

Optimizing Authenticated Garbling for Faster Secure Two-Party Computation

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

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

Mike Rosulek
Oregon State University
rosulekm@eecs.oregonstate.edu

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

June 14, 2018

Abstract

Wang et al. (CCS 2017) recently proposed a protocol for malicious secure two-party computation that represents the state-of-the-art with regard to concrete efficiency in both the single-execution and amortized settings, with or without preprocessing. We show here several optimizations of their protocol that result in a significant improvement in the overall communication and running time. Specifically:

- We show how to make the “authenticated garbling” at the heart of their protocol compatible with the half-gate optimization of Zahur et al. (Eurocrypt 2015). We also show how to avoid sending an information-theoretic MAC for each garbled row. These two optimizations give up to a $2.6\times$ improvement in communication, and make the communication of the online phase essentially equivalent to that of state-of-the-art *semi-honest* secure computation.
- We show various optimizations to their protocol for generating AND triples that, overall, result in a $1.5\times$ improvement in the communication and a $2\times$ improvement in the computation for that step.

1 Introduction

In recent years, we have witnessed amazing progress in secure two-party computation, in both the semi-honest and malicious settings. In the semi-honest case, there has been an orders-of-magnitude improvement in protocols based on Yao’s garbled circuit [Yao86] since the initial implementation by Malkhi et al. [MNPS04]. This has resulted from several important techniques, including oblivious-transfer extension [IKNP03], pipelining [HEKM11], hardware acceleration [BHKR13], free-XOR [KS08] and other improved garbling techniques [PSSW09, KMR14], etc. Similarly, the concrete efficiency of secure two-party computation in the *malicious* case has also improved tremendously in both the single-execution [LP07, NO09, SS11, LP11, KSS12, FJN⁺13, HKE13, Lin13, SS13, Bra13, AMPR14, WMK17, NST17, KNR⁺17, ZH17, WRK17a] and amortized [HKK⁺14, LR14, LR15, NO16, RR16] settings. Whereas initial implementations in the malicious case could

	One-way comm.		Two-way comm.	
	Dep. + online	Total	Dep. + online	Total
semi-honest	0.22	0.22	0.22	0.22
Single-execution setting				
[NST17]	0.22	15	0.22	15
[WRK17a]	0.57	3.43	0.57	6.29
[HIV17]	3.39	3.39	3.39	3.39
This work, v. 1	0.33	2.24	0.33	4.15
This work, v. 2	0.22	2.67	0.22	5.12
Amortized setting (1024 executions)				
[RR16]	1.60	1.60	3.20	3.20
[NST17]	0.22	6.6	0.22	6.6
[WRK17a]	0.57	2.57	0.57	4.57
[KNR ⁺ 17]	1.59	1.59	1.59	1.59
This work, v. 1	0.33	1.70	0.33	3.07
This work, v. 2	0.22	2.13	0.22	4.04

Table 1: **Communication complexity of different protocols (in MB) for evaluating an AES circuit.** One-way communication refers to the maximum communication one party sends to the other; two-way communication refers to the sum of both parties’ communication. The best prior number in each column is bolded for reference.

evaluate up to 1,000 gates at the rate of 1 gate/second [PSSW09], the current state-of-the-art protocol by Wang et al. [WRK17a] (the *WRK protocol*) can compute tens of millions of gates at a rate up to 700,000× faster. With this steady stream of improvements, it has become more and more difficult to squeeze out additional performance gains; as an illustrative example, Zahur et al. [ZRE15] introduced a highly non-trivial optimization (“half-gates”) just to reduce communication by 33%.

We show several improvements to the WRK protocol that, overall, improve its performance by 2–3×. Recall their protocol can be divided into three phases: a *function-independent phase* (Ind.) in which the parties know an upper bound on the number of gates in the circuit to be evaluated and the lengths of their inputs; a *function-dependent phase* (Dep.) in which the parties know the circuit, but not their inputs; and an *online phase* in which the parties evaluate the circuit on their respective inputs. Our results can be summarized as follows:

- We show how to make the “authenticated garbling” at the heart of the online phase of the WRK protocol compatible with the half-gate optimization of Zahur et al. We also show that it is possible to avoid sending an information-theoretic MAC for each garbled row. These two optimizations result in up to a 2.6× improvement in communication and, somewhat surprisingly, result in a protocol for malicious secure two-party computation in which the communication complexity of the online phase is essentially equivalent to that of state-of-the-art *semi-honest* secure computation.
- The function-dependent phase of the WRK protocol involves the computation of (shared)

“AND triples” between the parties. We show various optimizations of that step that result in a $1.5\times$ improvement in the communication and a $2\times$ improvement in the computation. Our optimizations also simplify the protocol significantly.

We can combine these improvements in various ways, and suggest in particular two instantiations of protocols with malicious security: one that minimizes the total communication across all phases, and one that trades off increased communication in the function-independent phase for reduced communication in the function-dependent phase. These protocols improve upon the state-of-the-art by a significant margin, as summarized in Table 1. For example, compared to the protocol of Nielsen et al. [NST17] we achieve the same communication across the function-dependent and online phases, but improve the total communication by more than $6\times$; compared to the prior work with the best total communication [HIV17], we achieve a $1.5\times$ improvement overall and, at the same time, push almost all communication to the function-independent preprocessing phase. (Our protocol also appears to be significantly better than that of Hazay et al. [HIV17] in terms of computation. See Section 6 for a more detailed discussion.)

The multi-party case. It is natural to wonder whether we can extend our improved technique for authenticated garbling to the multi-party case, i.e., to improve upon [WRK17b]. Unfortunately, we have not yet been able to do so. In Section 7, we discuss some of the difficulties that arise.

1.1 Outline

In Section 2 we provide some background about the WRK protocol. We provide the high-level intuition behind our improvements in Section 3. In Section 4, we describe in detail our optimizations of the online phase of the WRK protocol, and in Section 5 we discuss our optimizations of the preprocessing phase. In Section 6, we compare our resulting protocols to prior work.

2 Background

We begin by describing some general background, followed by an in-depth review of the authenticated-garbling technique introduced in [WRK17a]. In the section that follows, we give a high-level overview of our optimizations and improvements.

We use κ and ρ to denote the computational and statistical security parameters, respectively. We sometimes use “:=” to denote assignment.

Information-theoretic MACs. As in prior work, we authenticate bits using a particular information-theoretic MAC. Let $\Delta_B \in \{0, 1\}^\kappa$ be a value known to P_B that is chosen at the outset of the protocol. We say a bit b known to P_A is *authenticated to* P_B if P_B holds a key $K[b]$ and P_A holds the corresponding tag $M[b] = K[b] \oplus b\Delta_B$. We abstractly denote such a bit by $[b]_A$; i.e., for some fixed Δ_B , when we say the parties hold $[b]_A$ we mean that P_A holds $(b, M[b])$ and P_B holds $K[b]$ such that $M[b] = K[b] \oplus b\Delta_B$. We analogously let $[b]_B$ denote a bit b known to P_B and authenticated to P_A .

A pair of authenticated bits $[b_1]_A, [b_2]_B$, each known to a different party, form an *authenticated share* of $b_1 \oplus b_2$. We denote this by $\langle b_1 | b_2 \rangle$, where the value in the left slot is known to P_A , and the value in the right slot is known to P_B . Both authenticated bits and authenticated shares are XOR-homomorphic.

Authenticated bits can be computed efficiently based on oblivious transfer [NNOB12, NST17]. We abstract away the particular protocol used to generate authenticated bits, and design our

Functionality $\mathcal{F}_{\text{abit}}$

Honest case:

1. Upon receiving `init` from both parties, choose uniform $\Delta_A, \Delta_B \in \{0, 1\}^\kappa$; send Δ_A to P_A and Δ_B to P_B .
2. Upon receiving `(random, A)` from both parties, choose uniform $x \in \{0, 1\}$ and $K[x] \in \{0, 1\}^\kappa$, set $M[x] := K[x] \oplus x\Delta_B$, and send $(x, M[x])$ to P_A and $K[x]$ to P_B .
3. Upon receiving `(random, B)` from both parties, generate an authenticated bit for P_B in a manner symmetric to the above.

Corrupted parties: A corrupted party can specify the randomness used on its behalf by the functionality.

Global-key queries: A corrupted P_A (resp., P_B) can, at any time, send Δ , and is told whether $\Delta = \Delta_B$ (resp., $\Delta = \Delta_A$).

Figure 1: The authenticated-bits functionality.

protocols in the $\mathcal{F}_{\text{abit}}$ -hybrid model (cf. Figure 1) in which there is an ideal functionality that provides them.

Opening authenticated values. An authenticated bit $[b]_A$ known to P_A can be opened by having P_A send b and $M[b]$ to P_B , who then verifies that $M[b] = K[b] \oplus b\Delta_B$. As observed in prior work [DPSZ12], it is possible to open n authenticated bits with less than n times the communication. Specifically, P_A can open $[b_1]_A, \dots, [b_n]_A$ by sending b_1, \dots, b_n along with $h := H(M[b_1], \dots, M[b_n])$, where H is a hash function modeled as a random oracle. P_A then simply checks whether $h = H(K[b_1] \oplus b_1\Delta_B, \dots, K[b_n] \oplus b_n\Delta_B)$.

We let $\text{Open}([b_1]_A, \dots)$ denote the process of opening one or more authenticated bits in this way, and overload this notation so that $\text{Open}(\langle b_1 | b_2 \rangle)$ denotes the process of having each party open its portion of an authenticated share.

Circuit-dependent preprocessing. We consider boolean circuits with gates represented as a tuple $(\alpha, \beta, \gamma, T)$, where α and β are (the indices of) the input wires of the gate, γ is the output wire of the gate, and $T \in \{\oplus, \wedge\}$ is the type of the gate. We use \mathcal{W} to denote the output wires of all AND gates, $\mathcal{I}_1, \mathcal{I}_2$ to denote the input wires for each party, and \mathcal{O} to denote the output wires.

Wang et al. [WRK17a] introduced an ideal functionality called \mathcal{F}_{pre} (cf. Figure 2) that is used by the parties in a circuit-dependent, but input-*independent*, preprocessing phase. This functionality sets up information for the parties as follows:

1. For each wire w that is either an input wire of the circuit or an output wire of an AND gate, generate a random authenticated share $\langle r_w | s_w \rangle$. We refer to the value $\lambda_w \stackrel{\text{def}}{=} r_w \oplus s_w$ as the *mask* on wire w .
2. For the output wire γ of each XOR gate $(\alpha, \beta, \gamma, \oplus)$, generate a random authenticated share $\langle r_\gamma | s_\gamma \rangle$ whose value $r_\gamma \oplus s_\gamma$ is the XOR of the masks on the input wires α, β .

Functionality \mathcal{F}_{pre}

1. Choose uniform $\Delta_A, \Delta_B \in \{0, 1\}^\rho$. Send Δ_A to P_A and Δ_B to P_B .
2. For each wire $w \in \mathcal{W} \cup \mathcal{I}$, generate a random authenticated share $\langle r_w | s_w \rangle$.
3. For each gate $\mathcal{G} = (\alpha, \beta, \gamma, T)$, in topological order:
 - If $T = \oplus$, generate a random authenticated share $\langle r_\gamma | s_\gamma \rangle$ for which $r_\gamma \oplus s_\gamma = r_\alpha \oplus s_\alpha \oplus r_\beta \oplus s_\beta$.
 - If $T = \wedge$, generate a random authenticated share $\langle r_\gamma^* | s_\gamma^* \rangle$ for which $r_\gamma^* \oplus s_\gamma^* = (r_\alpha \oplus s_\alpha) \wedge (r_\beta \oplus s_\beta)$.

Figure 2: Preprocessing functionality for some fixed circuit.

3. For each AND gate $(\alpha, \beta, \gamma, \wedge)$, generate a random authenticated share $\langle r_\gamma^* | s_\gamma^* \rangle$ such that

$$r_\gamma^* \oplus s_\gamma^* = (r_\alpha \oplus s_\alpha) \wedge (r_\beta \oplus s_\beta).$$

We refer to a triple of authenticated shares $(\langle r_\alpha | s_\alpha \rangle, \langle r_\beta | s_\beta \rangle, \langle r_\gamma^* | s_\gamma^* \rangle)$ for which $r_\gamma^* \oplus s_\gamma^* = (r_\alpha \oplus s_\alpha) \wedge (r_\beta \oplus s_\beta)$ as an *authenticated AND triple*. These are just (authenticated) Beaver triples [Bea92] over the field \mathbb{F}_2 .

Authenticated garbling. We now describe the idea behind the authenticated garbling technique from the WRK protocol. We assume the reader is familiar with basic concepts of garbled circuits, e.g., point-and-permute [BMR90], free-XOR [KS08], etc.

Following the preprocessing phase described above, every wire w is associated with a secret mask λ_w , unknown to either party. If the actual value on that wire (when the circuit is evaluated on the parties' inputs) is z_w , then the *masked value* on that wire is defined to be $\hat{z}_w = z_w \oplus \lambda_w$. We focus on garbling a single AND gate $(\alpha, \beta, \gamma, \wedge)$. Assume P_A is the circuit garbler and P_B is the circuit evaluator. Say the garbled wire labels are $(L_{\alpha,0}, L_{\alpha,1})$ and $(L_{\beta,0}, L_{\beta,1})$ for wires α and β , respectively. Since we apply the free-XOR optimization, P_A also holds Δ such that $L_{w,0} \oplus L_{w,1} = \Delta$ for any wire w . The protocol inductively ensures that the evaluator P_B knows the wire labels $L_{\alpha, \hat{z}_\alpha}, L_{\beta, \hat{z}_\beta}$ and masked values $\hat{z}_\alpha, \hat{z}_\beta$ for both input wires. Note that the correct masked value for the output wire is then

$$\hat{z}_\gamma = (\lambda_\alpha \oplus \hat{z}_\alpha) \wedge (\lambda_\beta \oplus \hat{z}_\beta) \oplus \lambda_\gamma,$$

and we need to ensure that P_B learns this value.

To achieve this, P_A generates a garbled gate consisting of 4 rows (one for each $u, v \in \{0, 1\}$)

$$G_{u,v} = H(L_{\alpha,u}, L_{\beta,v}) \oplus (r_{u,v}, M[r_{u,v}], [L_{\gamma, \hat{z}_{u,v}}]),$$

with bit $\hat{z}_{u,v}$ defined as

$$\hat{z}_{u,v} = (\lambda_\alpha \oplus u) \wedge (\lambda_\beta \oplus v) \oplus \lambda_\gamma.$$

Here, $[L_{\gamma, \hat{z}_{u,v}}]$ is P_A 's share of the garbled label; $r_{u,v}$ is P_A 's share of the bit $\hat{z}_{u,v}$; and P_B holds the corresponding share $s_{u,v}$ such that $r_{u,v} \oplus s_{u,v} = \hat{z}_{u,v}$. The value $M[r_{u,v}]$ is the MAC authenticating

the underlying bit to P_B . Also note that the definition of $\hat{z}_{u,v}$ indicates that when $u = \hat{z}_\alpha$ and $v = \hat{z}_\beta$ then $\hat{z}_{u,v} = \hat{z}_\gamma$.

Suppose the evaluator P_B holds $(u, L_{\alpha,u})$ and $(v, L_{\beta,v})$, where $u = \hat{z}_\alpha$ and $v = \hat{z}_\beta$. Then P_B can evaluate this AND gate by decrypting $G_{u,v}$ to obtain $r_{u,v}$ and P_A 's share of $L_{\gamma,\hat{z}_{u,v}}$. After verifying the MAC on $r_{u,v}$, party P_B can combine these values with its own shares to reconstruct the masked output value $\hat{z}_{u,v}$ (that is, \hat{z}_γ) and its corresponding label $L_{\gamma,\hat{z}_{u,v}}$ (that is, $L_{\gamma,\hat{z}_\gamma}$).

Assuming that the authenticated bits and shares of the labels can be computed securely, the above protocol is secure against malicious adversaries. In particular, even if P_A cheats and causes P_B to abort during evaluation, any such abort depends only on the *masked* values on the wires. Since the masks are random and unknown to either party, this means that any abort is input-independent. The MACs checked by P_B ensure correctness, namely that evaluation has resulted in the correct (masked) output-wire value.

From authenticated shares to shared labels. Another important optimization in the WRK protocol is to compute shares of labels efficiently using authenticated shares. Assume the parties hold an authenticated share $\langle r | s \rangle$ of some mask $\lambda = s \oplus r$. It is then easy to compute a share of $\lambda\Delta_A$, since

$$\lambda\Delta_A = (r \oplus s)\Delta_A = \left(r\Delta_A \oplus K[s]\right) \oplus \left(M[s]\right).$$

Since P_A has r , Δ_A , and $K[s]$ while P_B has $M[s]$, the two parties can locally compute shares of $\lambda\Delta_A$ (namely, $[\lambda\Delta_A]$) given only $\langle r | s \rangle$.

We can use this fact to compute shares of labels for a secret masked bit efficiently. Assuming the global authentication key (i.e., Δ_A) is also used as the free-XOR shift, then it holds that $L_{\gamma,\hat{z}_{u,v}} = L_{\gamma,0} \oplus \hat{z}_{u,v}\Delta_A$. Therefore, the task of computing shares of labels reduces to the task of computing shares of $\hat{z}_{u,v}\Delta_A$, since $L_{\gamma,0}$ is known to P_A .

Notice that

$$\begin{aligned} \hat{z}_{u,v}\Delta_A &= ((\lambda_\alpha \oplus u) \wedge (\lambda_\beta \oplus v) \oplus \lambda_\gamma) \Delta_A \\ &= \lambda_\alpha \lambda_\beta \Delta_A \oplus u \lambda_\alpha \Delta_A \oplus v \lambda_\beta \Delta_A \oplus uv \Delta_A \oplus \lambda_\gamma \Delta_A. \end{aligned}$$

If the parties hold an authenticated AND triple $(\langle r_\alpha | s_\alpha \rangle, \langle r_\beta | s_\beta \rangle, \langle r_\gamma^* | s_\gamma^* \rangle)$ and a random authenticated share $\langle r_\gamma | s_\gamma \rangle$ such that $\lambda_\alpha = r_\alpha \oplus s_\alpha$, $\lambda_\beta = r_\beta \oplus s_\beta$, $\lambda_\alpha \wedge \lambda_\beta = r_\gamma^* \oplus s_\gamma^*$, and $\lambda_\gamma = r_\gamma \oplus s_\gamma$. The parties can then locally compute shares of $\lambda_\alpha \Delta_A$, $\lambda_\beta \Delta_A$, $\lambda_\gamma \Delta_A$, and $(\lambda_\alpha \wedge \lambda_\beta) \Delta_A$, and finally compute shares of $\hat{z}_{u,v} \Delta_A$ by linearly combining the above shares.

3 Overview of Our Optimizations

We separately discuss our optimizations for the authenticated garbling and the preprocessing phases. Details and proofs can be found in Sections 4 and 5.

3.1 Improving Authenticated Garbling

As a high level, the key ideas behind authenticated garbling are that (1) it is possible to share garbled circuits such that neither party knows how rows in the garbled tables are permuted (since no party knows the masks on the wires); moreover, (2) information-theoretic MACs can be used to ensure correctness of the garbled tables. In the original protocol by Wang et al., these two aspects

are tightly integrated: each garbled row includes an encryption of the corresponding MAC tag, so the evaluator only learns one such tag for each gate.

Here, we take a slightly different perspective on how authenticated garbling works. In particular, we (conceptually) divide the protocol into two parts:

- In the first part, the parties compute a shared garbled circuit, without any authentication, and let the evaluator reconstruct and evaluate that garbled circuit. We stress here that, even though there is no authentication, corrupting one or more garbled rows does not allow a selective-failure attack for the same reason as in the WRK protocol: any failure depends only on the *masked* wire values, but neither party knows those masks.

This part is achieved by the encrypted wire labels alone, which have the form $H(L_{\alpha,u}, L_{\beta,v}) \oplus [L_{\gamma, \hat{z}_{u,v}}]$. These require 4κ bits of communication per gate.

- In the second part, the evaluator holds masked wire values for every wire of the circuit. It then checks correctness of all these masked values. For example, it will ensure that for every AND gate, the underlying (real) values on the wires form an AND relationship. Such verification is needed for masked values that P_B obtains during the evaluation of the garbled circuit.

The WRK protocol achieves this by encrypting *authenticated shares* of the form $H(L_{\alpha,u}, L_{\beta,v}) \oplus (r_{u,v}, M[r_{u,v}])$ in each row of a garbled table. The evaluator decrypts one of the rows and checks the appropriate tag. These encrypted tags contribute 4ρ bits of communication per gate.

With this new way of viewing authenticated garbling, we can optimize each part independently. By doing so, we are able to reduce the communication of the first part to $2\kappa + 1$ bits per gate, and reduce the communication of the second part to 1 bit per gate. In the process, we also reduce the computation (in terms of hash evaluations) by about half. In the following, we discuss intuitively how these optimizations work.

Applying row-reduction techniques. In garbled circuits, *row reduction* refers to techniques that use fewer than four garbled rows per garbled gate [NPS99, PSSW09, ZRE15, GLNP15]. We review the simplest row-reduction technique here, describe the challenge of applying the technique to authenticated garbling, and then show how we overcome the challenge. This will serve as a warm-up to our final protocol that is compatible with the half-gate technique.

In classical garbling, a garbled AND gate can be written as (in our notation):

$$\begin{aligned} G_{0,0} &= H(L_{\alpha,0}, L_{\beta,0}) \oplus L_{\gamma, \hat{z}_{0,0}} = H(L_{\alpha,0}, L_{\beta,0}) \oplus L_{\gamma,0} \oplus \hat{z}_{0,0} \Delta_A \\ G_{0,1} &= H(L_{\alpha,0}, L_{\beta,1}) \oplus L_{\gamma, \hat{z}_{0,1}} = H(L_{\alpha,0}, L_{\beta,1}) \oplus L_{\gamma,0} \oplus \hat{z}_{0,1} \Delta_A \\ G_{1,0} &= H(L_{\alpha,1}, L_{\beta,0}) \oplus L_{\gamma, \hat{z}_{1,0}} = H(L_{\alpha,1}, L_{\beta,0}) \oplus L_{\gamma,0} \oplus \hat{z}_{1,0} \Delta_A \\ G_{1,1} &= H(L_{\alpha,1}, L_{\beta,1}) \oplus L_{\gamma, \hat{z}_{1,1}} = H(L_{\alpha,1}, L_{\beta,1}) \oplus L_{\gamma,0} \oplus \hat{z}_{1,1} \Delta_A. \end{aligned}$$

The idea behind GRR3 row reduction [NPS99] is to choose wire labels so $G_{0,0} = 0^\kappa$. That is, the garbler chooses

$$L_{\gamma,0} := H(L_{\alpha,0}, L_{\beta,0}) \oplus \hat{z}_{0,0} \Delta_A.$$

The garbler now needs to send only $(G_{0,1}, G_{1,0}, G_{1,1})$, reducing the communication from 4κ to 3κ bits. If the evaluator has input wires with masked values $(0,0)$, it can simply set $G_{0,0} = 0^\kappa$ and then proceed as before.

In authenticated garbling, the preprocessing results in shares of $\{\hat{z}_{u,v}\Delta_{\mathbf{A}}\}$. Hence, if $P_{\mathbf{A}}$ could compute $L_{\gamma,0}$ then the parties could locally compute shares of the $\{G_{u,v}\}$ (since $P_{\mathbf{A}}$ knows all the $L_{\alpha,u}, L_{\beta,v}$ values and their hashes). $P_{\mathbf{A}}$ could then send its shares to $P_{\mathbf{B}}$ to allow $P_{\mathbf{B}}$ to recover the entire garbled gate. Unfortunately, $P_{\mathbf{A}}$ cannot compute $L_{\gamma,0}$ because $P_{\mathbf{A}}$ does not know $\hat{z}_{0,0}$! Indeed, that value depends on the secret wire masks, unknown to either party.

Summarizing, row-reduction techniques in general compute one (or both) of the output-wire labels as a function of the input-wire labels **and** the secret masks, making them a challenge for authenticated garbling.

Our observation is that although $P_{\mathbf{A}}$ does not know $\hat{z}_{0,0}$, the garbling requires only $\hat{z}_{0,0}\Delta_{\mathbf{A}}$ for which the parties do have shares. Let $S_{\mathbf{A}}$ and $S_{\mathbf{B}}$ denote the parties' shares of this value, so that $S_{\mathbf{A}} \oplus S_{\mathbf{B}} = \hat{z}_{0,0}\Delta_{\mathbf{A}}$. Our main idea is for the parties to “shift” the entire garbling process by the value $S_{\mathbf{B}}$, as follows:

1. $P_{\mathbf{A}}$ computes $L_{\gamma,0} := H(L_{\alpha,0}, L_{\beta,0}) \oplus S_{\mathbf{A}}$. Note this value differs from the standard garbling value by a shift of $S_{\mathbf{B}}$. Intuitively, instead of choosing $L_{\gamma,0}$ so that $G_{0,0} = 0^{\kappa}$, we set implicitly set $G_{0,0} = S_{\mathbf{B}}$. Although $P_{\mathbf{A}}$ does not know $S_{\mathbf{B}}$, it only matters that the evaluator $P_{\mathbf{B}}$ knows it.
2. Based on this value of $L_{\gamma,0}$, the parties locally compute shares of the garbled gate $G_{0,1}, G_{1,0}, G_{1,1}$ defined above, and open them to $P_{\mathbf{B}}$.
3. When $P_{\mathbf{B}}$ evaluates the gate on input $L_{\alpha,u}, L_{\beta,v}$, if $(u, v) \neq (0, 0)$ then evaluation is the same as usual. If $(u, v) = (0, 0)$ then $P_{\mathbf{B}}$ sets $G_{0,0} = S_{\mathbf{B}}$. This is equivalent to $P_{\mathbf{B}}$ doing the usual evaluation but shifting the result by $S_{\mathbf{B}}$.

Using the half-gate technique. The state-of-the-art in semi-honest garbling is the half-gate construction of Zahur et al. [ZRE15]. It requires 2κ bits of communication per AND gate, while being compatible with free-XOR. We describe this idea, translated from the original work [ZRE15] to be written in terms of masks and masked wire values so as to match our notation.

The circuit garbler computes a garbled gate as:

$$\begin{aligned} G_0 &:= H(L_{\alpha,0}) \oplus H(L_{\alpha,1}) \oplus \lambda_{\beta}\Delta_{\mathbf{A}} \\ G_1 &:= H(L_{\beta,0}) \oplus H(L_{\beta,1}) \oplus L_{\alpha,0} \oplus \lambda_{\alpha}\Delta_{\mathbf{A}}, \end{aligned}$$

and computes the 0-label for that gate's output wire as:

$$L_{\gamma,0} := H(L_{\alpha,0}) \oplus H(L_{\beta,0}) \oplus (\lambda_{\alpha}\lambda_{\beta} \oplus \lambda_{\gamma})\Delta_{\mathbf{A}}.$$

If the evaluator $P_{\mathbf{B}}$ holds masked values u, v and corresponding labels $L_{\alpha,u}, L_{\beta,v}$, it computes:

$$\text{Eval}(u, v, L_{\alpha,u}, L_{\beta,v}) := H(L_{\alpha,u}) \oplus H(L_{\beta,v}) \oplus uG_0 \oplus v(G_1 \oplus L_{\alpha,u}).$$

This results in the value

$$\begin{aligned} \text{Eval}(u, v, L_{\alpha,u}, L_{\beta,v}) &= H(L_{\alpha,0}) \oplus H(L_{\beta,0}) \oplus (uv \oplus v\lambda_{\alpha} \oplus u\lambda_{\beta})\Delta_{\mathbf{A}} \\ &= H(L_{\alpha,0}) \oplus H(L_{\beta,0}) \oplus \left((u \oplus \lambda_{\alpha})(v \oplus \lambda_{\beta}) \oplus \lambda_{\alpha}\lambda_{\beta} \right)\Delta_{\mathbf{A}} \\ &= H(L_{\alpha,0}) \oplus H(L_{\beta,0}) \oplus (\hat{z}_{u,v} \oplus \lambda_{\alpha}\lambda_{\beta} \oplus \lambda_{\gamma})\Delta_{\mathbf{A}}, \end{aligned}$$

which is the correct output $L_{\gamma, \hat{z}_{u,v}} = L_{\gamma,0} \oplus \hat{z}_{u,v} \Delta_A$.

As before, this garbling technique is problematic for authenticated garbling, because the garbler P_A cannot compute $L_{\gamma,0}$ as specified. (P_A does not know the wire masks, so cannot compute the term $(\lambda_\alpha \lambda_\beta \oplus \lambda_\gamma) \Delta_A$.)

However, the parties hold¹ shares of this value; say, $S_A \oplus S_B = (\lambda_\alpha \lambda_\beta \oplus \lambda_\gamma) \Delta_A$. We can thus conceptually “shift” the entire garbling procedure by S_B to obtain the following interactive variant of half-gates:

1. P_A computes the output wire label as

$$L_{\gamma,0} := H(L_{\alpha,0}) \oplus H(L_{\beta,0}) \oplus S_A,$$

which is “shifted” by S_B from what the half-gates technique specifies.

2. The parties locally compute shares of G_0, G_1 as per the half-gates technique described above. These shares are opened to P_B , so P_B learns (G_0, G_1) .
3. To evaluate the gate on inputs $L_{\alpha,u}, L_{\beta,v}$, the evaluator P_B performs standard half-gates evaluation and then adds S_B as a correction value. This results in the correct output-wire label, since:

$$\begin{aligned} \text{Eval}(L_{\alpha,u}, L_{\beta,v}) \oplus S_B &= \text{Eval}(L_{\alpha,u}, L_{\beta,v}) \oplus (\lambda_\alpha \lambda_\beta \oplus \lambda_\gamma) \Delta_A \oplus S_A \\ &= H(L_{\alpha,0}) \oplus H(L_{\beta,0}) \oplus \hat{z}_{u,v} \Delta_A \oplus S_A \\ &= L_{\gamma,0} \oplus \hat{z}_{u,v} \Delta_A \\ &= L_{\gamma, \hat{z}_{u,v}}. \end{aligned}$$

Authentication almost for free. In the WRK scheme, suppose the actual values on the wires of an AND gate are $z_\alpha, z_\beta, z_\gamma$ with $z_\alpha \wedge z_\beta = z_\gamma$. During evaluation, P_B learn masked values $\hat{z}_\alpha = z_\alpha \oplus \lambda_\alpha$, $\hat{z}_\beta = z_\beta \oplus \lambda_\beta$, and $\hat{z}_\gamma = z_\gamma \oplus \lambda_\gamma$. For correctness it suffices to show that

$$\begin{aligned} z_\alpha \wedge z_\beta = z_\gamma &\iff (\hat{z}_\alpha \oplus \lambda_\alpha) \wedge (\hat{z}_\beta \oplus \lambda_\beta) = (\hat{z}_\gamma \oplus \lambda_\gamma) \\ &\iff \underbrace{(\hat{z}_\alpha \oplus \lambda_\alpha) \wedge (\hat{z}_\beta \oplus \lambda_\beta)}_{\hat{z}_{\alpha,\beta}} \oplus \lambda_\gamma = \hat{z}_\gamma. \end{aligned}$$

Note the parties already have authenticated shares of $\lambda_\alpha, \lambda_\beta, \lambda_\gamma$, and $(\lambda_\alpha \wedge \lambda_\beta)$, so they can also derive authenticated shares of related values.

In the WRK scheme the garbler P_A prepares an authenticated share (MAC) of $\hat{z}_{\alpha,\beta}$ corresponding to each of the 4 possible values of $\hat{z}_\alpha, \hat{z}_\beta$. It encrypts this share so that it can only be opened using the corresponding wire labels. P_B can then decrypt and verify the relevant $\hat{z}_{\alpha,\beta}$ value (and take it to be the masked output value \hat{z}_γ).

Our approach is to apply a technique suggested for the SPDZ protocol [DPSZ12]: evaluate the circuit without authentication and then perform batch authentication at the end. Thus, in our new protocol authentication works as follows:

1. P_B evaluates the circuit, obtaining (unauthenticated) masked values \hat{z}_α for every wire α .

¹Note that $(\lambda_\alpha \lambda_\beta \oplus \lambda_\gamma) = \hat{z}_{0,0}$, the same secret value as in the previous example.

2. P_B reveals the masked values of every wire (1 bit per wire). Revealing these to P_A does not affect privacy because the masks are hidden from both parties (except for certain input/output wires where one or both of the parties already know the underlying values).
3. P_A generates authenticated shares of only the relevant $\hat{z}_{\alpha,\beta}$ values and sends them. P_B verifies the authenticity of each share. This is equivalent to sending a MAC of P_A 's shares. As described in Section 2, this can be done by sending only a hash of the MACs.

This technique for authentication adds an extra round, but it makes the authentication almost free in terms of communication. P_B sends 1 bit per wire and P_A sends only a single hash value to authenticate.

Details of the optimizations described above can be found in Section 4.

3.2 Improving the Preprocessing Phase

We also improve the efficiency of preprocessing in the WRK protocol significantly; specifically: (1) we design a new protocol for generating so-called leaky-AND triples. Compared to the best previous protocol by Wang et al., it reduces the number of hash calls by $2.5\times$ and reduces communication by κ bits. (2) we propose a new function-dependent preprocessing protocol that can be computed much more efficiently. We remark that the second optimization is particularly suitable for RAM-model secure computation, where CPU circuits are fixed ahead of time.

To enable the above optimizations, we need $\text{lsb}(\Delta_A) = 1$ and $\text{lsb}(\Delta_B) = 0$, where $\text{lsb}(x)$ denotes the least significant bit of x .

A new leaky-AND protocol. The output of a leaky-AND protocol is a random authenticated AND triple $(\langle r_\alpha | s_\alpha \rangle, \langle r_\beta | s_\beta \rangle, \langle r_\gamma^* | s_\gamma^* \rangle)$ with one caveat: the adversary can choose to guess the value of $r_\alpha \oplus s_\alpha$. A correct guess remains undetected while an incorrect guess will be caught. (See Figure 4 for a formal definition.) The leaky-AND protocol by Wang et al. works in two steps. Two parties first run a protocol whose outputs are triples that are leaky without any correctness guarantee; then a checking procedure is run to ensure correctness. The leakage is later eliminated by bucketing. In our new protocol, we observe that these two steps can be computed at the same time, reducing the number of rounds as well as the amount of computation (i.e., H -evaluations). Moreover, computing and checking can be further improved by adopting ideas from the half-gate technique. Details are below.

Recall that in the half-gate approach, if a wire is associated with wire labels $(L_0, L_1 = L_0 \oplus \Delta_A)$, the first row of the gate computed by the garbler has the form

$$G = H(L_0) \oplus H(L_1) \oplus C,$$

for some C . An evaluator holding (b, L_b) can evaluate it as

$$\begin{aligned} E &= bG \oplus H(L_b) \\ &= b(H(L_0) \oplus H(L_1) \oplus C) \oplus H(L_b) \\ &= b(H(L_0) \oplus H(L_1)) \oplus H(L_b) \oplus bC \\ &= H(L_0) \oplus bC. \end{aligned} \tag{1}$$

Correctness ensures that $E \oplus H(L_0) = bC$, which means that after the evaluation the two parties hold shares of bC . Note that when free-XOR is used with shift Δ_A , then a pair of garbled labels

(L_0, L_1) and the IT-MAC for a bit (i.e., $(K[b], M[b])$) have the same structure. Therefore the above can be reformulated and extended as follows:

$$G = H(K[b]) \oplus H(M[b]) \oplus C_1$$

$$E = bG \oplus H(M[b]) \oplus bC_2$$

. Assuming the two parties have an authenticated bit $[b]_{\mathbf{B}}$, then $E \oplus H(K[b]) = b(C_1 \oplus C_2)$. If we view C_1 and C_2 as shares of some value $C = C_1 \oplus C_2$, then this can be interpreted as a way to select on a shared value such that the selection bit b is known only to one party and at the same time the output (namely, $bC = H(K[b]) \oplus E$) is still shared.

Now we are ready to present our protocol. We will start with a set of random authenticated bits $(\langle x_1 | x_2 \rangle, \langle y_1 | y_2 \rangle, \langle z_1 | r \rangle)$. We want the two parties to directly compute shares of

$$S = ((x_1 \oplus x_2) \wedge (y_1 \oplus y_2) \oplus z_1 \oplus r) (\Delta_{\mathbf{A}} \oplus \Delta_{\mathbf{B}}).$$

Assuming $\text{lsb}(\Delta_{\mathbf{A}} \oplus \Delta_{\mathbf{B}}) = 1$, revealing $d = \text{lsb}(S)$ allows the parties to “fix” these random authenticated shares to a valid triple (by computing $[z_2]_{\mathbf{B}} = [r]_{\mathbf{B}} \oplus d$). Once the parties hold shares of S (for example, $P_{\mathbf{A}}$ holds S_1 and $P_{\mathbf{B}}$ holds $S_2 = S \oplus S_1$), checking the correctness of d also becomes easy: d is valid if and only if $S_1 \oplus d\Delta_{\mathbf{A}}$ from $P_{\mathbf{A}}$ equals to $S_2 \oplus d\Delta_{\mathbf{B}}$ from $P_{\mathbf{B}}$. A wrong d can pass the equality check only if the adversary guesses the other party’s Δ value. Now the task is to compute shares of S , where S can be rewritten as

$$S = x_1(y_1 \oplus y_2)(\Delta_{\mathbf{A}} \oplus \Delta_{\mathbf{B}}) \oplus x_2(y_1 \oplus y_2)(\Delta_{\mathbf{A}} \oplus \Delta_{\mathbf{B}}) \oplus (z_1 \oplus r)(\Delta_{\mathbf{A}} \oplus \Delta_{\mathbf{B}}).$$

Here, we will focus on how to compute shares of

$$x_2(y_1\Delta_{\mathbf{A}} \oplus y_1\Delta_{\mathbf{B}} \oplus y_2\Delta_{\mathbf{A}} \oplus y_2\Delta_{\mathbf{B}}).$$

Now we apply the half-gate observation: $P_{\mathbf{A}}$ has $C_1 = y_1\Delta_{\mathbf{A}} \oplus K[y_2] \oplus M[y_1]$ and $P_{\mathbf{B}}$ has $C_2 = y_2\Delta_{\mathbf{B}} \oplus K[y_1] \oplus M[y_2]$, and we have

$$x_2(C_1 \oplus C_2) = x_2(y_1\Delta_{\mathbf{A}} \oplus y_1\Delta_{\mathbf{B}} \oplus y_2\Delta_{\mathbf{A}} \oplus y_1\Delta_{\mathbf{B}}).$$

Therefore, this value can be computed by $P_{\mathbf{A}}$ sending one ciphertext to $P_{\mathbf{B}}$. Given the above observations, the final protocol can be derived in a straightforward way. Overall this new approach improves communication by $1.2\times$ and improves computation by $2\times$.

For details and a security proof corresponding to the above, see Section 5.1.

New function-dependent preprocessing. Here we show how to further improve the efficiency of function-dependent preprocessing. Recall that in the WRK protocol, each AND triple is derived from B leaky-AND triples, for $B \approx \frac{\rho}{\log C}$; these triples are then used to multiply authenticated masked values for each AND gate of the circuit. Our observation is that we can reduce the number of authenticated shares needed per gate from $3B + 2$ to $3B - 1$. This idea was initially used by Araki et al. [ABF⁺17] in the setting of honest-majority three-party computation. See Section 5.2 for details.

Protocol Π_{2pc}

Inputs: P_A holds $x \in \{0,1\}^{\mathcal{I}_1}$ and P_B holds $y \in \{0,1\}^{\mathcal{I}_2}$. Parties agree on a circuit for a function $f : \{0,1\}^{\mathcal{I}_1} \times \{0,1\}^{\mathcal{I}_2} \rightarrow \{0,1\}^{\mathcal{O}}$.

1. P_A and P_B call \mathcal{F}_{pre} , which sends Δ_A to P_A , Δ_B to P_B , and sends $\{(r_w | s_w)\}_{w \in \mathcal{I} \cup \mathcal{W}}$, $\{(r_w^* | s_w^*)\}_{w \in \mathcal{W}}$ to P_A and P_B . For each $w \in \mathcal{I}_1 \cup \mathcal{I}_2$, P_A also picks a uniform κ -bit string $L_{w,0}$.
2. Following the topological order of the circuit, for each gate $\mathcal{G} = (\alpha, \beta, \gamma, T)$,

- If $T = \oplus$, P_A computes $L_{\gamma,0} := L_{\alpha,0} \oplus L_{\beta,0}$
- If $T = \wedge$, P_A computes $L_{\alpha,1} := L_{\alpha,0} \oplus \Delta_A$, $L_{\beta,1} := L_{\beta,0} \oplus \Delta_A$, and

$$\begin{aligned} G_{\gamma,0} &:= H(L_{\alpha,0}, \gamma) \oplus H(L_{\alpha,1}, \gamma) \oplus K[s_\beta] \oplus r_\beta \Delta_A \\ G_{\gamma,1} &:= H(L_{\beta,0}, \gamma) \oplus H(L_{\beta,1}, \gamma) \oplus K[s_\alpha] \oplus r_\alpha \Delta_A \oplus L_{\alpha,0} \\ L_{\gamma,0} &:= H(L_{\alpha,0}, \gamma) \oplus H(L_{\beta,0}, \gamma) \oplus K[s_\gamma] \oplus r_\gamma \Delta_A \oplus K[s_\gamma^*] \oplus r_\gamma^* \Delta_A \\ b_\gamma &:= \text{lsb}(L_{\gamma,0}) \end{aligned}$$

P_A sends $G_{\gamma,0}, G_{\gamma,1}, b_\gamma$ to P_B .

3. For each $w \in \mathcal{I}_2$, two parties compute $r_w := \text{Open}([r_w]_A)$. P_B then sends $y_w \oplus \lambda_w := y_w \oplus s_w \oplus r_w$ to P_A . Finally, P_A sends $L_{w, y_w \oplus \lambda_w}$ to P_B .
4. For each $w \in \mathcal{I}_1$, two parties compute $s_w := \text{Open}([s_w]_B)$. P_A then sends $x_w \oplus \lambda_w := x_w \oplus s_w \oplus r_w$ and $L_{w, x_w \oplus \lambda_w}$ to P_B .
5. 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 $G_0 := G_{\gamma,0} \oplus M[s_\beta]$, and $G_1 := G_{\gamma,1} \oplus M[s_\alpha]$. P_B evaluates the garbled table (G_0, G_1) to obtain the output label

$$\begin{aligned} L_{\gamma, z_\gamma \oplus \lambda_\gamma} &:= H(L_{\alpha, z_\alpha \oplus \lambda_\alpha}, \gamma) \oplus H(L_{\beta, z_\beta \oplus \lambda_\beta}, \gamma) \oplus M[s_\gamma] \oplus M[s_\gamma^*] \\ &\quad \oplus (z_\alpha \oplus \lambda_\alpha) G_0 \oplus (z_\beta \oplus \lambda_\beta) (G_1 \oplus L_{\alpha, z_\alpha \oplus \lambda_\alpha}) \end{aligned}$$

and $z_\gamma \oplus \lambda_\gamma := b_\gamma \oplus \text{lsb}(L_{\gamma, z_\gamma \oplus \lambda_\gamma})$

6. For each $w \in \mathcal{W}$, P_B sends $\hat{z}_w := z_w \oplus \lambda_w$ to P_A .
7. For each AND gates $(\alpha, \beta, \gamma, \wedge)$, both parties know $\hat{z}_\alpha = z_\alpha \oplus \lambda_\alpha$, $\hat{z}_\beta = z_\beta \oplus \lambda_\beta$, and $\hat{z}_\gamma = z_\gamma \oplus \lambda_\gamma$. Two parties compute authenticated share of bit c_γ defined as

$$c_\gamma = (\hat{z}_\alpha \oplus \lambda_\alpha) \wedge (\hat{z}_\beta \oplus \lambda_\beta) \oplus (\hat{z}_\gamma \oplus \lambda_\gamma).$$

Note that c_γ is a linear combination of λ_α , λ_β , λ_γ and $\lambda_\gamma^* = \lambda_\alpha \wedge \lambda_\beta$, therefore authenticated share of c_γ can be computed locally.

8. Two parties use Open to check that c_γ is 0 for all gates γ , and abort if any check fails.
9. For each $w \in \mathcal{O}$, two parties compute $r_w := \text{Open}([r_w]_A)$. P_B computes $z_w := (\lambda_w \oplus z_w) \oplus r_w \oplus s_w$.

Figure 3: The main protocol in the \mathcal{F}_{pre} hybrid model

4 Technical Details: Improved Authenticated Garbling

Since we already discussed the main intuition of the protocol in the previous section, we will present our main protocol in the \mathcal{F}_{pre} -hybrid model. Detailed protocol description is shown in Figure 3. Each step in the protocol can be summarized as follows:

1. Parties generate circuit preprocessing information using \mathcal{F}_{pre} .
2. P_A computes its own share of the garbled circuit and sends to P_B .
- 3-4. Parties process P_A and P_B 's input and let P_B learn the corresponding masked input wire values and garbled labels.
5. P_B locally reconstructs the garbled circuit and evaluates it.
- 6-8. P_B sends all masked wire values (including all input, output, and internal wires) to P_A ; two parties check the correctness of all masked wire values.
9. P_A reveals the masks of output wires to P_B , who can recover the output.

Note that steps 2 through 9 are performed in the online phase, with $2\kappa + 2$ bits of communication per AND gate, $\kappa + 1$ bits of communication per input bit, and 1 bit of communication per output bit.

4.1 Proof of Security

We start by stating our main theorem.

Theorem 1. *If H is modeled as a random oracle, the protocol in Figure 3 securely computes f against malicious adversaries in the \mathcal{F}_{pre} -hybrid model.*

Before proceeding to the formal proof, we first introduce two important lemmas. The first lemma addresses correctness of our distributed garbling scheme in the semi-honest case; the second lemma addresses correctness of the whole protocol when P_A is corrupted.

Lemma 1. *When both parties follow the protocol honestly then, after step 5, for each wire w in the circuit P_B holds $(z_w \oplus \lambda_w, L_{w, z_w \oplus \lambda_w})$.*

Proof. We prove this by induction on the gates in the circuit.

Base case. It is easy to verify from step 3 and step 4 that the lemma holds for input wires.

Induction step. XOR-gates are trivial and so focus on an AND gate $(\alpha, \beta, \gamma, \wedge)$. First, the garbled tables are computed distributively, therefore we first write down the table after P_B merged its own share as follows. Note that we ignore the gate id (γ) for simplicity.

$$\begin{aligned}
 G_0 &= H(L_{\alpha,0}) \oplus H(L_{\alpha,1}) \oplus K[s_\beta] \oplus r_\beta \Delta_A \oplus M[s_\beta] \\
 &= H(L_{\alpha,0}) \oplus H(L_{\alpha,1}) \oplus \lambda_\beta \Delta_A \\
 G_1 &= H(L_{\beta,0}) \oplus H(L_{\beta,1}) \oplus K[s_\alpha] \oplus r_\alpha \Delta_A \oplus M[s_\alpha] \oplus L_{\alpha,0} \\
 &= H(L_{\beta,0}) \oplus H(L_{\beta,1}) \oplus \lambda_\alpha \Delta_A \oplus L_{\alpha,0}.
 \end{aligned}$$

P_A locally computes the output garbled label for 0 values, namely $L_{\gamma,0}$ as:

$$L_{\gamma,0} := H(L_{\alpha,0}) \oplus H(L_{\beta,0}) \oplus K[s_\gamma] \oplus r_\gamma \Delta_A \oplus K[s_\gamma^*] \oplus r_\gamma^* \Delta_A.$$

P_B , who holds $(z_\alpha \oplus \lambda_\alpha, L_{\alpha, z_\alpha \oplus \lambda_\alpha})$ and $(z_\beta \oplus \lambda_\beta, L_{\beta, z_\beta \oplus \lambda_\beta})$ by the induction hypothesis, evaluates the circuit as follows:

$$\begin{aligned} L_{\gamma, z_\gamma \oplus \lambda_\gamma} &:= H(L_{\alpha, z_\alpha \oplus \lambda_\alpha}) \oplus H(L_{\beta, z_\beta \oplus \lambda_\beta}) \oplus (z_\alpha \oplus \lambda_\alpha) G_0 \\ &\quad \oplus (z_\beta \oplus \lambda_\beta) (G_1 \oplus L_{\alpha, z_\alpha \oplus \lambda_\alpha}) \oplus M[s_\gamma] \oplus M[s_\gamma^*]. \end{aligned}$$

Observe that

$$\begin{aligned} &(z_\alpha \oplus \lambda_\alpha) G_0 \oplus H(L_{\alpha, z_\alpha \oplus \lambda_\alpha}) \\ &= (z_\alpha \oplus \lambda_\alpha) (H(L_{\alpha,0}) \oplus H(L_{\alpha,1}) \oplus \lambda_\beta \Delta_A) \oplus H(L_{\alpha, z_\alpha \oplus \lambda_\alpha}) \\ &= (z_\alpha \oplus \lambda_\alpha) (H(L_{\alpha,0}) \oplus H(L_{\alpha,1}) \oplus \lambda_\beta \Delta_A) \oplus (z_\alpha \oplus \lambda_\alpha) (H(L_{\alpha,0}) \oplus H(L_{\alpha,1})) \oplus H(L_{\alpha,0}) \\ &= H(L_{\alpha,0}) \oplus \lambda_\beta (z_\alpha \oplus \lambda_\alpha) \Delta_A, \end{aligned}$$

and

$$\begin{aligned} &(z_\beta \oplus \lambda_\beta) (G_1 \oplus L_{\alpha, z_\alpha \oplus \lambda_\alpha}) \oplus H(L_{\beta, z_\beta \oplus \lambda_\beta}) \\ &= (z_\beta \oplus \lambda_\beta) (H(L_{\beta,0}) \oplus H(L_{\beta,1}) \oplus \lambda_\alpha \Delta_A \oplus (z_\alpha \oplus \lambda_\alpha) \Delta_A) \oplus H(L_{\beta, z_\beta \oplus \lambda_\beta}) \\ &= (z_\beta \oplus \lambda_\beta) (H(L_{\beta,0}) \oplus H(L_{\beta,1}) \oplus z_\alpha \Delta_A) \oplus (z_\beta \oplus \lambda_\beta) (H(L_{\beta,0}) \oplus H(L_{\beta,1})) \oplus H(L_{\beta,0}) \\ &= H(L_{\beta,0}) \oplus (\lambda_\beta \oplus z_\beta) z_\alpha \Delta_A. \end{aligned}$$

Therefore, we conclude that

$$\begin{aligned} &L_{\gamma,0} \oplus L_{\gamma, z_\gamma \oplus \lambda_\gamma} \\ &= H(L_{\alpha,0}) \oplus H(L_{\beta,0}) \oplus H(L_{\alpha, z_\alpha \oplus \lambda_\alpha}) \oplus H(L_{\beta, z_\beta \oplus \lambda_\beta}) \oplus (z_\alpha \oplus \lambda_\alpha) G_0 \\ &\quad \oplus (z_\beta \oplus \lambda_\beta) (G_1 \oplus L_{\alpha, z_\alpha \oplus \lambda_\alpha}) \oplus \lambda_\gamma \Delta_A \oplus (\lambda_\alpha \wedge \lambda_\beta) \Delta_A \\ &= (\lambda_\alpha \oplus z_\alpha) \lambda_\beta \Delta_A \oplus (\lambda_\beta \oplus z_\beta) z_\alpha \Delta_A \oplus \lambda_\gamma \Delta_A \oplus (\lambda_\alpha \wedge \lambda_\beta) \Delta_A \\ &= ((z_\alpha \wedge z_\beta) \oplus \lambda_\gamma) \Delta_A = (z_\gamma \oplus \lambda_\gamma) \Delta_A. \end{aligned}$$

This means that, with respect to P_A 's definition of $L_{\gamma, z_\gamma \oplus \lambda_\gamma}$, P_B 's label is always correct. The masked value is correct because the least-significant bit of Δ_A is 1; thus,

$$\begin{aligned} b_\gamma \oplus \text{lsb}(L_{\gamma, z_\gamma \oplus \lambda_\gamma}) &= \text{lsb}(L_{\gamma,0}) \oplus \text{lsb}(L_{\gamma, z_\gamma \oplus \lambda_\gamma}) \\ &= \text{lsb}(L_{\gamma,0} \oplus L_{\gamma, z_\gamma \oplus \lambda_\gamma}) \\ &= \text{lsb}((z_\gamma \oplus \lambda_\gamma) \Delta_A) = z_\gamma \oplus \lambda_\gamma. \end{aligned}$$

□

Lemma 2. Let $x \stackrel{\text{def}}{=} \hat{x}_w \oplus \lambda_w$ and $y \stackrel{\text{def}}{=} \hat{y}_w \oplus \lambda_w$, where \hat{x}_w is what P_B sends in step 3, \hat{y}_w is what P_A sends in step 4, and λ_w is defined by \mathcal{F}_{pre} . If P_A is malicious, then P_B either aborts or outputs $f(x, y)$.

Proof. After step 5, P_B obtains a set of masked values $z_w \oplus \lambda_w$ for all wires w in the circuit. In the following, we will show that if these masked values are not correct, then P_B will abort with all but negligible probability.

Again we will prove by induction. Note that the lemma holds for all wires $w \in \mathcal{I}_1 \cup \mathcal{I}_2$, according to how x, y are defined, as well as for XOR-gates. In the following, we will focus on an AND gate $(\alpha, \beta, \gamma, \wedge)$. Now, according to induction hypothesis, we already know that P_B hold correct values of $(z_\alpha \oplus \lambda_\alpha, z_\beta \oplus \lambda_\beta)$.

Recall that the checking is done by computing

$$c = (\hat{z}_\alpha \oplus \lambda_\alpha) \wedge (\hat{z}_\beta \oplus \lambda_\beta) \oplus (\hat{z}_\gamma \oplus \lambda_\gamma).$$

The correctness of input masked values means that

$$c = z_\alpha \wedge z_\beta \oplus \hat{z}_\gamma \oplus \lambda_\gamma.$$

Since **Open** does not abort, $c = 0$, which means that $\hat{z}_\gamma = z_\alpha \wedge z_\beta \oplus \lambda_\gamma = z_\gamma \oplus \lambda_\gamma$. This means that the output masked wire value is also correct. \square

Given the above two lemmas, the proof of security of our main protocol is relatively easy. We provide all details below.

Proof. We consider separately a malicious P_A and P_B .

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. \mathcal{S} plays the role of \mathcal{F}_{pre} and records all values that \mathcal{F}_{pre} sends to two parties.
2. \mathcal{S} receives all values that \mathcal{A} sends.
3. \mathcal{S} acts as an honest P_B using input $y := 0$.
4. For each wire $w \in \mathcal{I}_1$, \mathcal{S} receives \hat{x}_w and computes $x_w := \hat{x}_w \oplus r_w \oplus s_w$, where r_w, s_w are the values used by \mathcal{F}_{pre} in the previous steps.
6. \mathcal{S} picks random bits for all \hat{z}_w and send them to \mathcal{A} .
- 7–9. \mathcal{S} acts as an honest P_B If an honest P_B would abort, \mathcal{S} aborts; otherwise \mathcal{S} computes the input x of \mathcal{A} . from the output of \mathcal{F}_{pre} and the values \mathcal{A} sent. \mathcal{S} then sends x to \mathcal{F} .

We show that the joint distribution of the outputs of \mathcal{A} and the honest P_B in the real world is indistinguishable from the joint distribution of the outputs of \mathcal{S} and P_B in the ideal world. We prove this by considering a sequence of experiments, the first of which corresponds to the execution of our protocol and the last of which corresponds to execution in the ideal world, and showing that successive experiments are computationally indistinguishable.

Hybrid₁. This is the hybrid-world protocol, where we imagine \mathcal{S} playing the role of an honest P_B using P_B 's actual input y , while also playing the role of \mathcal{F}_{pre} .

Hybrid₂. Same as **Hybrid₁**, except that in step 6, for each wire $w \in \mathcal{I}_1$ the simulator \mathcal{S} receives \hat{x}_w and computes $x_w := \hat{x}_w \oplus r_w \oplus s_w$, where r_w, s_w are the values used by \mathcal{F}_{pre} . If an honest P_B would abort in any later step, \mathcal{S} sends **abort** to \mathcal{F} ; otherwise it sends $x = \{x_w\}_{w \in \mathcal{I}_1}$ to \mathcal{F} .

The distributions on the view of \mathcal{A} in **Hybrid₁** and **Hybrid₂** are identical. The output P_B gets are the same due to Lemma 1 and Lemma 2.

Hybrid₃. Same as **Hybrid₂**, except that \mathcal{S} uses $y' = 0$ in step 3 and ignore what \mathcal{A} sends back. Then in step 6, \mathcal{S} sends random bits instead of the value for $z_w \oplus \lambda_w$.

The distributions on the view of \mathcal{A} in **Hybrid₃** and **Hybrid₂** are again identical (since the $\{s_w\}_{w \in \mathcal{I}_2}$ are uniform).

Note that **Hybrid₃** corresponds to the ideal-world execution described earlier. 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. \mathcal{S} plays the role of \mathcal{F}_{pre} and records all values sent to both parties.
2. \mathcal{S} acts as an honest P_A and send the shared garbled tables to P_B .
3. For each wire $w \in \mathcal{I}_2$, \mathcal{S} receives \hat{y}_w and computes $y_w := \hat{y}_w \oplus r_w \oplus s_w$, where r_w, s_w are the values used by \mathcal{F}_{pre} in the previous steps.
4. \mathcal{S} acts as an honest P_A using input $x = 0$.
- 6–8. \mathcal{S} acts as an honest P_A . If an honest P_A would abort, \mathcal{S} abort.
9. \mathcal{S} sends y computed in step 3 to \mathcal{F} , which returns $z = f(x, y)$. \mathcal{S} then computes $z' := f(0, y)$ and defines $r'_w = z_w \oplus z'_w \oplus r_w$ for each $w \in \mathcal{O}$. \mathcal{S} then acts as an honest P_A and opens values r'_w to \mathcal{A} . If an honest P_A would abort, \mathcal{S} outputs whatever \mathcal{A} outputs.

We now show that the distribution on the view of \mathcal{A} in the real world is indistinguishable from the distribution on the view of \mathcal{A} in the ideal world. (Note P_A has no output.)

Hybrid₁. This is the hybrid-world protocol, where \mathcal{S} acts as an honest P_A using P_A 's actual input x , while playing the role of \mathcal{F}_{pre} .

Hybrid₂. Same as **Hybrid₁**, except that in step 3, \mathcal{S} receives \hat{y}_w and computes $y_w := \hat{y}_w \oplus r_w \oplus s_w$, where r_w, s_w are the values used by \mathcal{F}_{pre} . If an honest P_A abort in any step, send **abort** to \mathcal{F} .

Hybrid₃. Same as **Hybrid₂**, except that in step 4, \mathcal{S} acts as an honest P_A with input $x = 0$. \mathcal{S} sends x computed in step 3 to \mathcal{F} , which returns $z = f(x, y)$. \mathcal{S} then computes $z' := f(0, y)$ and defines $r'_w = z_w \oplus z'_w \oplus r_w$ for each $w \in \mathcal{O}$. \mathcal{S} then acts as an honest P_A and opens values r'_w to \mathcal{A} . If an honest P_A would abort, \mathcal{S} outputs whatever \mathcal{A} outputs.

The distributions on the view of \mathcal{A} in **Hybrid₃** and **Hybrid₂** are identical.

Note that **Hybrid₃** is identical to the ideal-world execution. □

Functionality $\mathcal{F}_{\text{Land}}$

Honest case:

1. Generate uniform $\langle x_1 | x_2 \rangle, \langle y_1 | y_2 \rangle, \langle z_1 | z_2 \rangle$ such that $z_1 \oplus z_2 = (x_1 \oplus x_2) \wedge (y_1 \oplus y_2)$, and send the respective shares to the two parties.
2. P_A can choose to send $(P_1, p_2, P_3) \in \{0, 1\}^\kappa \times \{0, 1\} \times \{0, 1\}^\kappa$. The functionality checks

$$P_3 \oplus x_2 P_1 = (p_2 \oplus x_2 \text{lsb}(P_1)) \Delta_B.$$

If the check fails, the functionality sends fail to both parties and abort. (P_B can do the same symmetrically.)

Corrupted parties: A corrupted party gets to specify the randomness used on its behalf by the functionality.

Figure 4: Functionality $\mathcal{F}_{\text{Land}}$ for computing a leaky AND triple.

5 Technical Details: Improved Preprocessing

In this section, we provide details for our two optimizations of the preprocessing phase. The first optimization improves the efficiency to compute a leaky AND gate. Leaky AND gate is a key component towards a preprocessing with full security. This functionality ($\mathcal{F}_{\text{Land}}$) outputs triples with guaranteed correctness but the adversary can choose to guess the x value from the honest party: an incorrect guess will be caught immediately; while a correct guess remain undetected.

The second optimization focuses on how to combine leaky triples in a more efficient way. In particular, we observe that a recent optimization in the honest-majority secret sharing protocol by Araki et al. [ABF⁺17], can be applied to our setting too. As a result, we can roughly reduce the bucket size by one.

5.1 Improved Leaky AND

Before giving the details, we point out a minor difference in the leaky-AND functionality ($\mathcal{F}_{\text{Land}}$) as compared to [WRK17a]. As shown in Figure 4, instead of letting \mathcal{A} directly learn the value of x , the functionality allows \mathcal{A} to send a query in a form of (P_1, p_2, P_3) and return if $P_3 \oplus x_2 P_1 = (p_2 \oplus x_2 \text{lsb}(P_1)) \Delta_B$. It can be seen that this special way is no more than a query on x and two queries on Δ , and the \mathcal{A} cannot learn any information on y or z .

The main intuition of the protocol is already discussed in Section 3.2. We will proceed to present the protocol, in Figure 5.

Theorem 2. *The protocol in Figure 5 securely realizes $\mathcal{F}_{\text{Land}}$ in the $(\mathcal{F}_{\text{abit}}, \mathcal{F}_{\text{eq}})$ -hybrid model.*

Proof. As the first step, we will show that the protocol is correct if both parties are honest. We recall that

1. $G_1 := H(\mathcal{K}[x_2] \oplus \Delta_A) \oplus H(\mathcal{K}[x_2]) \oplus C_A$

Protocol Π_{Land}

Protocol:

1. P_A and P_B obtain random authenticated shares $(\langle x_1 | x_2 \rangle, \langle y_1 | y_2 \rangle, \langle z_1 | r \rangle)$.
 P_A locally computes $C_A := y_1 \Delta_A \oplus K[y_2] \oplus M[y_1]$, and
 P_B locally computes $C_B := y_2 \Delta_B \oplus M[y_2] \oplus K[y_1]$.
2. P_A sends $G_1 := H(K[x_2] \oplus \Delta_A) \oplus H(K[x_2]) \oplus C_A$ to P_B .
 P_B computes $E_1 := x_2 G_1 \oplus H(M[x_2]) \oplus x_2 C_B$.
3. P_B sends $G_2 := H(K[x_1] \oplus \Delta_B) \oplus H(K[x_1]) \oplus C_B$ to P_A .
 P_A computes $E_2 := x_1 G_2 \oplus H(M[x_1]) \oplus x_1 C_A$.
4. P_A computes $S_1 := H(K[x_2]) \oplus E_2 \oplus (z_1 \Delta_A \oplus K[r] \oplus M[z_1])$, P_B computes $S_2 := H(K[x_1]) \oplus E_1 \oplus (r \Delta_B \oplus M[r] \oplus K[z_1])$. P_A sends $\text{lsb}(S_1)$ to P_B ; P_B sends $\text{lsb}(S_2)$ to P_A . Both parties compute $d := \text{lsb}(S_1) \oplus \text{lsb}(S_2)$.
5. P_A sends $L_1 := S_1 \oplus d \Delta_A$ to \mathcal{F}_{eq} , P_B sends $L_2 := S_2 \oplus d \Delta_B$ to \mathcal{F}_{eq} . If \mathcal{F}_{eq} returns 0, parties abort, otherwise, they compute $[z_2]_B := [r]_B \oplus d$.

Figure 5: Our improved leaky-AND protocol.

2. $G_2 := H(K[x_1] \oplus \Delta_B) \oplus H(K[x_1]) \oplus C_B$
3. $C_A := y_1 \Delta_A \oplus K[y_2] \oplus M[y_1]$
4. $C_B := y_2 \Delta_B \oplus M[y_2] \oplus K[y_1]$

Note that

$$E_1 \oplus H(K[x_2]) = x_2 G_1 \oplus H(M[x_2]) \oplus x_2 C_B \oplus H(K[x_2]).$$

When $x_2 = 0$, we have

$$\begin{aligned} E_1 \oplus H(K[x_2]) &= x_2 G_1 \oplus H(M[x_2]) \oplus x_2 C_B \oplus H(K[x_2]) \\ &= H(M[x_2]) \oplus H(K[x_2]) \\ &= 0 = x_2(C_A \oplus C_B). \end{aligned}$$

When $x_2 = 1$, we have

$$\begin{aligned} E_1 \oplus H(K[x_2]) &= x_2 G_1 \oplus H(M[x_2]) \oplus x_2 C_B \oplus H(K[x_2]) \\ &= x_2(G_1 \oplus C_B) \oplus H(M[x_2]) \oplus H(K[x_2]) \\ &= x_2(G_1 \oplus C_B) \oplus H(K[x_2] \oplus \Delta_A) \oplus H(K[x_2]) \\ &= x_2(C_A \oplus C_B). \end{aligned}$$

Therefore,

$$E_1 \oplus H(K[x_2]) = x_2(C_A \oplus C_B)$$

$$\begin{aligned}
&= x_2(y_1\Delta_A \oplus K[y_2] \oplus M[y_1] \oplus y_2\Delta_B \oplus M[y_2] \oplus K[y_1]) \\
&= x_2(y_1\Delta_A \oplus y_2\Delta_A \oplus y_1\Delta_B \oplus y_2\Delta_B) \\
&= x_2(y_1 \oplus y_2)(\Delta_A \oplus \Delta_B).
\end{aligned}$$

Similarly,

$$E_2 \oplus H(K[x_1]) = x_1(y_1 \oplus y_2)(\Delta_A \oplus \Delta_B).$$

Taking these two equations, we know that

$$\begin{aligned}
S_1 \oplus S_2 &= (E_1 \oplus H(K[x_2])) \oplus (E_2 \oplus H(K[x_1])) \\
&\quad \oplus (z_1\Delta_A \oplus K[r] \oplus M[z_1] \oplus r\Delta_B \oplus M[r] \oplus K[z_1]) \\
&= (x_1 \oplus x_2)(y_1 \oplus y_2)(\Delta_A \oplus \Delta_B) \\
&\quad \oplus (z_1\Delta_A \oplus K[z_1] \oplus M[z_1] \oplus r\Delta_B \oplus K[r] \oplus M[r]) \\
&= (x_1 \oplus x_2)(y_1 \oplus y_2)(\Delta_A \oplus \Delta_B) \\
&\quad \oplus (z_1\Delta_A \oplus z_1\Delta_B \oplus r\Delta_B \oplus r\Delta_A) \\
&= (x_1 \oplus x_2)(y_1 \oplus y_2)(\Delta_A \oplus \Delta_B) \oplus (z_1 \oplus r)(\Delta_A \oplus \Delta_B) \\
&= ((x_1 \oplus x_2) \wedge (y_1 \oplus y_2) \oplus z_1 \oplus r)(\Delta_A \oplus \Delta_B).
\end{aligned}$$

Since $\text{lsb}(\Delta_A \oplus \Delta_B) = 1$, it holds that

$$d = \text{lsb}(S_1 \oplus S_2) = (x_1 \oplus x_2) \wedge (y_1 \oplus y_2) \oplus z_1 \oplus r.$$

Therefore, $(x_1 \oplus x_2) \wedge (y_1 \oplus y_2) = d \oplus z_1 \oplus r = z_1 \oplus z_2$.

Now we will focus on the security of the protocol in the malicious setting. First note that the protocol is symmetric, therefore we only need to focus on the case of a malicious P_A . The local computation of both parties is deterministic, with all inputs sent from $\mathcal{F}_{\text{abit}}$. Therefore, all messages sent during the protocol can be anticipated (emulated) by \mathcal{S} after \mathcal{S} sending out the shares. This is not always possible if \mathcal{A} uses local random coins or if \mathcal{A} has private inputs. This fact significantly reduces the difficulty of the proof. Intuitively, \mathcal{S} will be able to immediately catch \mathcal{A} cheating by comparing what it sends with what it would have sent (which \mathcal{S} knows by locally emulating). The majority of the work then is to extract \mathcal{A} 's attempt to perform a selective failure attack.

Define a simulator \mathcal{S} as follows.

- 0a. \mathcal{S} interacts with $\mathcal{F}_{\text{Land}}$ and obtains P_A 's share of $(\langle x_1 | x_2 \rangle, \langle y_1 | y_2 \rangle, \langle z_1 | z_2 \rangle)$. \mathcal{S} also gets Δ_A from $\mathcal{F}_{\text{abit}}$. \mathcal{S} randomly picks Δ_B and P_B 's share of $(\langle x_1 | x_2 \rangle, \langle y_1 | y_2 \rangle, \langle z_1 | z_2 \rangle)$ in a way that makes it consistent with P_A 's share. \mathcal{S} now randomly picks d and computes $[r]_B := [z_2]_B \oplus d$.
- 0b. Using values $(\langle x_1 | x_2 \rangle, \langle y_1 | y_2 \rangle, \langle z_1 | r \rangle)$ from both parties, \mathcal{S} locally emulates all messages sent by each party, namely (G_1, d_1, L_1) sent by an honest P_A and (G_2, d_2, L_2) sent by an honest P_B .
 1. \mathcal{S} plays the role of $\mathcal{F}_{\text{abit}}$ and sends out $(\langle x_1 | x_2 \rangle, \langle y_1 | y_2 \rangle, \langle z_1 | r \rangle)$ as defined above.
 2. \mathcal{S} acts as an honest P_B and receive G'_1 sent by \mathcal{A} . \mathcal{S} computes $P_1 = G'_1 \oplus G_1$.
 3. \mathcal{S} randomly picks a G_2 and send it to \mathcal{A} .
 4. \mathcal{S} acts as an honest P_B and receives d'_1 . \mathcal{S} computes $p_2 := d'_1 \oplus d_1$.

Protocol Π_{pre}

Inputs: Two parties agree on a circuit for a function $f : \{0, 1\}^{\mathcal{I}_1} \times \{0, 1\}^{\mathcal{I}_2} \rightarrow \{0, 1\}^{\mathcal{O}}$.

Protocol:

1. Two parties initialize $\mathcal{F}_{\text{abit}}$, which sends Δ_{A} to P_{A} and Δ_{B} to P_{B} .
2. For each wire $w \in \mathcal{I}_1 \cup \mathcal{I}_2 \cup \mathcal{W}$, two parties obtain an authenticated share $\langle r_w | s_w \rangle$ from $\mathcal{F}_{\text{abit}}$.
3. For each gate $\mathcal{G} = (\alpha, \beta, \gamma, \oplus)$, two parties compute $\langle r_\gamma | s_\gamma \rangle := \langle s_\alpha | r_\alpha \rangle \oplus \langle r_\beta | s_\beta \rangle$.
4. For each gate $\mathcal{G} = (\alpha, \beta, \gamma, \wedge)$, two parties have $(\langle r_\alpha | s_\alpha \rangle, \langle r_\beta | s_\beta \rangle)$, and run step 2 to step 5 in Π_{Land} to obtain $\langle r_\gamma^* | s_\gamma^* \rangle$, such that $r_\gamma^* \oplus s_\gamma^* = (r_\alpha \oplus s_\alpha) \wedge (r_\beta \oplus s_\beta)$
5. P_{A} and P_{B} call $\mathcal{F}_{\text{Land}}$ to obtain $(B - 1)|\mathcal{C}|$ number of leaky AND triples $(\langle x_1 | x_2 \rangle, \langle y_1 | y_2 \rangle, \langle z_1 | z_2 \rangle)$.
6. Two parties perform secure coin-flipping to determine a random permutation and permute the triples obtained in step 4. For each AND gate $\mathcal{G} = (\alpha, \beta, \gamma, \wedge)$ in the circuit, perform secure merging for $B - 1$ times.
 - (a) Obtain the next triple in the permuted list, namely $(\langle x_1 | x_2 \rangle, \langle y_1 | y_2 \rangle, \langle z_1 | z_2 \rangle)$
 - (b) Compute $\langle d_1 | d_2 \rangle := \langle y_1 | y_2 \rangle \oplus \langle r_\beta | s_\beta \rangle$, and $d := \text{Open}(\langle d_1 | d_2 \rangle)$.
 - (c) Update triple: $\langle r_\alpha | s_\alpha \rangle := \langle r_\alpha | s_\alpha \rangle \oplus \langle x_1 | x_2 \rangle$, $\langle r_\gamma^* | s_\gamma^* \rangle := \langle r_\gamma^* | s_\gamma^* \rangle \oplus \langle z_1 | z_2 \rangle \oplus d \langle x_1 | x_2 \rangle$.

Figure 6: Protocol Π_{pre} instantiating \mathcal{F}_{pre} in the $(\mathcal{F}_{\text{abit}}, \mathcal{F}_{\text{Land}})$ -hybrid model.

5. \mathcal{S} plays the role of \mathcal{F}_{eq} and obtain L_1 . \mathcal{S} computes $P_3 = L_1' \oplus L_1$. \mathcal{S} sends (P_1, p_2, P_3) to $\mathcal{F}_{\text{Land}}$ as the selective failure attack query. If $\mathcal{F}_{\text{Land}}$ abort, \mathcal{S} plays the role of \mathcal{F}_{eq} and aborts. If the value d in the protocol equals to r defined in step 0a, \mathcal{F}_{eq} returns 0; otherwise \mathcal{F}_{eq} returns 1.
6. \mathcal{S} sends (P_1, p_2, P_3) to $\mathcal{F}_{\text{Land}}$ as the selective failure query. If $\mathcal{F}_{\text{Land}}$ returns fail, \mathcal{S} sends 0 to \mathcal{A} as the output of \mathcal{F}_{eq} .

Note that messages that \mathcal{S} sends to \mathcal{A} in the protocol are changed from (G_2, d_2, L_2) to $(G_2, d_2 \oplus x_2 \text{lsb}(P_1), L_2 \oplus x_2 P_1 \oplus d' \Delta_{\text{B}})$, where $d' = p_2 \oplus x_2 \cdot \text{lsb}(P_1)$ and the equality checking in step 5 changed from comparing $L_1 = L_2$ to

$$L_1 \oplus P_3 = L_2 \oplus x_2 P_1 \oplus (p_2 \oplus x_2 \text{lsb}(P_1)) \Delta_{\text{B}},$$

that is

$$P_3 \oplus x_2 P_1 = (p_2 \oplus x_2 \text{lsb}(P_1)) \Delta_{\text{B}}.$$

This is the same form as the selective failure query in $\mathcal{F}_{\text{Land}}$. □

	One-way Communication (Max)				Two-way Communication			
	Ind. (MB)	Dep. (MB)	Online (KB)	Total (MB)	Ind. (MB)	Dep. (MB)	Online (KB)	Total (MB)
Single execution								
[NST17]	15	0.22	16	15	15	0.22	16	15
[WRK17a]	2.9	0.57	4.9	3.4	5.7	0.57	6.0	6.3
[HIV17]	-	3.4	≥ 4.9	3.4	-	3.4	≥ 4.9	3.4
This work, v. 1	1.9	0.33	5.0	2.2	3.8	0.33	5.0	4.2
This work, v. 2	2.5	0.22	5.0	2.7	4.9	0.22	5.0	5.1
Amortized cost over 1024 executions								
[RR16]	-	1.6	17	1.6	-	3.2	17	3.2
[NST17]	6.4	0.22	16	6.6	6.4	0.22	16	6.6
[KNR ⁺ 17]	-	1.6	19	1.6	-	1.6	19	1.6
[WRK17a]	2.0	0.57	4.9	2.6	4.0	0.57	6.0	4.6
This work, v. 1	1.4	0.33	5.0	1.7	2.7	0.33	5.0	3.1
This work, v. 2	1.9	0.22	5.0	2.1	3.8	0.22	5.0	4.0

Table 2: **Communication complexity of different protocols for evaluating AES, rounded to two significant figures.** As in Table 1, one-way communication refers to the maximum communication one party sends to the other; two-way communication refers to the sum of both parties’ communication. The best prior number in each column is bolded for reference.

5.2 Improved Function-Dependent Preprocessing

In this section, we will focus on improving the preprocessing in the Leaky AND triple generation ($\mathcal{F}_{\text{Land}}$) hybrid model. The main observation is that in the protocol of WRK, each wire is associated with a mask (in the authenticated share format). Then the AND of input masks are computed using one AND triple. This is a waste of randomness, since we also directly construct all triples in place for all wires. Note that the idea is similar to Araki et al. [ABF⁺17]. The detailed protocol is presented in Figure 6

Note that although the above optimization aims to reduce the overall cost of the protocol, but it turns out that even in this case, most of the computation and communication (including computation of all authenticated bits as well as all leaky-AND triples in step 5) can be still done in the function-independent phase. The function-dependent cost is increased by only κ bits per AND gate only. Therefore, here we have an option to trade-off between total communication and communication in the offline stage. By increasing the function-dependent cost by κ bits per gate, we reduce bucket size by 1. We believe both versions can be useful depending on the application, and the concrete cost of both versions of the protocol are presented in the performance section.

6 Performance

In this section, we discuss the concrete efficiency of our protocol. We consider two variants of our protocol that optimize the cost of different phases: The first version of our protocol is optimized to minimize the total communication; the second version is optimized to minimize the communication in the function-dependent phase. (The cost of the online phase is identical in both versions.)

	Ind.	Dep.	Online	Total
WRK	$10B$	8	2	$10B + 10$
This work, v. 1	$4B - 4$	8	2	$4B + 6$
This work, v. 2	$4B$	4	2	$4B + 6$

Table 3: **Number of H -evaluations.** We align the security parameters in both protocols and set $B = \rho / \log C + 1$ for a fair comparison.

6.1 Communication Complexity

Table 2 shows the communication complexity of recent two-party computation protocols in the malicious setting. Numbers for these protocols are obtained from the respective papers, while numbers for our protocol are calculated. We tabulate both one-way communication and total communication. If parties’ data can be sent at the same time over a full-duplex network, then one-way communication is a better reflection of the running time. In general, for a circuit that requires a bucket size of B , we can obtain an estimation of the concrete communication cost: our first version has function dependent cost of 3κ per gate, and function independent cost of $(4B - 2)\kappa + (3B - 1)\rho$ per gate; our second version has a function dependent cost of 2κ per gate, and a function independent cost of $(4B + 2)\kappa + (3B + 2)\rho$ per gate.

We see that our protocol and the protocol by Nielsen et al. [NST17] are the only ones that, considering the function-dependent phase and the online phase, have cost similar to that of the state-of-the-art semi-honest garbled-circuit protocol. In other words, *the overhead induced by malicious security can be completely pushed to the preprocessing stage*. Compared to the protocol by Nielsen et al., we are able to reduce the communication in the preprocessing stage by $6\times$ in the single-execution setting, and by $3.4\times$ in the amortized setting. Our protocol also has the best total communication complexity in both settings, excepting the work of [RR16, KNR⁺17] which are 6% better but do not support function-independent preprocessing.

6.2 Computational Complexity

Since the WRK protocol represents the state-of-the-art as far as implementations are concerned, we compare the computational complexity of our protocol to theirs. We also include a comparison to the more recent protocol by Hazay et al. [HIV17] (the *HIV protocol*), which has not yet been implemented.

Comparing to the WRK protocol. Our protocol follows the same high-level approach as the WRK protocol. Almost all H -evaluations in our protocol can be accelerated using fixed-key AES, as done in [BHKR13]. We tabulate the number of H -evaluations for both protocols in Table 3. Due to our improved $\mathcal{F}_{\text{Land}}$, we are able to achieve a 2–2.5 \times improvement.

Comparing to the HIV protocol. As noted by the authors, the HIV protocol has polylogarithmic computational overhead compared to semi-honest garbled circuits. This is due to their use of the MPC-based zero-knowledge proof by Ames et al. [AHIV17]. On the other hand, in our protocol, the computation is linear in the circuit size. Furthermore, almost all cryptographic operations in our protocol can be accelerated using hardware AES instructions.

Taking an AES circuit as example, the ZK protocol by Ames et al. for a circuit of that size has

a prover running time of around 70 ms and a verifier running time of around 30 ms. Therefore, even if we ignore the cost of computing and sending the garbled circuit, the oblivious transfers, and other operations, the end-to-end running time of the HIV protocol will still be at least 100 ms. On the other hand, the entire WRK protocol runs within 17 ms for the same circuit. As our protocol results in at least a $2\times$ improvement, our protocol will be at least an order of magnitude faster than the HIV protocol.

7 Challenges in Extending to the Multi-Party Case

Wang et al. [WRK17b] have also shown how to extend their authenticated-garbling protocol to the multi-party case. In this section, we discuss the challenges involved in applying our new techniques to that setting. Note that Ben-Efraim et al. [BE17] recently proposed new techniques for multi-party garbling, making it compatible with some of the half-gate optimizations. Despite being based on half-gates, they still require 4 garbled rows per AND gate, and thus their work still leaves open the question of reducing the communication complexity of the online phase in the multi-party case.

In the multi-party WRK protocol, there are $n - 1$ garbling parties and one evaluating party. For each wire, each garbler chooses their own set of wire labels (called “subkeys”). As in the 2-party case, the preprocessing defines some authenticated bits, and as a result all parties can locally compute additive shares of *any garbler’s* subkey corresponding to *any authenticated value*.

In each gate, each garbler P_i generates standard Yao garbled gate consisting of 4 rows. Each row of P_i ’s gate is encrypted by only P_i ’s subkeys, and the payload of the row is P_i ’s shares of *all garblers’* subkeys. That way, the evaluator can decrypt the correct row of everyone’s garbled gates, obtain everyone’s shares of everyone’s subkeys, and combine them to get everyone’s appropriate subkey for the output wire.

Now suppose we modify things so each garbler generates a half-gates-style garbled gate instead of a standard Yao garbled gate. The half-gate uses garbler P_i ’s subkeys as its “keys” and encodes P_i ’s shares of all subkeys as its “payloads”. Now the protocol may not be secure against an adversary corrupting the evaluator and a garbler. In particular, half-gates garbling defines $G_0 = H(L_{\alpha,0}) \oplus H(L_{\alpha,1}) \oplus \lambda_\beta \Delta$. When P_i is acting as garbler, these $L_{\alpha,u}$ values correspond to P_i ’s subkeys. Now suppose P_i colludes with the evaluator. If the evaluator comes to learn G_0 (which is necessary to evaluate the gate in half of the cases), then the adversary can learn the secret mask λ_β since it is the only unknown term in G_0 . Clearly revealing the secret wire mask breaks the privacy of the protocol. This is not a problem with Yao garbled gates, where each row can be written as $G_{u,v} = H(L_{\alpha,u}, L_{\beta,v}) \oplus [\text{payload already known to garbler}]$. The secret masks do not appear in the garbled table, except indirectly through the payloads (subkey shares).

It is even unclear if row-reduction can be made possible. In the multi-party setting, the garbler has no control over the “payload” (i.e., output wire label) of the garbled gate when using row-reduction. Indeed, this is what makes it possible to reduce the size of a garbled gate. This is not a problem in the two-party case, where there is only one garbler who has control over all garbled gates and all wire labels. He generates a garbled table, and then computes his output wire label (subkey) as a function of the payload in the table. However, in the multi-party case, P_i generates a half-gate whose payloads include P_i ’s shares of P_j ’s subkeys! We would need P_j ’s choice of subkeys to depend on the payloads of P_i ’s garbling (for all i and j !). It is not clear how this can be done, and even if it were possible it would apparently require additional rounds proportional to the depth of the circuit.

Acknowledgments

This material is based on work supported by NSF awards #1111599, #1563722, #1564088, and #1617197. Portions of this work were also supported by DARPA and SPAWAR under contract N66001-15-C-4065. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes not withstanding any copyright notation thereon. The views, opinions, and/or findings expressed are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government. The authors would like to thank Ruiyu Zhu for his helpful comments.

References

- [ABF⁺17] Toshinori Araki, Assi Barak, Jun Furukawa, Tamar Lichter, Yehuda Lindell, Ariel Nof, Kazuma Ohara, Adi Watzman, and Or Weinstein. Optimized honest-majority MPC for malicious adversaries - breaking the 1 billion-gate per second barrier. In *2017 IEEE Symposium on Security and Privacy*, pages 843–862, San Jose, CA, USA, May 22–26, 2017. IEEE Computer Society Press.
- [AHIV17] Scott Ames, Carmit Hazay, Yuval Ishai, and Muthuramakrishnan Venkatasubramanian. Liger: Lightweight sublinear arguments without a trusted setup. In *ACM CCS 17*, pages 2087–2104. ACM Press, 2017.
- [AMPR14] Arash Afshar, Payman Mohassel, Benny Pinkas, and Ben Riva. Non-interactive secure computation based on cut-and-choose. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 387–404, Copenhagen, Denmark, May 11–15, 2014. Springer, Heidelberg, Germany.
- [BE17] Aner Ben-Efraim. On multiparty garbling of arithmetic circuits. Cryptology ePrint Archive, Report 2017/1186, 2017. <https://eprint.iacr.org/2017/1186>.
- [Bea92] Donald Beaver. Efficient multiparty protocols using circuit randomization. In Joan Feigenbaum, editor, *CRYPTO'91*, volume 576 of *LNCS*, pages 420–432, Santa Barbara, CA, USA, August 11–15, 1992. Springer, Heidelberg, Germany.
- [BHKR13] Mihir Bellare, Viet Tung Hoang, Sriram Keelveedhi, and Phillip Rogaway. Efficient garbling from a fixed-key blockcipher. In *2013 IEEE Symposium on Security and Privacy*, pages 478–492, Berkeley, CA, USA, May 19–22, 2013. IEEE Computer Society Press.
- [BMR90] Donald Beaver, Silvio Micali, and Phillip Rogaway. The round complexity of secure protocols. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing*, pages 503–513. ACM, 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 - (extended abstract). In Kazuo Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 441–463, Bangalore, India, December 1–5, 2013. Springer, Heidelberg, Germany.

- [DPSZ12] Ivan Damgård, Valerio Pastro, Nigel P. Smart, and Sarah Zakarias. Multiparty computation from somewhat homomorphic encryption. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 643–662, Santa Barbara, CA, USA, August 19–23, 2012. Springer, Heidelberg, Germany.
- [FJN⁺13] Tore Kasper Frederiksen, Thomas Pelle Jakobsen, Jesper Buus Nielsen, Peter Sebastian Nordholt, and Claudio Orlandi. MiniLEGO: Efficient secure two-party computation from general assumptions. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 537–556, Athens, Greece, May 26–30, 2013. Springer, Heidelberg, Germany.
- [GLNP15] Shay Gueron, Yehuda Lindell, Ariel Nof, and Benny Pinkas. Fast garbling of circuits under standard assumptions. In Indrajit Ray, Ninghui Li, and Christopher Kruegel, editors, *ACM CCS 15*, pages 567–578, Denver, CO, USA, October 12–16, 2015. ACM Press.
- [HEKM11] Yan Huang, David Evans, Jonathan Katz, and Lior Malka. Faster secure two-party computation using garbled circuits. In *USENIX Security 2011*, 2011.
- [HIV17] Carmit Hazay, Yuval Ishai, and Muthuramakrishnan Venkatasubramanian. Actively secure garbled circuits with constant communication overhead in the plain model. In *TCC 2017, Part II*, *LNCS*, pages 3–39. Springer, Heidelberg, Germany, March 2017.
- [HKE13] Yan Huang, Jonathan Katz, and David Evans. Efficient secure two-party computation using symmetric cut-and-choose. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 18–35, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Heidelberg, Germany.
- [HKK⁺14] Yan Huang, Jonathan Katz, Vladimir Kolesnikov, Ranjit Kumaresan, and Alex J. Malozemoff. Amortizing garbled circuits. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 458–475, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Heidelberg, Germany.
- [IKNP03] Yuval Ishai, Joe Kilian, Kobbi Nissim, and Erez Petrank. Extending oblivious transfers efficiently. In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 145–161, Santa Barbara, CA, USA, August 17–21, 2003. Springer, Heidelberg, Germany.
- [KMR14] Vladimir Kolesnikov, Payman Mohassel, and Mike Rosulek. FlexOR: Flexible garbling for XOR gates that beats free-XOR. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 440–457, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Heidelberg, Germany.
- [KNR⁺17] Vladimir Kolesnikov, Jesper Buus Nielsen, Mike Rosulek, Ni Trieu, and Roberto Trifiletti. DUPLO: Unifying cut-and-choose for garbled circuits. In *ACM CCS 17*, pages 3–20. ACM Press, 2017.
- [KS08] Vladimir Kolesnikov and Thomas Schneider. Improved garbled circuit: Free XOR gates and applications. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M.

- Halldórsson, Anna Ingólfssdóttir, and Igor Walukiewicz, editors, *ICALP, Part II*, volume 5126 of *LNCS*, pages 486–498, Reykjavik, Iceland, July 7–11, 2008. Springer, Heidelberg, Germany.
- [KSS12] Benjamin Kreuter, Abhi Shelat, and Chih-Hao Shen. Billion-gate secure computation with malicious adversaries. In *USENIX Security 2012*, 2012.
- [Lin13] Yehuda Lindell. Fast cut-and-choose based protocols for malicious and covert adversaries. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 1–17, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Heidelberg, Germany.
- [LP07] Yehuda Lindell and Benny Pinkas. An efficient protocol for secure two-party computation in the presence of malicious adversaries. In Moni Naor, editor, *EUROCRYPT 2007*, volume 4515 of *LNCS*, pages 52–78, Barcelona, Spain, May 20–24, 2007. Springer, Heidelberg, Germany.
- [LP11] Yehuda Lindell and Benny Pinkas. Secure two-party computation via cut-and-choose oblivious transfer. In Yuval Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 329–346, Providence, RI, USA, March 28–30, 2011. Springer, Heidelberg, Germany.
- [LR14] Yehuda Lindell and Ben Riva. Cut-and-choose Yao-based secure computation in the online/offline and batch settings. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 476–494, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Heidelberg, Germany.
- [LR15] Yehuda Lindell and Ben Riva. Blazing fast 2PC in the offline/online setting with security for malicious adversaries. In Indrajit Ray, Ninghui Li, and Christopher Kruegel, editors, *ACM CCS 15*, pages 579–590, Denver, CO, USA, October 12–16, 2015. ACM Press.
- [MNPS04] Dahlia Malkhi, Noam Nisan, Benny Pinkas, and Yaron Sella. Fairplay—a secure two-party computation system. In *USENIX Security 2004*, 2004.
- [NNOB12] Jesper Buus Nielsen, Peter Sebastian Nordholt, Claudio Orlandi, and Sai Sheshank Burra. A new approach to practical active-secure two-party computation. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 681–700, Santa Barbara, CA, USA, August 19–23, 2012. Springer, Heidelberg, Germany.
- [NO09] Jesper Buus Nielsen and Claudio Orlandi. LEGO for two-party secure computation. In Omer Reingold, editor, *TCC 2009*, volume 5444 of *LNCS*, pages 368–386. Springer, Heidelberg, Germany, March 15–17, 2009.
- [NO16] Jesper Buus Nielsen and Claudio Orlandi. Cross and clean: Amortized garbled circuits with constant overhead. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part I*, volume 9985 of *LNCS*, pages 582–603, Beijing, China, October 31 – November 3, 2016. Springer, Heidelberg, Germany.

- [NPS99] Moni Naor, Benny Pinkas, and Reuban Sumner. Privacy preserving auctions and mechanism design. In *1st ACM Conference on Electronic Commerce*, 1999.
- [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.
- [PSSW09] Benny Pinkas, Thomas Schneider, Nigel P. Smart, and Stephen C. Williams. Secure two-party computation is practical. In Mitsuru Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 250–267, Tokyo, Japan, December 6–10, 2009. Springer, Heidelberg, Germany.
- [RR16] Peter Rindal and Mike Rosulek. Faster malicious 2-party secure computation with online/offline dual execution. In *USENIX Security 2016*, 2016.
- [SS11] Abhi Shelat and Chih-Hao Shen. Two-output secure computation with malicious adversaries. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 386–405, Tallinn, Estonia, May 15–19, 2011. Springer, Heidelberg, Germany.
- [SS13] Abhi Shelat and Chih-Hao Shen. Fast two-party secure computation with minimal assumptions. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 13*, pages 523–534, Berlin, Germany, November 4–8, 2013. ACM Press.
- [WMK17] Xiao Wang, Alex J. Malozemoff, and Jonathan Katz. Faster secure two-party computation in the single-execution setting. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part II*, volume 10211 of *LNCS*, pages 399–424, Paris, France, May 8–12, 2017. Springer, Heidelberg, Germany.
- [WRK17a] Xiao Wang, Samuel Ranellucci, and Jonathan Katz. Authenticated garbling and efficient maliciously secure two-party computation. In *ACM CCS 17*, pages 21–37. ACM Press, 2017.
- [WRK17b] Xiao Wang, Samuel Ranellucci, and Jonathan Katz. Global-scale secure multiparty computation. In *ACM CCS 17*, pages 39–56. ACM Press, 2017.
- [Yao86] Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *27th FOCS*, pages 162–167, Toronto, Ontario, Canada, October 27–29, 1986. IEEE Computer Society Press.
- [ZH17] Ruiyu Zhu and Yan Huang. JIMU: Faster LEGO-based secure computation using additive homomorphic hashes. In *ASIACRYPT 2017, Part II*, LNCS, pages 529–572. Springer, Heidelberg, Germany, December 2017.
- [ZRE15] Samee Zahur, Mike Rosulek, and David Evans. Two halves make a whole - reducing data transfer in garbled circuits using half gates. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 220–250, Sofia, Bulgaria, April 26–30, 2015. Springer, Heidelberg, Germany.