1	1	$H(L_{\alpha,1}, L_{\beta,1}, \gamma, 11) \oplus (r_{11}, M[r_{11}], R_{11} \oplus L_{\gamma, \bar{z}_{11}})$	$(s_{11} = \bar{z}_{11} \oplus r_{11}, K[r_{11}], R_{11})$

Once P_A sends its shares of the garbled circuit to P_B , the garbled circuit can be evaluated as before. Now, however, P_B will verify the MAC on P_A 's share of each masked bit that it learns. This limits P_A to only being able to cause P_B to abort; as before, though, any such abort will occur independent of P_B 's actual input.

Efficient realization. Although the above idea is powerful, it still remains to design an efficient protocol that allows the parties to distributively compute shares of a garbled table of the above form even when one of the parties is malicious. One key observation is that P_A 's shares of the wire labels need not be authenticated; in the worst-case, incorrect values used by P_A will cause an *input-independent* abort. Note further that, for example,

$$\begin{aligned}
L_{\gamma, \bar{z}_{00}} &= L_{\gamma,0} \oplus \bar{z}_{00} \Delta_A \\
&= L_{\gamma,0} \oplus (r_{00} \oplus s_{00}) \Delta_A \\
&= L_{\gamma,0} \oplus r_{00} \Delta_A \oplus s_{00} \Delta_A \\
&= (L_{\gamma,0} \oplus r_{00} \Delta_A \oplus K[s_{00}]) \oplus (K[s_{00}] \oplus s_{00} \Delta_A).
\end{aligned}$$

Our next key observation is that if s_{00} is an authenticated bit known to \mathcal{P}_B , then \mathcal{P}_A can locally compute $L_{\gamma,0} \oplus r_{00}\Delta_A \oplus K[s_{00}]$; then the other share $K[s_{00}] \oplus s_{00}\Delta_A$ is just the MAC on s_{00} that \mathcal{P}_B already knows! Thus, we can rewrite the garbled table as follows:

$x \oplus \lambda_\alpha$	$y \oplus \lambda_\beta$	\mathcal{P}_A 's share of Garbled Table	\mathcal{P}_B 's share of Garbled Table
0	0	$H(L_{\alpha,0}, L_{\beta,0}, \gamma, 00) \oplus (r_{00}, M[r_{00}], L_{\gamma,0} \oplus r_{00}\Delta_A \oplus K[s_{00}])$	$(s_{00} = \bar{z}_{00} \oplus r_{00}, K[r_{00}], M[s_{00}])$
0	1	$H(L_{\alpha,0}, L_{\beta,1}, \gamma, 01) \oplus (r_{01}, M[r_{01}], L_{\gamma,0} \oplus r_{01}\Delta_A \oplus K[s_{01}])$	$(s_{01} = \bar{z}_{01} \oplus r_{01}, K[r_{01}], M[s_{01}])$
1	0	$H(L_{\alpha,1}, L_{\beta,0}, \gamma, 10) \oplus (r_{10}, M[r_{10}], L_{\gamma,0} \oplus r_{10}\Delta_A \oplus K[s_{10}])$	$(s_{10} = \bar{z}_{10} \oplus r_{10}, K[r_{10}], M[s_{10}])$
1	1	$H(L_{\alpha,1}, L_{\beta,1}, \gamma, 11) \oplus (r_{11}, M[r_{11}], L_{\gamma,0} \oplus r_{11}\Delta_A \oplus K[s_{11}])$	$(s_{11} = \bar{z}_{11} \oplus r_{11}, K[r_{11}], M[s_{11}])$

(The $\{R_{ij}\}$ values are no longer needed since the $\{s_{ij}\}$ are unknown to \mathcal{P}_A , and that is enough to hide the masks from \mathcal{P}_A .) Shares of the table then become easy to compute in a distributed fashion.

Final optimization. One final optimization is based on the simple observation that the entries in the truth table are linearly dependent. More precisely,

$$\begin{aligned} \bar{z}_{00} &= (\lambda_\alpha \wedge \lambda_\beta) \oplus \lambda_\gamma \\ \bar{z}_{01} &= (\lambda_\alpha \wedge \bar{\lambda}_\beta) \oplus \lambda_\gamma = \bar{z}_{00} \oplus \lambda_\alpha \\ \bar{z}_{10} &= (\bar{\lambda}_\alpha \wedge \lambda_\beta) \oplus \lambda_\gamma = \bar{z}_{00} \oplus \lambda_\beta \\ \bar{z}_{11} &= (\bar{\lambda}_\alpha \wedge \bar{\lambda}_\beta) \oplus \lambda_\gamma = \bar{z}_{01} \oplus \lambda_\beta \oplus 1 \end{aligned}$$

Therefore, in order to jointly compute the above table, the parties just need to get MACs on shares of the masks $\lambda_\alpha, \lambda_\beta, \lambda_\gamma$, and then compute the MACs on shares of the bit $\lambda_\alpha \wedge \lambda_\beta$.

4 Framework for Our Protocols

4.1 Protocol in the \mathcal{F}_{Pre} -Hybrid Model

In Figure 2, we give the complete description of our main protocol in the \mathcal{F}_{Pre} -hybrid model. Note that the calls to \mathcal{F}_{Pre} can be performed in parallel, so the protocol runs in constant rounds. Since we show below that \mathcal{F}_{Pre} can be realized efficiently by constant-round protocols, this gives a protocol (in the plain model) with overall constant round complexity.

Although our protocol calls \mathcal{F}_{Pre} in the function-dependent preprocessing phase, it is easy to push this to the function-independent phase using standard techniques. The protocol can be easily extended to support reactive computation. We leave it as future work to figure out further details.

4.2 Instantiation of \mathcal{F}_{Pre}

In the following, we discuss various ways \mathcal{F}_{Pre} can be instantiated.

TinyOT-based instantiation. We can instantiate \mathcal{F}_{Pre} using TinyOT. (We describe some improvements to the TinyOT protocol in Section 6.) In fact, rather than using TinyOT itself to realize the \mathcal{F}_{Pre} functionality, we can instantiate \mathcal{F}_{Pre} directly based on the $\mathcal{F}_{\text{DEAL}}$ functionality defined in the TinyOT paper. One technical issue is that the $\mathcal{F}_{\text{DEAL}}$ functionality defined there includes a “global key query” for technical reasons. This can be added to our \mathcal{F}_{Pre} functionality without affecting the proof much. We will provide further details in the full version.

IPS-based instantiation. We can use the IPS protocol to realize the \mathcal{F}_{Pre} functionality. In the function-dependent preprocessing phase, we need to produce a sharing of λ_i for each wire i ,

Protocol Π_{2pc}

Inputs: In the function-dependent phase, the parties agree on a circuit for a function $f : \{0, 1\}^{|\mathcal{I}_1|} \times \{0, 1\}^{|\mathcal{I}_2|} \rightarrow \{0, 1\}^{|\mathcal{O}|}$. In the input-processing phase, P_A holds $x \in \{0, 1\}^{|\mathcal{I}_1|}$ and P_B holds $y \in \{0, 1\}^{|\mathcal{I}_2|}$.

Function-independent preprocessing:

1. P_A and P_B send init to \mathcal{F}_{Pre} , which sends Δ_A to P_A and Δ_B to P_B .
2. For each wire $w \in \mathcal{I}_1 \cup \mathcal{I}_2 \cup \mathcal{W}$, parties P_A and P_B send random to \mathcal{F}_{Pre} . In return, \mathcal{F}_{Pre} sends $(r_w, M[r_w], K[s_w])$ to P_A and $(s_w, M[s_w], K[r_w])$ to P_B , where $\lambda_w = s_w \oplus r_w$. P_A also picks a uniform κ -bit string $L_{w,0}$.

Function-dependent preprocessing:

3. For each gate $\mathcal{G} = (\alpha, \beta, \gamma, \oplus)$, P_A computes $(r_\gamma, M[r_\gamma], K[s_\gamma]) := (r_\alpha \oplus r_\beta, M[r_\alpha] \oplus M[r_\beta], K[s_\alpha] \oplus K[s_\beta])$ and $L_{\gamma,0} := L_{\alpha,0} \oplus L_{\beta,0}$. P_B computes $(s_\gamma, M[s_\gamma], K[r_\gamma]) := (s_\alpha \oplus s_\beta, M[r_\beta] \oplus M[r_\alpha], K[r_\alpha] \oplus K[r_\beta])$.

4. Then, for each gate $\mathcal{G} = (\alpha, \beta, \gamma, \wedge)$:

(a) P_A (resp., P_B) sends **(and, $(r_\alpha, M[r_\alpha], K[s_\alpha]), (r_\beta, M[r_\beta], K[s_\beta])$)** (resp., **(and, $(s_\alpha, M[s_\alpha], K[r_\alpha]), (s_\beta, M[s_\beta], K[r_\beta])$)**) to \mathcal{F}_{Pre} . In return, \mathcal{F}_{Pre} sends $(r_\sigma, M[r_\sigma], K[s_\sigma])$ to P_A and $(s_\sigma, M[s_\sigma], K[r_\sigma])$ to P_B , where $s_\sigma \oplus r_\sigma = \lambda_\alpha \wedge \lambda_\beta$.

(b) P_A computes the following locally:

$$\begin{aligned} (r_{\gamma,0}, M[r_{\gamma,0}], K[s_{\gamma,0}]) &:= (r_\sigma \oplus r_\gamma, & M[r_\sigma] \oplus M[r_\gamma], & K[s_\sigma] \oplus K[s_\gamma] &) \\ (r_{\gamma,1}, M[r_{\gamma,1}], K[s_{\gamma,1}]) &:= (r_\sigma \oplus r_\gamma \oplus r_\alpha, & M[r_\sigma] \oplus M[r_\gamma] \oplus M[r_\alpha], & K[s_\sigma] \oplus K[s_\gamma] \oplus K[s_\alpha] &) \\ (r_{\gamma,2}, M[r_{\gamma,2}], K[s_{\gamma,2}]) &:= (r_\sigma \oplus r_\gamma \oplus r_\beta, & M[r_\sigma] \oplus M[r_\gamma] \oplus M[r_\beta], & K[s_\sigma] \oplus K[s_\gamma] \oplus K[s_\beta] &) \\ (r_{\gamma,3}, M[r_{\gamma,3}], K[s_{\gamma,3}]) &:= (r_\sigma \oplus r_\gamma \oplus r_\alpha \oplus r_\beta, & M[r_\sigma] \oplus M[r_\gamma] \oplus M[r_\alpha] \oplus M[r_\beta], & K[s_\sigma] \oplus K[s_\gamma] \oplus K[s_\alpha] \oplus K[s_\beta] \oplus \Delta_A &) \end{aligned}$$

(c) P_B computes the following locally:

$$\begin{aligned} (s_{\gamma,0}, M[s_{\gamma,0}], K[r_{\gamma,0}]) &:= (s_\sigma \oplus s_\gamma, & M[s_\sigma] \oplus M[s_\gamma], & K[r_\sigma] \oplus K[r_\gamma] &) \\ (s_{\gamma,1}, M[s_{\gamma,1}], K[r_{\gamma,1}]) &:= (s_\sigma \oplus s_\gamma \oplus s_\alpha, & M[s_\sigma] \oplus M[s_\gamma] \oplus M[s_\alpha], & K[r_\sigma] \oplus K[r_\gamma] \oplus K[r_\alpha] &) \\ (s_{\gamma,2}, M[s_{\gamma,2}], K[r_{\gamma,2}]) &:= (s_\sigma \oplus s_\gamma \oplus s_\beta, & M[s_\sigma] \oplus M[s_\gamma] \oplus M[s_\beta], & K[r_\sigma] \oplus K[r_\gamma] \oplus K[r_\beta] &) \\ (s_{\gamma,3}, M[s_{\gamma,3}], K[r_{\gamma,3}]) &:= (s_\sigma \oplus s_\gamma \oplus s_\alpha \oplus s_\beta \oplus 1, & M[s_\sigma] \oplus M[s_\gamma] \oplus M[s_\alpha] \oplus M[s_\beta], & K[r_\sigma] \oplus K[r_\gamma] \oplus K[r_\alpha] \oplus K[r_\beta] &) \end{aligned}$$

(d) P_A computes $L_{\alpha,1} := L_{\alpha,0} \oplus \Delta_A$ and $L_{\beta,1} := L_{\beta,0} \oplus \Delta_A$, and then sends the following to P_B :

$$\begin{aligned} G_{\gamma,0} &:= H(L_{\alpha,0}, L_{\beta,0}, \gamma, 0) \oplus (r_{\gamma,0}, M[r_{\gamma,0}], L_{\gamma,0} \oplus K[s_{\gamma,0}] \oplus r_{\gamma,0} \Delta_A) \\ G_{\gamma,1} &:= H(L_{\alpha,0}, L_{\beta,1}, \gamma, 1) \oplus (r_{\gamma,1}, M[r_{\gamma,1}], L_{\gamma,0} \oplus K[s_{\gamma,1}] \oplus r_{\gamma,1} \Delta_A) \\ G_{\gamma,2} &:= H(L_{\alpha,1}, L_{\beta,0}, \gamma, 2) \oplus (r_{\gamma,2}, M[r_{\gamma,2}], L_{\gamma,0} \oplus K[s_{\gamma,2}] \oplus r_{\gamma,2} \Delta_A) \\ G_{\gamma,3} &:= H(L_{\alpha,1}, L_{\beta,1}, \gamma, 3) \oplus (r_{\gamma,3}, M[r_{\gamma,3}], L_{\gamma,0} \oplus K[s_{\gamma,3}] \oplus r_{\gamma,3} \Delta_A) \end{aligned}$$

Input processing:

5. For each $w \in \mathcal{I}_1$, P_A sends $(r_w, M[r_w])$ to P_B , who checks that $(r_w, M[r_w], K[r_w])$ is valid. P_B then sends $y_w \oplus \lambda_w := s_w \oplus y_w \oplus r_w$ to P_A . Finally, P_A sends $L_{w, y_w \oplus \lambda_w}$ to P_B .
6. For each $w \in \mathcal{I}_2$, P_B sends $(s_w, M[s_w])$ to P_A , who checks that $(s_w, M[s_w], K[s_w])$ is valid. P_A then sends $x_w \oplus \lambda_w := s_w \oplus x_w \oplus r_w$ and $L_{w, x_w \oplus \lambda_w}$ to P_B .

Circuit evaluation:

7. P_B evaluates the circuit in topological order. For each gate $\mathcal{G} = (\alpha, \beta, \gamma, T)$, P_B initially holds $(z_\alpha \oplus \lambda_\alpha, L_{\alpha, z_\alpha \oplus \lambda_\alpha})$ and $(z_\beta \oplus \lambda_\beta, L_{\beta, z_\beta \oplus \lambda_\beta})$, where z_α, z_β are the underlying values of the wires.

(a) If $T = \oplus$, P_B computes $z_\gamma \oplus \lambda_\gamma := (z_\alpha \oplus \lambda_\alpha) \oplus (z_\beta \oplus \lambda_\beta)$ and $L_{\gamma, z_\gamma \oplus \lambda_\gamma} := L_{\alpha, z_\alpha \oplus \lambda_\alpha} \oplus L_{\beta, z_\beta \oplus \lambda_\beta}$.

(b) If $T = \wedge$, P_B computes $i := 2(z_\alpha \oplus \lambda_\alpha) + (z_\beta \oplus \lambda_\beta)$ followed by $(r_{\gamma,i}, M[r_{\gamma,i}], L_{\gamma,0} \oplus K[s_{\gamma,i}] \oplus r_{\gamma,i} \Delta_A) := G_{\gamma,i} \oplus H(L_{\alpha, z_\alpha \oplus \lambda_\alpha}, L_{\beta, z_\beta \oplus \lambda_\beta}, \gamma, i)$. Then P_B checks that $(r_{\gamma,i}, M[r_{\gamma,i}], K[r_{\gamma,i}])$ is valid and, if so, computes $z_\gamma \oplus \lambda_\gamma := (s_{\gamma,i} \oplus r_{\gamma,i})$ and $L_{\gamma, z_\gamma \oplus \lambda_\gamma} := (L_{\gamma,0} \oplus K[s_{\gamma,i}] \oplus r_{\gamma,i} \Delta_A) \oplus M[s_{\gamma,i}]$.

Output determination:

8. For each $w \in \mathcal{O}$, P_A sends $(r_w, M[r_w])$ to P_B , who checks that $(r_w, M[r_w], K[r_w])$ is valid. If so, P_B computes $z_w := (\lambda_w \oplus z_w) \oplus r_w \oplus s_w$.

Figure 2: Our main protocol in the \mathcal{F}_{Pre} -hybrid model.

and a sharing of $\lambda_\sigma = (\lambda_\alpha \wedge \lambda_\beta) \oplus \lambda_\gamma$ for each and-gate $(\alpha, \beta, \gamma, \wedge)$. These can be realized by a constant-depth circuit with $O((\kappa + \rho) \cdot |C|)$ gates. To evaluate a circuit of depth d and size

ℓ , the IPS protocol has communication complexity $O(\ell) + \text{poly}(\kappa, d, \log \ell)$ and $O(d)$ rounds of communication. When applied to our setting, this translates to a communication complexity of $O((\kappa + \rho) \cdot |C|) + \text{poly}(\kappa, \log |C|)$; for sufficiently large circuit size, the leading term is $O((\kappa + \rho) \cdot |C|)$. By using the IPS protocol we thus gain asymptotic improvements in communication complexity in the OT-hybrid model.

Using a trusted server. It is straightforward to instantiate \mathcal{F}_{Pre} using a trusted server. By applying the technique of Mohassel et al. [MOR16], the offline phase can also be decoupled from the identity of other party; we refer to their paper for further details.

5 Proof

5.1 Proof Intuition

The intuition of the protocol in Section 3 also provides some intuition of the security of the protocol. Here we provide more intuition on some key issues.

Selective-failure type attack. In most garbled circuit protocols, selective failure attacks can be launched on the input (by corrupting some garbled keys sent to oblivious transfer), as well as some internal wires (by corrupting some garbled rows in the garbled table). Input selective failure attack is usually prevented using an XOR-tree, while internal wire selective failure is prevented by various of ways, for example cut-and-choose and input recovery.

We argue that such selective failure attacks launched by P_A does not help P_A learn any information in our protocol. The key observation is that all wires, including input wires and internal wires are masked with some random masks (λ values) not known to P_A . Therefore the best that P_A could learn is masked wire values, which appears random to P_A .

Correctness of the garbled circuit. Note that the garbled circuit in our protocol is not guaranteed to be computed correctly (This does not lead to an attack as explained above). However P_B is still able to evaluate the circuit if P_B does not abort. The key reason is that all permutation bits in the truth table are masked. Therefore, P_A cannot change the logic of the garbled table without breaking an IT-MAC.

5.2 The Main Proof

Theorem 5.1. *The protocol in Figure 2 securely instantiate $\mathcal{F}_{2\text{pc}}$ in the \mathcal{F}_{Pre} -hybrid model with security $\text{negl}(\kappa)$*

Proof. We consider separately the case where P_A or P_B is malicious.

Malicious P_A . Let \mathcal{A} be an adversary corrupting P_A . We construct a simulator \mathcal{S} that runs \mathcal{A} as a subroutine and plays the role of P_A in the ideal world involving an ideal functionality \mathcal{F} evaluating f . \mathcal{S} is defined as follows.

- 1-4 \mathcal{S} interacts with \mathcal{A} acting as an honest P_B .
- 5 \mathcal{S} interacts with \mathcal{A} acting as an honest P_B using input $y = 0$.
- 6 \mathcal{S} interacts with \mathcal{A} acting as an honest P_B . For each wire $w \in \mathcal{I}_1$ \mathcal{S} receives $x_w \oplus \lambda_w$ and computes $x_w = (x_w \oplus \lambda_w) \oplus r_w \oplus s_w$. \mathcal{S} sends x to $\mathcal{F}_{2\text{pc}}$.

7-8 \mathcal{S} interacts with \mathcal{A} acting as an honest P_B . If an honest P_B would abort, \mathcal{S} outputs whatever \mathcal{A} outputs and aborts; otherwise \mathcal{S} sends continue to \mathcal{F}_{2pc} .

We now show that the joint distribution over the outputs of \mathcal{A} and the honest P_B in the real world is indistinguishable from the joint distribution over the outputs of \mathcal{S} and P_B in the ideal world.

Hybrid₁. Same as the hybrid-world protocol, where \mathcal{S} plays the role of an honest P_B using the actual input y .

Hybrid₂. Same as **Hybrid₁**, except that in step 6, for each wire $w \in \mathcal{I}_1$ \mathcal{S} receives $x_w \oplus \lambda_w$ and computes $x_w = (x_w \oplus \lambda_w) \oplus r_w \oplus s_w$. \mathcal{S} sends x to \mathcal{F}_{2pc} . If an honest P_B would abort, \mathcal{S} outputs whatever \mathcal{A} outputs and aborts; otherwise \mathcal{S} sends continue to \mathcal{F}_{2pc} .

The view of the two **Hybrids** are exactly the same. According to Lemma 5.1, P_B will learn the same output in both **Hybrids** with all but negligible probability.

Hybrid₃. Same as **Hybrid₂**, except that \mathcal{S} computes $\{s_w\}_{w \in \mathcal{I}_2}$ as follows: \mathcal{S} first randomly pick $\{u_w\}_{w \in \mathcal{I}_2}$, and then computes $s_w := u_w \oplus y_w$.

The two **Hybrids** have exactly the same view.

Hybrid₄. Same as **Hybrid₃**, except that \mathcal{S} uses $y = 0$ as inputs throughout the protocol.

Note that although the distribution of y in **Hybrid₃** and **Hybrid₄** are different, the distribution of $s_w \oplus y_w$ are exactly the same. The view of the two **Hybrids** are therefore the same, we will show that P_B aborts with the same probability in two **Hybrids**.

Observe that the only place where P_B 's abort can possibly depends on y is in step 7(b). However, this abort depends on which row is selected to decrypt, that is the value of $\lambda_\alpha \oplus z_\alpha$ and $\lambda_\beta \oplus z_\beta$, which are chosen independently random in both **Hybrids**.

As **Hybrid₄** is the ideal-world execution, this completes the proof for a malicious P_A .

Malicious P_B . Let \mathcal{A} be an adversary corrupting P_B . We construct a simulator \mathcal{S} that runs \mathcal{A} as a subroutine and plays the role of P_B in the ideal world involving an ideal functionality \mathcal{F} evaluating f . \mathcal{S} is defined as follows.

- 1-4 \mathcal{S} interacts with \mathcal{A} acting as an honest P_A and plays the functionality of \mathcal{F}_{Pre} . If an honest P_A would abort, \mathcal{S} output whatever \mathcal{A} outputs and aborts.
- 5 \mathcal{S} interacts with \mathcal{A} acting as an honest P_A , receives $y_w \oplus \lambda_w$ from \mathcal{A} , and computes $y_w := (y_w \oplus \lambda_w) \oplus s_w \oplus r_w$, where s_w, r_w are values \mathcal{S} used when playing the role of \mathcal{F}_{Pre} . \mathcal{S} sends y to \mathcal{F}_{2pc} , which sends $z = f(x, y)$ to \mathcal{S} .
- 6 \mathcal{S} interacts with \mathcal{A} acting as an honest P_A using input $x = 0$. If an honest P_A would abort, \mathcal{S} output whatever \mathcal{A} outputs and aborts.
- 8 \mathcal{S} computes $z' = f(0, y)$. For each $w \in \mathcal{O}$, if $z'_w = z_w$, \mathcal{S} sends $(r_w, M[r_w])$; otherwise, \mathcal{S} sends $(r_w \oplus 1, M[r_w] \oplus \Delta_B)$.

We now show that the joint distribution over the outputs of \mathcal{A} and the honest P_A in the real world is indistinguishable from the joint distribution over the outputs of \mathcal{S} and P_A in the ideal world.

Hybrid₁. Same as the hybrid-world protocol, where \mathcal{S} plays the role of an honest P_A using the actual input x .

Hybrid₂. Same as **Hybrid₁**, except that, in step 5, \mathcal{S} receives $y_w \oplus \lambda_w$ from \mathcal{A} , and computes $y_w := (y_w \oplus \lambda_w) \oplus s_w \oplus r_w$, where s_w, r_w are values \mathcal{S} used when playing the role of \mathcal{F}_{Pre} . \mathcal{S} then sends y to $\mathcal{F}_{2\text{PC}}$, and receives $z = f(x, y)$.

P_A does not have output; further the view of \mathcal{A} does not change between two **Hybrids**.

Hybrid₃. Same as **Hybrid₂**, except that in step 6, \mathcal{S} uses $x = 0$ as input and in step 8, \mathcal{S} computes $z' = f(0, y)$. For each $w \in \mathcal{O}$, if $z'_w = z_w$, \mathcal{S} sends $(r_w, \mathsf{M}[r_w])$; otherwise, \mathcal{S} sends $(r_w \oplus 1, \mathsf{M}[r_w] \oplus \Delta_B)$.

\mathcal{A} has no knowledge of r_w , therefore r_w and $r_w \oplus 1$ are indistinguishable.

Note that since \mathcal{S} uses different values for x between two **Hybrids**, we also need to show that the garbled rows P_B opened are indistinguishable between two **Hybrids**. According to Lemma 5.2, P_B is able to open only one garble rows in each garbled table $G_{\gamma,i}$. Therefore, given that $\{\lambda_w\}_{w \in \mathcal{I}_1 \cup \mathcal{W}}$ values are not known to P_B , masked values and garbled keys are indistinguishable between two **Hybrids**.

As **Hybrid₃** is the ideal-world execution, the proof is complete. \square

Lemma 5.1. *Consider an \mathcal{A} corrupting P_A and denote $x_w := (x_w \oplus \lambda_w) \oplus s_w \oplus r_w$, where $x_w \oplus \lambda_w$ is the value \mathcal{A} sent to P_B , s_w, r_w are the values from \mathcal{F}_{Pre} . With probability all but negligible, P_B either aborts or learns $z = f(x, y)$.*

Proof. Define z_w^* as the correct wire values computed using x defined above and y , z_w as the actually wire values P_B holds in the evaluation.

We will first show that P_B learns $\{z_w^* \oplus \lambda_w = z_w^* \oplus \lambda_w\}_{w \in \mathcal{O}}$ by induction on topology of the circuit.

Base step: It is obvious that $\{z_w^* \oplus \lambda_w = z_w \oplus \lambda_w\}_{w \in \mathcal{I}_1 \cup \mathcal{I}_2}$, unless \mathcal{A} is able to forge an IT-MAC.

Induction step: Now we show that for a gate $(\alpha, \beta, \gamma, T)$, if P_B has $\{z_w^* \oplus \lambda_w = z_w \oplus \lambda_w\}_{w \in \{\alpha, \beta\}}$, then P_B also obtains $z_\gamma^* \oplus \lambda_\gamma = z_\gamma \oplus \lambda_\gamma$.

- $T = \oplus$: It is true according to the following: $z_\gamma^* \oplus \lambda_\gamma = (z_\alpha^* \oplus \lambda_\alpha) \oplus (z_\beta^* \oplus \lambda_\beta) = (z_\alpha \oplus \lambda_\alpha) \oplus (z_\beta \oplus \lambda_\beta) \oplus z_\gamma \oplus \lambda_\gamma$
- $T = \wedge$: According to the protocol, P_B will open the garbled row defined by $i := 2(z_\alpha \oplus \lambda_\alpha) + (z_\beta \oplus \lambda_\beta)$. If P_B learns $z_\gamma \oplus \lambda_\gamma \neq z_\gamma^* \oplus \lambda_\gamma$, then it means that P_B learns $r_{\gamma,i}^* \neq r_{\gamma,i}$. However, this would mean that \mathcal{A} forge a valid IT-MAC, happening with negligible probability.

Now we know that P_B learns correct masked output. P_B can therefore learn correct output $f(x, y)$ unless \mathcal{A} is able to flip $\{r_w\}_{w \in \mathcal{O}}$, which, again, happens with negligible probability. \square

Lemma 5.2. *Consider an \mathcal{A} corrupting P_B , with negligible, probability, P_B learns both garbled keys for some wire.*

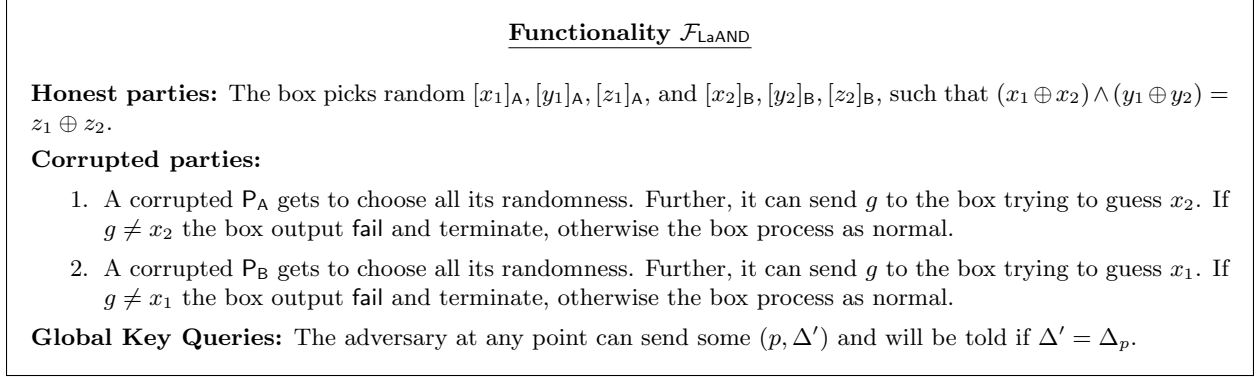


Figure 3: Functionality $\mathcal{F}_{\text{LaAND}}$ for leaky AND triple generation.

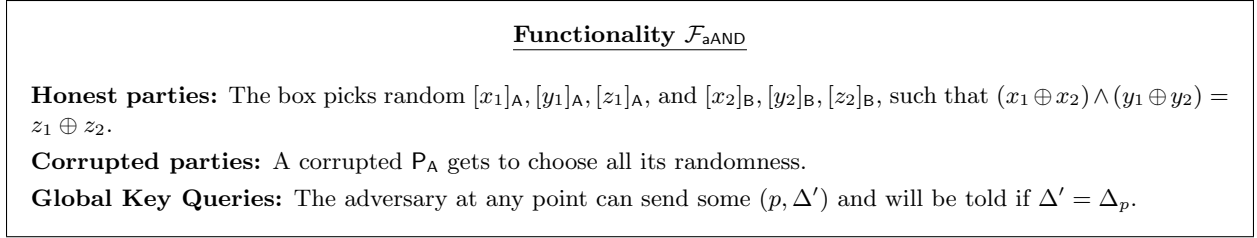


Figure 4: Functionality $\mathcal{F}_{\text{aAND}}$ for generating AND triples

Proof. The proof is very similar to the proof of security for garbled circuits in the semi-honest setting.

Base step: P_B can only learn one garbled keys for each input wire, since P_A only sends one garbled wire, and P_B cannot learn Δ_A in the protocol.

Induction step: It is obvious that P_B cannot learn the other label for an XOR gate and we will focus on AND gates.

Note that P_B only learns one garbled keys for input wire α and β . However, each row is encrypted using different combinations of $\{L_{\alpha,b}\}_{b \in \{0,1\}}$ and $\{L_{\beta,b}\}_{b \in \{0,1\}}$. In order for P_B to open two rows in the garbled table, P_B needs to learn both garbled keys for some input wire, which contradict with assumptions in the induction step. \square

6 Improved TinyOT protocol

In this section, we describe an improvement to the TinyOT protocol. For a bucket size of $B = \frac{\rho}{\log|\mathcal{C}|} + 1$, the original protocol requires $14B + 2$ authenticated bits for each AND gate. In the following, we will introduce an improved version where only $6B$ authenticated bits are needed for each AND gate. For a circuit of size 2^{20} , with $\rho = 40$, this is an improvement of $2.4\times$.

Assuming that two parties hold $[x_1]_A, [y_1]_A, [x_2]_B, [y_2]_B$. In the original TinyOT protocol, to compute $(x_1 \oplus x_2)(y_1 \oplus y_2)$, P_A and P_B compute $[x_1y_1]_A, [x_2y_2]_B, [x_1y_2 + r]_A$ and $[x_2y_1 + r]_B$ separately, with some random $r \in \{0, 1\}$, using various authenticated constructions proposed in their paper. Computing each entry separately incurs a lot of unnecessary cost. We observe that it is possible to compute a whole AND gate directly. Similar to the original TinyOT protocol, we

Protocol Π_{LaAND}

1. P_A and P_B obtain random authenticated bits $[x_1]_A, [y_1]_A, [z_1]_A, [x_2]_B, [y_2]_B, [r]_B$.
2. P_A parses $K[x_2] := [x_2]_B, K[y_2] := [y_2]_B$ and sends the following four bits to P_B .

$$G_{0,0} := \text{Lsb}(H(K[x_2], K[y_2])) \oplus (0 \oplus x_1) \wedge (0 \oplus y_1) \oplus z_1$$

$$G_{1,0} := \text{Lsb}(H(K[x_2] \oplus \Delta_A, K[y_2])) \oplus (1 \oplus x_1) \wedge (0 \oplus y_1) \oplus z_1$$

$$G_{0,1} := \text{Lsb}(H(K[x_2], K[y_2] \oplus \Delta_A)) \oplus (0 \oplus x_1) \wedge (1 \oplus y_1) \oplus z_1$$

$$G_{1,1} := \text{Lsb}(H(K[x_2] \oplus \Delta_A, K[y_2] \oplus \Delta_A)) \oplus (1 \oplus x_1) \wedge (1 \oplus y_1) \oplus z_1$$
3. P_B parses $(x_2, M[x_2]) := [x_2]_B, (y_2, M[y_2]) := [y_2]_B$ and computes $z_2 := \text{Lsb}(H(M[x_2], M[y_2])) \oplus G_{x_2, y_2}$. P_B announces $d := r \oplus z_2$ to P_A . Two parties compute $[z_2]_B = [r]_B \oplus d$.
4. P_B checks the correctness as follows:
 - (a) P_B computes:

$$T_0 := H(K[x_1], K[z_1] \oplus z_2 \Delta_B)$$

$$U_0 := T_0 \oplus H(K[x_1] \oplus \Delta_B, K[y_1] \oplus K[z_1] \oplus (y_2 \oplus z_2) \Delta_B)$$

$$T_1 := H(K[x_1], K[y_1] \oplus K[z_1] \oplus (y_2 \oplus z_2) \Delta_B)$$

$$U_1 := T_1 \oplus H(K[x_1] \oplus \Delta_B, K[z_1] \oplus z_2 \Delta_B)$$
 - (b) P_B sends U_{x_2} to P_A .
 - (c) P_A randomly picks a κ -bit string R and computes

$$V_0 := H(M[x_1], M[z_1]) \qquad V_1 := H(M[x_1], M[z_1] \oplus M[y_1])$$

$$W_{0,0} := H(K[x_2] \oplus V_0 \oplus R) \qquad W_{0,1} := H(K[x_2] \oplus \Delta_A \oplus V_1 \oplus R)$$

$$W_{1,0} := H(K[x_2] \oplus V_1 \oplus U \oplus R) \qquad W_{1,1} := H(K[x_2] \oplus \Delta_A \oplus V_0 \oplus U \oplus R)$$
 - (d) P_A sends $W_{x_1,0}, W_{x_1,1}$ to P_B and sends R to \mathcal{F}_{EQ} .
 - (e) P_B computes $R' := W_{x_1, x_2} \oplus H(M[x_2]) \oplus T_{x_2}$ and sends R' to \mathcal{F}_{EQ} .
5. P_A checks the correctness as follows:
 - (a) P_A computes:

$$T_0 := H(K[x_2], K[z_2] \oplus z_1 \Delta_A)$$

$$U_0 := T_0 \oplus H(K[x_2] \oplus \Delta_A, K[y_2] \oplus K[z_2] \oplus (y_1 \oplus z_1) \Delta_A)$$

$$T_1 := H(K[x_2], K[y_2] \oplus K[z_2] \oplus (y_1 \oplus z_1) \Delta_A)$$

$$U_1 := T_1 \oplus H(K[x_2] \oplus \Delta_A, K[z_2] \oplus z_1 \Delta_A)$$
 - (b) P_A sends U_{x_1} to P_B .
 - (c) P_B randomly picks a κ -bit string R and computes

$$V_0 := H(M[x_2], M[z_2]) \qquad V_1 := H(M[x_2], M[z_2] \oplus M[y_2])$$

$$W_{0,0} := H(K[x_1] \oplus V_0 \oplus R) \qquad W_{0,1} := H(K[x_1] \oplus \Delta_B \oplus V_1 \oplus R)$$

$$W_{1,0} := H(K[x_1] \oplus V_1 \oplus U \oplus R) \qquad W_{1,1} := H(K[x_1] \oplus \Delta_B \oplus V_0 \oplus U \oplus R)$$
 - (d) P_B sends $W_{x_2,0}, W_{x_2,1}$ to P_A and sends R to \mathcal{F}_{EQ} ,
 - (e) P_A computes $R' := W_{x_2, x_1} \oplus H(M[x_1]) \oplus T_{x_1}$ and sends R' to \mathcal{F}_{EQ} .

Figure 5:

propose a “leaky AND” protocol (Π_{LaAND}), where adversary is allowed to perform selective failure attack on one input, and later use bucketing and combining to eliminate such leakage (Π_{aAND}). In the following, we will first discuss intuition of the protocol. The full protocol description is in Figure 5 and Figure 6.

Compute the triple in the honest case. The first step of the protocol is to generate the triple securely assuming that both parties are honest. Since x_1, y_1, z_1, x_2, y_2 are all random, we just need

Protocol Π_{aAND}

1. P_A and P_B call $\mathcal{F}_{\text{LaAND}}$ $\ell' = \ell B$ times and obtains $\{[x_1^i]_A, [y_1^i]_A, [z_1^i]_A, [x_2^i]_B, [y_2^i]_B, [z_2^i]_B\}_{i=1}^{\ell'}$.
 2. P_A and P_B randomly partition all objects into ℓ buckets, each with B objects.
 3. For each bucket, two parties combine B Leaky ANDs into one non-leaky AND. To combine two leaky ANDs, namely $([x_1^i]_A, [y_1^i]_A, [z_1^i]_A, [x_2^i]_B, [y_2^i]_B, [z_2^i]_B)$ and $[x_1'']_A, [y_1'']_A, [z_1'']_A, [x_2'']_B, [y_2'']_B, [z_2'']_B$
 - (a) Two parties reveal $d' := y_1^i \oplus y_1'', d'' = y_2^i \oplus y_2''$ with their MAC checked, and compute $d := d' \oplus d''$.
 - (b) Set $[x_1]_A := [x_1^i]_A \oplus [x_1'']_A, [x_2]_B := [x_2^i]_B \oplus [x_2'']_B, [y_1]_A := [y_1^i]_A, [y_2]_A := [y_2^i]_A, [z_1]_A := [z_1^i]_A \oplus [z_1'']_A \oplus d[x_1'']_A, [z_2]_B := [z_2^i]_B \oplus [z_2'']_B \oplus d[x_2'']_B$.
- Two parties iterate all B leaky objects, by taking the resulted object and combine with the next element.

Figure 6: Protocol Π_{aAND} instantiating $\mathcal{F}_{\text{aAND}}$.

P_B to learn $z_2 = (x_1 \oplus x_2) \wedge (y_1 \oplus y_2) \oplus z_1$. Our idea is to generate a garbled table for AND. We observe that if we treat $K[x_2], K[y_2]$ as zero garbled label, then $M[x_2], M[y_2]$ are garbled labels representing the underlying values, that is we do not need oblivious transfer to let P_B obtain the label. Further, authenticity is not needed in our case, which means we do not need P_B to learn the whole label, as long as P_B learns the output. Inspired by these, our construction only requires 4 bits in order for P_B to learn z_2 (step 1 to 3 in Figure 5).

Verifying the correctness. The above steps are not enough for malicious security: a malicious P_A can cheat by sending incorrect garbled tables and a malicious P_B can announce a incorrect d in step 3. Therefore, both parties needs to check the correctness of the output. In the protocol, we designed a verification protocol that check the correctness while allowing a malicious party to perform a selective failure attack on x values.

The initial idea was to adopt the check from TinyOT to our case. If $x_2 \oplus x_1 = 0$, then we want to check that $z_2 = z_1$; if $x_2 \oplus x_1 = 1$, then to check $y_1 \oplus z_1 = y_2 \oplus z_2$. However, an obvious problem is that no party knows the value of $x_1 \oplus x_2$. To solve this problem, when P_B checks the correctness, we let P_B construct the checking depending on the value of x_2 . P_A will perform the checking twice, as if x_2 is 0 and 1.

For example, using notations in the protocol, when $x_1 = 0$, P_A computes V_0, V_1 . P_A and P_B should have performed an equality check between V_{x_2} and T_{x_2} . All different cases (depending on the value of x_1 and x_2) are summarized in the following table.

	$x_1 = 0$	$x_1 = 1$
$x_2 = 0$	$V_0 = T_0$	$V_0 \oplus U_0 = T_0$
$x_2 = 1$	$V_1 = T_1$	$V_1 \oplus U_1 = T_1$

However, P_A should not learn x_2 , while P_B should not learn $V_{1 \oplus x_2}$. One idea is to let P_A “encrypt” the response (V_0, V_1) such that P_B can only learn the response for the value of x_2 (V_{x_2}), then P_B can compare locally. (This is possible because P_B ’s bit x_2 is authenticated by P_A). However, the problem is that P_A is not able to learn the outcome of the comparison. To solve this, we let P_A send encrypted $V_0 \oplus R$ and $V_1 \oplus R$ for some random R such that P_B learns $V_{x_2} \oplus R$, and learn R from it. Now P_A and P_B can check the equality on R using the \mathcal{F}_{EQ} functionality in the TinyOT paper that allows both parties get the outcome. Note that this allows P_A to perform an additional selective failure attack on x_2 , by sending some corrupted encrypted values. This does not introduce

additional leakage, since x_2 is allowed to be learnt by \mathcal{A} anyway. Now \mathcal{A} is allowed to guess x_2 twice, once in step 4 and once in step 5. If the guesses are inconsistent, it is guaranteed to abort.

Combining leaky ANDs. The above check is vulnerable to selective failure attack, from which a malicious party can learn the value of x_1/x_2 with a risk of caught with one-half probability. In order to get rid of the leakage, bucketing is performed similar to TinyOT. Here, the key is to devise a way to combine leaky objects. Assuming that two triple are $([x'_1]_{\mathcal{A}}, [y'_1]_{\mathcal{A}}, [z'_1]_{\mathcal{A}}, [x'_2]_{\mathcal{B}}, [y'_2]_{\mathcal{B}}, [z'_2]_{\mathcal{B}})$ and $([x''_1]_{\mathcal{A}}, [y''_1]_{\mathcal{A}}, [z''_1]_{\mathcal{A}}, [x''_2]_{\mathcal{B}}, [y''_2]_{\mathcal{B}}, [z''_2]_{\mathcal{B}})$. Note that for each triple, only x_1, x_2 can be leaked. Therefore, one natural way is to set $[x_1]_{\mathcal{A}} := [x'_1]_{\mathcal{A}} \oplus [x''_1]_{\mathcal{A}}, [x_2]_{\mathcal{B}} := [x'_2]_{\mathcal{B}} \oplus [x''_2]_{\mathcal{B}}$. By doing this, $[x_1]_{\mathcal{A}}, [x_2]_{\mathcal{B}}$ are non-leaky as long as one triple is non-leaky. We can also set $[y_1]_{\mathcal{A}} := [y'_1]_{\mathcal{A}}, [y_2]_{\mathcal{B}} := [y'_2]_{\mathcal{B}}$ and reveal the bit $d := y'_1 \oplus y'_2 \oplus y''_1 \oplus y''_2$, since y 's bits are all private. Now observe that

$$\begin{aligned}
(x_1 \oplus x_2)(y_1 \oplus y_2) &= (x'_1 \oplus x'_2 \oplus x''_1 \oplus x''_2)(y'_1 \oplus y'_2) \\
&= (x'_1 \oplus x'_2)(y'_1 \oplus y'_2) \oplus (x''_1 \oplus x''_2)(y'_1 \oplus y'_2) \\
&= (x'_1 \oplus x'_2)(y'_1 \oplus y'_2) \oplus (x''_1 \oplus x''_2)(y''_1 \oplus y''_2) \oplus (x''_1 \oplus x''_2)(y'_1 \oplus y'_2 \oplus y''_1 \oplus y''_2) \\
&= (z'_1 \oplus z'_2) \oplus (z''_1 \oplus z''_2) \oplus d(x''_1 \oplus x''_2) \\
&= (z'_1 \oplus z''_1 \oplus dx''_1) \oplus (z'_2 \oplus z''_2 \oplus dx''_2)
\end{aligned}$$

Therefore, we could just set $[z_1]_{\mathcal{A}} := [z'_1]_{\mathcal{A}} \oplus [z''_1]_{\mathcal{A}} \oplus d[x''_1]_{\mathcal{A}}, [z_2]_{\mathcal{A}} := [z'_2]_{\mathcal{A}} \oplus [z''_2]_{\mathcal{A}} \oplus d[x''_2]_{\mathcal{A}}$. The security of this bucketing and merging can be proved as in [NNOB12, Appendix I].

6.1 Proof Sketch

In the following, we will discuss from a high-level view how the proof works for the new TinyOT protocol. We will focus on the security of Π_{LaAND} protocol, since the security of Π_{aAND} is fairly straightforward given the proof in the original paper [NNOB12].

Lemma 6.1. *The protocol in Figure 5 securely implements the functionality in Figure 3 against corrupted $\mathcal{P}_{\mathcal{A}}$ in the $(\mathcal{F}_{\text{abit}}, \mathcal{F}_{\text{EQ}})$ -Hybrid model.*

Proof. We will construct a simulator as follows:

1 \mathcal{S} interacts with \mathcal{A} and receive $(x_1, M[x_1]), (y_1, M[y_1]), (z_1, M[z_1]), K[x_2], K[y_2], K[r], \Delta_{\mathcal{A}}$ that \mathcal{A} sent to $\mathcal{F}_{\text{abit}}$. \mathcal{S} picks a random bit s , sets $K[z_2] := K[r] \oplus s\Delta_{\mathcal{A}}$, and sends $(x_1, M[x_1]), (y_1, M[y_1]), (z_1, M[z_1]), K[x_2], K[y_2], K[z_2], \Delta_{\mathcal{A}}$ to $\mathcal{F}_{\text{LaAND}}$, which sends $(x_2, M[x_2]), (y_2, M[y_2]), (z_2, M[z_2]), K[x_1], K[y_1], K[z_1], \Delta_{\mathcal{B}}$ to $\mathcal{P}_{\mathcal{B}}$.

2-3 \mathcal{S} randomly picks one row of $G_{i,j}$ and check its correctness. If it is not computed correctly, \mathcal{S} aborts; otherwise, \mathcal{S} announce $s \oplus z_2$.

4 \mathcal{S} sends a random U^* to \mathcal{A} , and receives some W_0, W_1 and computes some R_0, R_1 , such that, if $x_1 = 0$, $W_0 := H(K[x_2]) \oplus V_0 \oplus R_1, W_1 := H(K[x_2] \oplus \Delta_{\mathcal{A}}) \oplus V_1 \oplus R_2$; otherwise, $W_0 := H(K[x_2]) \oplus V_1 \oplus U^* \oplus R_1$ and $W_1 := H(K[x_2] \oplus \Delta_{\mathcal{A}}) \oplus V_0 \oplus U^* \oplus R_2$.

\mathcal{S} also obtains R that \mathcal{A} sent to \mathcal{F}_{EQ} . If R does not equal to either R_0 or R_1 , \mathcal{S} aborts; otherwise \mathcal{S} computes g_1 such that $R \neq R_{g_1}$ for some $g_1 \in \{0, 1\}$.

5 \mathcal{S} receives U , picks random W_0^*, W_1^* and sends them to \mathcal{A} . \mathcal{S} obtains R' that \mathcal{A} sent to \mathcal{F}_{EQ} .

- If both U, R' are honestly computed, \mathcal{S} proceeds as normal.

- If U is not honestly computed and that $R' = W_{x_1}^* \oplus H(M[x_1]) \oplus T_{x_1}$ is honestly computed, \mathcal{S} set $g_2 = 0$
- If either of the following is true: 1) $x_1 = 0$ and $R' = W_{x_1}^* \oplus H(M[x_1]) \oplus U \oplus H(K[x_1] \oplus \Delta_B, K[y_1] \oplus (y_2 \oplus z_2) \Delta_B)$; 2) $x_1 = 1$ and $R' = W_{x_1}^* \oplus H(M[x_1]) \oplus U \oplus H(K[x_1] \oplus \Delta_B, K[z_1] \oplus z_2 \Delta_B)$, \mathcal{S} sets $g_2 = 1$.
- Otherwise \mathcal{S} aborts.

6 If $g_1 \neq g_2$, \mathcal{S} aborts; otherwise, \mathcal{S} sends g_1 to $\mathcal{F}_{\text{LaAND}}$. If $\mathcal{F}_{\text{LaAND}}$ abort, \mathcal{S} aborts.

Note that the first 3 steps are perfect simulation. In step 4, U^* is sent in the simulation, while U_{x_2} is sent. This is a perfect simulation unless both of the input to Random Oracle in U_{x_2} get queried. This does not happen during the protocol, since Δ_B is not known to \mathcal{A} . In step 5, W_0^*, W_1^* are sent in the simulation, while $W_{x_2,0}, W_{x_2,0}$ are sent in the real protocol. This is also a perfect simulation unless \mathcal{P}_A gets Δ_B : both R and one of $H(K[x_1])$ and $H(K[x_1] \oplus \Delta_B)$ are random.

Another difference is that \mathcal{P}_B always aborts in the simulation if G_{x_2, y_2} is not honestly computed. This is also the case in the real protocol unless \mathcal{A} learns Δ_B . □

Lemma 6.2. *The protocol in Figure 5 securely implements the functionality in Figure 3 against corrupted \mathcal{P}_B in the $(\mathcal{F}_{\text{abit}}, \mathcal{F}_{\text{EQ}})$ -Hybrid model.*

Proof. We will construct a simulator as follows:

1. \mathcal{S} interacts with \mathcal{A} and receive $(x_2, M[x_2]), (y_2, M[y_2]), (r, M[r]), K[x_1], K[y_1], K[z_1], \Delta_B$ that \mathcal{A} sent to $\mathcal{F}_{\text{abit}}$. \mathcal{S} picks a random bit s , sets $(z_2, M[z_2]) := (r \oplus s, M[z_2] \oplus s \Delta_B)$, and sends $(x_2, M[x_2]), (y_2, M[y_2]), (z_2, M[z_2]), K[x_1], K[y_1], K[z_1]$ to $\mathcal{F}_{\text{LaAND}}$, which sends $(x_1, M[x_1]), (y_1, M[y_1]), (z_1, M[z_1]), K[x_2], K[y_2], K[z_2]$ to \mathcal{P}_B .

2-3 \mathcal{S} sends \mathcal{A} four random bits.

4-5 The simulation are the symmetric to the simulation for malicious \mathcal{P}_A .

The first three steps are perfect simulation; the proof for step 4 and 5 are the same as the proof for malicious \mathcal{P}_A (with order of steps switched). □

7 Extensions and Optimizations

Reducing the size of the garbled table. In the original protocol, all MAC keys are κ -bit values, which may not be necessary. For ρ -bit statistical security, $M[r_{00}]$ encrypted in step 4(d) only needs to be of size ρ bit. This reduces the size of a garbled table from 8κ bits to $4(\kappa + \rho)$ bits.

Partial garbled row reduction (PGRR). After applying the above optimization, the size of a garbled table is $4(\kappa + \rho + 1)$. However, we observe that randomness in $L_{\gamma,0}$ are not utilized, which means we can potentially perform Garbled Row Reduction but only on part of the first row. In particular, instead of picking $L_{\gamma,0}$ randomly, it will be set as $L_{\gamma,0} = H(L_{\alpha,0}, L_{\beta,0}, \gamma, 0)[0 : \kappa]$, where $[0 : \kappa]$ means obtaining the lower κ bits.

Note that this optimization does not increase the round trip of the protocol, because there is no interaction needed at the time to send the garbled tables.

Bucket size	3	4	5
$\rho = 40$	280K	3.1K	320
$\rho = 64$	1.2G	780K	21K
$\rho = 80$	300G	32M	330K

Table 2: Least number of AND gates needed in the bucketing, for different bucket size and statistical security parameter.

Pushing computation to earlier phases. In our protocol, it requires online communication complexity $|\mathcal{I}|(\kappa + \rho) + |\mathcal{O}|\rho$. It is easy to eliminate the term $|\mathcal{O}|\rho$, by sending encryption of $\{(r_w, \mathbf{M}[r_w])\}_{w \in \mathcal{O}}$ and commitment of the key used in the encryption. There is no increase in communication, round trip or total computation. Further, IT-MACs in step 5 and step 6 can also be sent in the function independent phase. Further, IT-MAC associated with input masks can also be sent in the function dependent phase. The resulting online phase has communication $|\mathcal{I}|\kappa$.

Extending to a two-output protocol. Our protocol can be extended to a two-output version such that generator also gets an output. Denoting \mathcal{O}_1 as the set of output wire-indices for \mathbf{P}_A , then after step 7, \mathbf{P}_B learns $\{\mathbf{L}_{w,z \oplus \lambda_w}\}_{w \in \mathcal{O}_1}$. Instead of following step 8, \mathbf{P}_B sends $\{s_w, \mathbf{M}[s_w], \mathbf{L}_{w,z \oplus \lambda_w}\}_{w \in \mathcal{O}_1}$ to \mathbf{P}_A , who check the validity of s_w , computes λ_w , and computes z using $\mathbf{L}_{w,z \oplus \lambda_w}$ and λ_w . \mathbf{P}_B is not able to obtain \mathbf{P}_A 's input, since the values are masked by some value unknown to \mathbf{P}_B ; \mathbf{P}_B also cannot flip the output, which requires either knowing Δ_1 or forging an IT-MAC.

More TinyOT optimizations. In the following we will briefly discuss more optimizations when \mathcal{F}_{Pre} is instantiated using our TinyOT protocol.

1. For clarity, R was chosen randomly in Π_{LaAND} . It is possible to perform garbled row reduction, such that $W_{0,0}, W_{1,0}$ are zero. This saves two ciphertext per leaky AND.
2. For the value of R 's and U 's, only ρ bits of the values are needed to be sent.

8 Evaluation

We are planning to implement our protocol. In the following, we will discuss the communication complexity of our scheme compared to other schemes as well as the cost at each stage. Throughout this section, all numbers will be given with $\kappa = 128, \rho = 40$.

We count the communication of TinyOT based on optimization mentioned in previous sections. After applying the optimization from Nielsen et al. [NST17] for the authenticated bit, the communication for each authenticated bit is 21 bytes and the communication for authenticated AND is about $94B$ bytes from each party, where B is the bucket size.

In Table 2 we calculated the smallest number of gates needed in the TinyOT protocol in order to make the bucketing works. The calculation is based on the formula in Appendix B of their paper, which is tighter. In Table 3 we compare the communication complexity of our protocol with other related works. Similar to previous papers, only one way communication is counted. As we can see, our total communication is $3\times$ to $6\times$ less than Nielsen et al.'s protocol. Further, our cost with single execution is also twice less than the cost of Nielsen et al.'s with 1024 circuits. Note that

Protocol	#execution	Ind. Process	Dep. Process	Online
[RR16]	32	-	3.75 MB	25.76 kB
	128	-	2.5 MB	21.31 kB
	1024	-	1.56 MB	16.95 kB
[NST17]	1	14.94 MB	226.86 kB	16.13 kB
	32	8.74 MB	226.86 kB	16.13 kB
	128	7.22 MB	226.86 kB	16.13 kB
	1024	6.42 MB	226.86 kB	16.13 kB
This paper	1	2.56 MB	464.3 kB	4.1 kB
	32	2.56 MB	464.3 kB	4.1 kB
	128	1.92 MB	464.3 kB	4.1 kB
	1024	1.92 MB	464.3 kB	4.1 kB

Table 3: Comparison of communication with previous protocols for an AES circuit. Data are counted as the amount sent to evaluator. Total communication will be roughly doubled for [RR16] and this paper.

for protocols based on cut-and-choose, the total communication to send 40 AES garbled circuit is 8.7 MB, which is already higher than the total communication of ours in the single execution setting.

We also observe that our function dependent processing is higher than Nielsen et al. this is due to that we need to send $3\kappa + 4\rho$ bits per gate while they only need to send 2κ bits. On the other hand, our online communication is extremely small: it is about $4\times$ smaller than in the protocol of Nielsen et al. and $4\text{--}6\times$ smaller than in the protocol of Rindal and Rosulek.

Acknowledgments

This material is based upon work supported by NSF awards #1111599 and #1563722; Samuel Ranellucci is also supported by NSF award #1564088.

References

- [AMPR14] Arash Afshar, Payman Mohassel, Benny Pinkas, and Ben Riva. Non-interactive secure computation based on cut-and-choose. In *Advances in Cryptology—Eurocrypt 2014*, volume 8441 of *LNCS*, pages 387–404. Springer, 2014.
- [BMR90] D. Beaver, S. Micali, and P. Rogaway. The round complexity of secure protocols. In *22nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 503–513. ACM Press, 1990.
- [Bra13] Luís T. A. N. Brandão. Secure two-party computation with reusable bit-commitments, via a cut-and-choose with forge-and-lose technique. In *Advances in Cryptology—Asiacrypt 2013, Part II*, volume 8270 of *LNCS*, pages 441–463. Springer, 2013.

- [CKMZ14] Seung Geol Choi, Jonathan Katz, Alex J. Malozemoff, and Vassilis Zikas. Efficient three-party computation from cut-and-choose. In *Advances in Cryptology—Crypto 2014, Part II*, volume 8617 of *LNCS*, pages 513–530. Springer, 2014.
- [DI05] Ivan Damgård and Yuval Ishai. Constant-round multiparty computation using a black-box pseudorandom generator. In *Advances in Cryptology—Crypto 2005*, volume 3621 of *LNCS*, pages 378–394. Springer, 2005.
- [DPSZ12] Ivan Damgård, Valerio Pastro, Nigel P. Smart, and Sarah Zakarias. Multiparty computation from somewhat homomorphic encryption. In *Advances in Cryptology—Crypto 2012*, volume 7417 of *LNCS*, pages 643–662. Springer, 2012.
- [FJN⁺13] Tore Kasper Frederiksen, Thomas Pelle Jakobsen, Jesper Buus Nielsen, Peter Sebastian Nordholt, and Claudio Orlandi. MiniLEGO: Efficient secure two-party computation from general assumptions. In *Advances in Cryptology—Eurocrypt 2013*, volume 7881 of *LNCS*, pages 537–556. Springer, 2013.
- [FJN14] Tore Kasper Frederiksen, Thomas P. Jakobsen, and Jesper Buus Nielsen. Faster maliciously secure two-party computation using the GPU. In *9th Intl. Conf. on Security and Cryptography for Networks (SCN)*, volume 8642 of *LNCS*, pages 358–379. Springer, 2014.
- [FJNT15] Tore Kasper Frederiksen, Thomas P. Jakobsen, Jesper Buus Nielsen, and Roberto Trifiletti. TinyLEGO: An interactive garbling scheme for maliciously secure two-party computation. Cryptology ePrint Archive, Report 2015/309, 2015. <http://eprint.iacr.org/2015/309>.
- [GMW87] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game, or a completeness theorem for protocols with honest majority. In *19th Annual ACM Symposium on Theory of Computing (STOC)*, pages 218–229. ACM Press, 1987.
- [HKE13] Yan Huang, Jonathan Katz, and David Evans. Efficient secure two-party computation using symmetric cut-and-choose. In *Advances in Cryptology—Crypto 2013, Part II*, volume 8043 of *LNCS*, pages 18–35. Springer, 2013.
- [HKK⁺14] Yan Huang, Jonathan Katz, Vladimir Kolesnikov, Ranjit Kumaresan, and Alex J. Malozemoff. Amortizing garbled circuits. In *Advances in Cryptology—Crypto 2014, Part II*, volume 8617 of *LNCS*, pages 458–475. Springer, 2014.
- [HZ15] Yan Huang and Ruiyu Zhu. Revisiting LEGOs: Optimizations, analysis, and their limit. Cryptology ePrint Archive, Report 2015/1038, 2015. <http://eprint.iacr.org/2015/1038>.
- [IPS08] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. Founding cryptography on oblivious transfer—efficiently. In *Advances in Cryptology—Crypto 2008*, volume 5157 of *LNCS*, pages 572–591. Springer, 2008.
- [KOS16] Marcel Keller, Emmanuela Orsini, and Peter Scholl. MASCOT: Faster malicious arithmetic secure computation with oblivious transfer. In *23rd ACM Conf. on Computer and Communications Security (CCS)*, pages 830–842. ACM Press, 2016.

- [KS08] Vladimir Kolesnikov and Thomas Schneider. Improved garbled circuit: Free XOR gates and applications. In *35th Intl. Colloquium on Automata, Languages, and Programming (ICALP), Part II*, volume 5126 of *LNCS*, pages 486–498. Springer, 2008.
- [Lin13] Yehuda Lindell. Fast cut-and-choose based protocols for malicious and covert adversaries. In *Advances in Cryptology—Crypto 2013, Part II*, volume 8043 of *LNCS*, pages 1–17. Springer, 2013.
- [LOP11] Yehuda Lindell, Eli Oxman, and Benny Pinkas. The IPS compiler: Optimizations, variants and concrete efficiency. In *Advances in Cryptology—Crypto 2011*, volume 6841 of *LNCS*, pages 259–276. Springer, 2011.
- [LP07] Yehuda Lindell and Benny Pinkas. An efficient protocol for secure two-party computation in the presence of malicious adversaries. In *Advances in Cryptology—Eurocrypt 2007*, volume 4515 of *LNCS*, pages 52–78. Springer, 2007.
- [LP11] Yehuda Lindell and Benny Pinkas. Secure two-party computation via cut-and-choose oblivious transfer. In *8th Theory of Cryptography Conference—TCC 2011*, volume 6597 of *LNCS*, pages 329–346. Springer, 2011. Available at <http://eprint.iacr.org/2010/284>.
- [LPSY15] Yehuda Lindell, Benny Pinkas, Nigel P. Smart, and Avishay Yanai. Efficient constant round multi-party computation combining BMR and SPDZ. In *Advances in Cryptology—Crypto 2015, Part II*, volume 9216 of *LNCS*, pages 319–338. Springer, 2015.
- [LR14] Yehuda Lindell and Ben Riva. Cut-and-choose Yao-based secure computation in the online/offline and batch settings. In *Advances in Cryptology—Crypto 2014, Part II*, volume 8617 of *LNCS*, pages 476–494. Springer, 2014.
- [LR15] Yehuda Lindell and Ben Riva. Blazing fast 2PC in the offline/online setting with security for malicious adversaries. In *22nd ACM Conf. on Computer and Communications Security (CCS)*, pages 579–590. ACM Press, 2015.
- [MOR16] Payman Mohassel, Ostap Orobets, and Ben Riva. Efficient server-aided 2PC for mobile phones. *Proc. Privacy Enhancing Technologies*, 2016(2):82–99, 2016.
- [NNOB12] Jesper Buus Nielsen, Peter Sebastian Nordholt, Claudio Orlandi, and Sai Sheshank Burra. A new approach to practical active-secure two-party computation. In *Advances in Cryptology—Crypto 2012*, volume 7417 of *LNCS*, pages 681–700. Springer, 2012.
- [NO09] Jesper Buus Nielsen and Claudio Orlandi. LEGO for two-party secure computation. In *6th Theory of Cryptography Conference—TCC 2009*, volume 5444 of *LNCS*, pages 368–386. Springer, 2009.
- [NST17] Jesper Nielsen, Thomas Schneider, and Roberto Trifiletti. Constant-round maliciously secure 2PC with function-independent preprocessing using LEGO. In *Network and Distributed System Security Symposium (NDSS)*, 2017.

- [RR16] Peter Rindal and Mike Rosulek. Faster malicious 2-party secure computation with online/offline dual execution. In *Proc. 25th USENIX Security Symposium*, pages 297–314. USENIX Association, 2016.
- [SS11] Abhi Shelat and Chih-Hao Shen. Two-output secure computation with malicious adversaries. In *Advances in Cryptology—Eurocrypt 2011*, volume 6632 of *LNCS*, pages 386–405. Springer, 2011.
- [WMK16] Xiao Wang, Alex J. Malozemoff, and Jonathan Katz. Faster two-party computation secure against malicious adversaries in the single-execution setting. Cryptology ePrint Archive, Report 2016/762, 2016. <http://eprint.iacr.org/2016/762>.
- [Yao86] Andrew C.-C. Yao. How to generate and exchange secrets. In *27th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 162–167. IEEE, 1986.