

Global-Scale Secure Multiparty Computation

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

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

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

ABSTRACT

We propose a new, constant-round protocol for multi-party computation of boolean circuits that is secure against an arbitrary number of malicious corruptions. At a high level, we extend and generalize recent work of Wang et al. in the two-party setting. Namely, we design an efficient preprocessing phase that allows the parties to generate authenticated information; we then show how to use this information to distributively construct a *single* “authenticated” garbled circuit that is evaluated by one party.

Our resulting protocol improves upon the state-of-the-art both asymptotically and concretely. We validate these claims via several experiments demonstrating both the *efficiency* and *scalability* of our protocol:

- **Efficiency:** For three-party computation over a LAN, our protocol requires only 95 ms to evaluate AES. This is roughly a 700× improvement over the best prior work, and only 2.5× slower than the best known result in the *two*-party setting. In general, for n -party computation our protocol improves upon prior work (which was never implemented) by a factor of more than $230n$, e.g., an improvement of 3 orders of magnitude for 5-party computation.
- **Scalability:** We successfully executed our protocol with a large number of parties located all over the world, computing (for example) AES with 128 parties across 5 continents in under 3 minutes. Our work represents the largest-scale demonstration of secure computation to date.

CCS CONCEPTS

• **Theory of computation** → **Cryptographic protocols**;

KEYWORDS

Multi-party Computation; Secure Computation; Garbled Circuit

1 INTRODUCTION

Secure multi-party computation (MPC) allows a set of parties to jointly perform a distributed computation while ensuring correctness, privacy of the parties’ inputs, and more. MPC protocols can be classified in various ways depending on the native computations

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CCS '17, October 30-November 3, 2017, Dallas, TX, USA

© 2017 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

ACM ISBN 978-1-4503-4946-8/17/10...\$15.00

<https://doi.org/10.1145/3133956.3133979>

Paper	Comm./Comp. Complexity	Rounds
[12]	$O(C B^3n)$	$O(d)$
[20]	$O(C B^2n)$	$O(d)$
[31] + [28]	$O(C \kappa n^2)$	$O(1)$
[25] (concurrent)	$O(C B^2n)$	$O(1)$
This paper	$O(C Bn)$	$O(1)$

Table 1: Asymptotic complexity (per party) for n -party MPC protocols for boolean circuits, secure against an arbitrary number of malicious corruptions. Here, κ (resp., ρ) is the computational (resp., statistical) security parameter, $|C|$ is the circuit size, d is the depth of the circuit, and $B = O(\rho/\log |C|)$.

they support and the class of adversarial behavior they tolerate. With regard to the first of these, protocols are typically designed to compute either boolean circuits or arithmetic circuits over a large field. Although these models are equivalent in terms of their expressive power, for many natural computational tasks (e.g., comparisons, divisions, bit-wise operations, etc.) an arithmetic circuits can be much larger than the corresponding boolean circuit, as well as more cumbersome to design.

With regard to security, some protocols tolerate only semi-honest adversaries that are assumed to follow the prescribed protocol but then try to learn additional information from the transcript of the execution. In contrast, the stronger malicious model does not make any assumptions about the behavior of the corrupted parties. Finally, some protocols assume an honest majority (i.e., are secure only as long as strictly fewer than 1/2 of the parties are corrupted), while others are secure for any number of corruptions (i.e., even if only of the parties is honest).

In this work we focus on MPC protocols tolerating any number of malicious corruptions. Previous implementations of MPC protocols in this model rely on some variant of the secret-sharing paradigm introduced by Goldreich, Micali, and Wigderson [23]. At a high level, this technique requires the parties to maintain the invariant of holding a linear secret sharing of the values on the circuit wires, along with some sort of authentication information on those shares. Linear gates in the circuit (e.g., XOR, ADD) can be processed locally, while non-linear operations (e.g., AND, MULT) are handled by having the parties interact with each other to maintain the desired invariant. The most notable example of a protocol in this framework is perhaps SPDZ [18, 19, 28], which supports arithmetic circuits; protocols for boolean circuits have also been designed [20, 30].

Although this approach can lead to protocols with reasonable efficiency when run over a LAN, it suffers the inherent drawback of leading to round complexity *linear* in the depth of the circuit being evaluated. This can have a significant impact on the overall efficiency when the parties running the protocol are geographically

Setting	Setup	Indep.	Depen.	Online	Total
3PC-LAN	36	47	12	2	95
128PC-LAN	390	2727	11670	1870	16657
14PC-Worldwide	8711	9412	1947	250	20320
128PC-Worldwide	88056	30796	22659	2316	143827

Table 2: Selected performance results for our protocol. All results are in milliseconds, based on a statistical security of 2^{-40} . We consider the following settings (see Section 6 for more details):

- (a) **3PC-LAN**: three-party computation over a LAN;
- (b) **128PC-LAN**: 128-party computation over a LAN;
- (c) **14PC-Worldwide**: 14-party computation over a WAN, with parties located in 14 different cities across five continents;
- (d) **128PC-Worldwide**: 128-party computation over a WAN, with parties located in 8 different cities across five continents (each city with 16 parties).

separated or when the number of parties is high, and the communication latency dominates the cost of the execution. For example, the communication latency between parties located in the U.S. and Europe is around 75 ms even with the dedicated network provided by Amazon EC2. If such parties are evaluating, say, SHA-256 (which has a circuit depth of about 4,000), then a linear-round protocol requires 300,000 ms just for the back-and-forth interaction between those parties, not even counting the time required for performing local cryptographic operations or transmitting any data.

Constant-round protocols tolerating any number of malicious corruptions have also been designed. The basic approach here, first proposed by Beaver, Micali, and Rogaway [3], is to have the parties run a linear-round secure-computation protocol to compute a garbled circuit [38] for the function f of interest; the parties can then evaluate that garbled circuit using a constant number of additional rounds. Since the circuit for computing the garbling of f has depth independent of f , the overall number of rounds is constant. Although several recent papers have explored this approach [14, 31, 32], these investigations have remained largely theoretical since the overall cost of even the best protocol using this approach is asymptotically worse than the nonconstant-round protocols mentioned above; see Table 1. In fact, prior implementations of constant-round MPC consider only the *semi-honest* setting [5, 6].

1.1 Our Contributions

In this paper, we take a significant step towards practical MPC tolerating an unbounded number of malicious corruptions. To this end, we propose a new, constant-round protocol for multi-party computation of boolean circuits secure in this setting. Our protocol extends and generalizes the recent work of Wang et al. [37] in the two-party setting. Specifically, we first design an optimized, multi-party version of their TinyOT protocol so as to enable n parties to generate certain authenticated information as part of a preprocessing phase. Next, generalizing their main protocol, we show how to use this information to distributively construct a *single* “authenticated” garbled circuit that is evaluated by a single party. An overview of our entire protocol appears in Section 3, with details in the remainder of the paper.

Our protocol improves upon the state-of-the-art both asymptotically (cf. Table 1) and concretely (see Section 6.4), and in particular it allows us to give the *first* implementation of a constant-round MPC protocol with malicious security. Our experiments demonstrate that our protocol is both *efficient* and *scalable*:

- **Efficiency**: For three-party computation over a LAN, our protocol requires only 95 ms to securely evaluate AES. This is roughly a $700\times$ improvement over the best prior work, and only $2.5\times$ slower than the best known result in the *two*-party setting [37]. In general, for n -party computation our protocol improves upon the best prior work [31, 32] (which was not implemented) by a factor of more than $200n$.
- **Scalability**: We successfully executed our protocol with many parties located all over the world, computing (for example) AES with 128 parties across 5 continents in under 3 minutes. To the best of our knowledge, our work represents the largest-scale demonstration of secure computation to date, even considering weaker adversarial models.

Selected performance results for our protocol are reported in Table 2. Following [34], we divide execution into different phases:

- **Setup**. Here, the parties generate information that can be used for computation of multiple, possibly different functions. For example, base-OT can be performed in this phase.
- **Function-independent preprocessing**. Here, the parties begin execution of the protocol for a particular computation. At this point, the parties only need to know an upper bound on the size of the circuit that will be computed.
- **Function-dependent preprocessing**. Here the parties know the function being computed, but need not know their inputs.
- **Online phase**. The parties evaluate the function on their inputs.

1.2 Related Work

Implementations of MPC protocols. The first implementations of generic MPC assumed a semi-honest adversary corrupting a minority of the parties. Early work in this area includes FairplayMP [5] for boolean circuits, and VIFF [15] and SEPIA [11] for arithmetic circuits. Implementations of protocols handling an arbitrary number of corrupted parties, but still in the semi-honest setting, were shown by Choi et al. [13] and Ben-Efraim et al. [6], the latter running in a constant number of rounds.

There are fewer implementations of MPC protocols handling *malicious* attackers. Jakobsen et al. [26] developed the first such system. SPDZ and its subsequent improvements [18, 19, 28, 29] greatly improved the efficiency. As noted earlier, all existing implementations of MPC tolerating malicious attackers have round complexity linear in the depth of the circuit.

Another line of work has specifically targeted *three*-party computation. Implementations here include Sharemind [8, 9], the sugar-beet auction run by Bogetoft et al. [10], and the recent work of Araki et al. [2]; these each tolerate only semi-honest behavior. Mohassel et al. [33] and Furukawa et al. [21] tolerate a malicious attacker corrupting only one party.

Constant-round MPC. Techniques for building constant-round MPC protocols have been studied in various settings. As noted

above, FairplayMP [5] and Ben-Efraim et al. [6] implemented this approach in the semi-honest setting. Other researchers have proposed approaches without providing an implementation, possibly because an implementation would be too complex or because the concrete efficiency of the resulting protocol would be uncompetitive with nonconstant-round protocols. As examples, Damgård and Ishai [16] proposed a protocol making black-box use of the underlying cryptographic primitives, and Choi et al. [14] looked at the three-party setting with malicious corruption of two parties. Lindell et al. [31, 32] considered optimizations of the BMR approach in the malicious setting. We compare the efficiency of our protocol with relevant prior work in Section 6.4.

Concurrent work. Independent of our work, Hazay et al. [25] proposed and implemented a constant-round MPC protocol with malicious security against any number of corruptions. The complexities of the online the function-dependent preprocessing phases of their protocol are similar to ours, though the underlying technique is very different. On the other hand, our function-independent processing phase is more efficient than theirs both asymptotically and concretely, and appears to be $3 \times -5 \times$ faster in practice. See Section 6.4 for a more detailed comparison.

2 NOTATION AND PRELIMINARIES

We use κ and ρ to denote the computational and statistical security parameters, respectively. We use $=$ to denote equality and $:=$ to denote assignment.

A circuit is viewed as a list of gates of the form $(\alpha, \beta, \gamma, T)$, where this represents a gate with input wires α and β , output wire γ , and gate type $T \in \{\oplus, \wedge\}$. Parties are denoted by P_1, \dots, P_n . We use I_i to denote the set of input-wire indices for P_i , \mathcal{W} to denote the set of output-wire indices for all AND gates, and \mathcal{O} to denote the set of output-wire indices of the circuit. (We assume all parties learn the output.) \mathcal{M} is used to denote the set of all corrupted parties, with $\mathcal{H} = [n] \setminus \mathcal{M}$ denoting the set of all honest parties.

Our protocol operates by having the parties distributively construct a garbled circuit that is evaluated by one of the parties; we let P_1 be the circuit evaluator.

Authenticated bits. Information-theoretic message authentication codes (IT-MACs) for authenticating bits were first used for efficient secure computation by Nielsen et al. [35] in the two-party setting. The idea can be extended to the multiparty setting as follows. Each player P_i holds a global MAC key $\Delta_i \in \{0, 1\}^K$. When P_i holds a bit x authenticated by P_j , this means that P_j is given a random key $K_j[x] \in \{0, 1\}^K$ and P_i is given the MAC tag $M_j[x] := K_j[x] \oplus x\Delta_j$. We let $[x]^i$ denote an authenticated bit where the value of x is known to P_i , and is authenticated to all other parties. In more detail, $[x]^i$ means that $(x, \{M_k[x]\}_{k \neq i})$ is given to P_i , and $K_j[x]$ is given to P_j for $j \neq i$.

Note that $[x]^i$ is XOR-homomorphic: given two authenticated bits $[x]^i, [y]^i$ known to the same party P_i , it is possible to locally compute the authenticated bit $[z]^i$ with $z = x \oplus y$ as follows:

- P_i computes $z := x \oplus y$, and $\{M_j[z] := M_j[x] \oplus M_j[y]\}_{j \neq i}$;
- P_j (for $j \neq i$) computes $K_j[z] := K_j[x] \oplus K_j[y]$.

Parties can also locally negate $[x]^i$, resulting in $[z]^i$ with $z = \bar{x}$:

- P_i computes $z := x \oplus 1$ and $\{M_j[z] := M_j[x]\}_{j \neq i}$;
- P_j (for $j \neq i$) computes $K_j[z] := K_j[x] \oplus \Delta_j$.

We let $\mathcal{F}_{\text{aBit}}^n$ denote an ideal functionality that distributes authenticated bits to the parties. (For details, see Figure 4 in Section 5.1). We also discuss an efficient instantiation of this functionality in Section 5.1.

Note that the above representation assumes that P_i uses a single global MAC key Δ_i . In cases where other keys are used, we explicitly add a subscript to the representation, i.e., we use $K_i[x]_{G_i}$ and $M_i[x]_{G_i} = K_i[x]_{G_i} \oplus xG_i$ to denote the key and MAC tag in this case.

Authenticated shares. In the above construction, x is known to one party. To generate an authenticated *shared* bit x , where x is not known to any party, we generate XOR-shares for x (i.e., shares $\{x^i\}$ with $\bigoplus_i x^i = x$) and then distribute the authenticated bits $\{[x^i]^i\}$. We let $\langle x \rangle$ denote the collection of these authenticated shares for x ; that is, $\langle x \rangle$ means that each party P_i holds $(x^i, \{M_j[x^i], K_i[x^j]\}_{j \neq i})$.

We let $\mathcal{F}_{\text{aShare}}$ denote an ideal functionality that distributes authenticated shared bits to the parties (Details see Figure 6 in Section 5.2). We discuss an efficient instantiation of this functionality in Section 5.2.

Definition of security. We use the standard notion of (standalone) security against an unbounded number of malicious parties. As our protocol is described, only P_A receives output; however, it would not be difficult to modify the protocol (using standard techniques) to support arbitrary outputs for different parties. In that case our protocol would achieve the notion of secure computation with designated abort [22, 24].

3 OVERVIEW OF OUR MAIN PROTOCOL

Our main protocol is designed in the \mathcal{F}_{Pre} -hybrid model (see Figure 1). At a high level, \mathcal{F}_{Pre} generates authenticated shares on random bits x, y, z such that $z = x \wedge y$. (We refer to these as *AND triples*.) Our main protocol then uses those authenticated shares to distributively construct a single, “authenticated” garbled circuit that is evaluated by one of the parties. In the remainder of this section, we describe our main protocol; further details are in Section 4. In Section 5 we then discuss how to efficiently realize \mathcal{F}_{Pre} .

3.1 Two-Party Authenticated Garbling

Wang et al. [37] recently proposed a maliciously secure protocol for two-party computation using a two-party version of \mathcal{F}_{Pre} . Before explaining how we extend it to the multiparty setting, we briefly describe their protocol partially based on their presentation.

Their protocol is based on a standard garbled circuit, where each wire α is associated with a random “mask” $\lambda_\alpha \in \{0, 1\}$ known to P_A . If the actual wire value when the circuit is evaluated is x , then the masked value observed by the circuit evaluator (namely, P_B) will be $\hat{x} = x \oplus \lambda_\alpha$. 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 . P_B also learns $L_{\alpha,\hat{x}}$ for that wire. Let H be a hash function modeled as a random oracle. The garbled table for an AND gate $(\alpha, \beta, \gamma, \wedge)$ is given by:

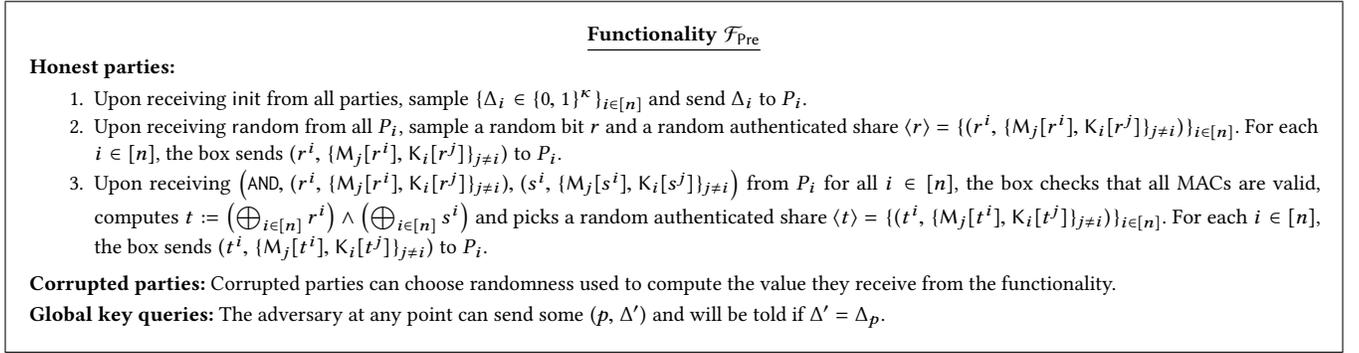


Figure 1: The multi-party preprocessing functionality.

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

P_B , holding $(\hat{x}, L_{\alpha, \hat{x}})$ and $(\hat{y}, L_{\beta, \hat{y}})$, evaluates it by picking the (\hat{x}, \hat{y}) -th row and decrypting using the garbled labels it holds, thus obtaining $(\hat{z}, L_{\gamma, \hat{z}})$. This is not secure because a malicious party can perform a selective failure attack to learn, e.g., \hat{x} . However, Wang et al. observe that it can be prevented if two parties hold *secret shares* of the garbled table: e.g., for the first row, P_A knows $(r_{00}, L_{\gamma, \hat{z}_{00}}^A)$, while P_B knows $(s_{00} = \hat{z}_{00} \oplus r_{00}, L_{\gamma, \hat{z}_{00}}^B)$, where $L_{\gamma, z} = L_{\gamma, z}^A \oplus L_{\gamma, z}^B$. Once P_A sends its shares of all the garbled gates, P_B can XOR those shares with its own and then evaluate the garbled circuit as before.

Informally, it ensures *privacy* against a malicious P_A since the results of any changes P_A makes to the garbled circuit are *independent* of P_B 's inputs. However, P_A can still affect *correctness* by, e.g., flipping the masked value in a row. This is solved using information-theoretic MACs on P_A 's share of masked bits. The shares of the garbled table now take the form as shown in the table below.

$\hat{x} \hat{y}$	P_A 's share of garbled table	P_B 's share of garbled table
0 0	$H(L_{\alpha,0}, L_{\beta,0}, \gamma, 00) \oplus (r_{00}, M[r_{00}], L_{\gamma, \hat{z}_{00}}^A)$	$(s_{00} = \hat{z}_{00} \oplus r_{00}, K[r_{00}], L_{\gamma, \hat{z}_{00}}^B)$
0 1	$H(L_{\alpha,0}, L_{\beta,1}, \gamma, 01) \oplus (r_{01}, M[r_{01}], L_{\gamma, \hat{z}_{01}}^A)$	$(s_{01} = \hat{z}_{01} \oplus r_{01}, K[r_{01}], L_{\gamma, \hat{z}_{01}}^B)$
1 0	$H(L_{\alpha,1}, L_{\beta,0}, \gamma, 10) \oplus (r_{10}, M[r_{10}], L_{\gamma, \hat{z}_{10}}^A)$	$(s_{10} = \hat{z}_{10} \oplus r_{10}, K[r_{10}], L_{\gamma, \hat{z}_{10}}^B)$
1 1	$H(L_{\alpha,1}, L_{\beta,1}, \gamma, 11) \oplus (r_{11}, M[r_{11}], L_{\gamma, \hat{z}_{11}}^A)$	$(s_{11} = \hat{z}_{11} \oplus r_{11}, K[r_{11}], L_{\gamma, \hat{z}_{11}}^B)$

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, which will occur independently of P_B 's actual input.

Finally they observe that if setting $\Delta = \Delta_A$ then $L_{\gamma, \hat{z}_{00}}$ can be efficiently secret-shared:

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

Recall that P_A knows $L_{\gamma, 0}$, r_{00} and Δ_A . Their key insight is that if s_{00} is an authenticated bit known to P_B , then P_A can locally compute the share $L_{\gamma, \hat{z}_{00}}^A := L_{\gamma, 0} \oplus r_{00} \Delta_A \oplus K[s_{00}]$ from the information it

has, and then the other share $L_{\gamma, \hat{z}_{00}}^B := K[s_{00}] \oplus s_{00} \Delta_A$ is just the MAC on s_{00} that P_B already holds. Table 3 shows the garbled table based on this observation.

3.2 Extension to the Multiparty Setting

It is not trivial to extend the above protocol to the multiparty setting. The main challenge is that even when $n - 1$ parties are corrupted, we still need to make sure that the adversary cannot learn any information about the honest party's inputs.

Attempted ideas. One idea, adopted by Choi et al. [14] in the three-party setting, is to let $n - 1$ parties jointly compute a garbled circuit that the remaining party will evaluate. However, if the $n - 1$ garblers are corrupt, there is no guarantee about the correctness of the garbled circuit they generate. For that reason, Choi et al. had to use cut-and-choose to check correctness of a random subset of ρ garbled circuits, which imposes a huge overhead.

To avoid this additional cut-and-choose, we would like all parties to be involved in the garbled-circuit generation, as in the BMR protocol [3]. However, state-of-the-art protocols based on BMR that are maliciously secure against corruption of $n - 1$ parties require either $O(n)$ somewhat homomorphic encryptions [32] or $O(n)$ SPDZ multiplication subprotocols [31] per AND gate both of which are relatively inefficient. We aim instead to use "simpler" TinyOT-like functionalities as we explain next.

Multiparty TinyOT: BDOZ-style vs. SPDZ-style. We observe that in the existing literature, there are mainly two flavors on how authenticated shares are constructed.

- **BDOZ-style** [7]: For a secret bit x , each party holds a share of x . For each ordered pair of parties (P_i, P_j) , P_i authenticates its own share (namely x^i) to P_j .
- **SPDZ-style** [19]: Each party holds a share of a global MAC key. For a secret bit x , each party holds a share of x and a share of the MAC on x .

Note that these protocols are constructed for arithmetic circuits, but these representations also apply to binary circuits. Existing papers prefer SPDZ-style shares to BDOZ-style shares, because SPDZ-style shares are smaller and thus more efficient to operate on. Indeed, existing papers that investigated protocols for multi-party TinyOT are all based on SPDZ-style shares [12, 20, 30].

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

Table 3: Final construction of an authenticated garbled table for an AND gate.

Our key observation is that such SPDZ-style AND triple, although efficient for interactive MPC protocols, are not suitable for our use to construct constant-round MPC protocols. In particular, in the SPDZ-style shares, each parties knows Δ_i as a share of the global key $\Delta = \bigoplus_i \Delta_i$. For each bit x , they holds shares of $x\Delta$. Since Δ is not known to any party, it is not directly related to any garbled circuit. On the contrary, in the BDOZ-style protocols, each party holds $(x^i, \{M_j[x^i], K_i[x^j]\}_{j \neq i})$, as we have already described in Section 2. In this case, they essentially hold shares of $x\Delta_i$ for all $i \in [n]$, because:

$$\begin{aligned}
x\Delta_i &= \left(\bigoplus_j x^j \right) \Delta_i = x^i \Delta_i \oplus \left(\bigoplus_{j \neq i} x^j \Delta_i \right) \\
&= x^i \Delta_i \oplus \left(\bigoplus_{j \neq i} M_i[x^j] \oplus K_i[x^j] \right) \\
&= \left(x^i \Delta_i \oplus \bigoplus_{j \neq i} K_i[x^j] \right) \oplus \bigoplus_{j \neq i} M_i[x^j]
\end{aligned}$$

Here, P_i knows the first value, while each P_j with $j \neq i$ knows $M_i[x^j]$. In other word, a BDOZ-style share of a bit x can be used to construct shares of $x\Delta_i$ for each i . This can further be used to construct shares of garbled labels, *if we use the same Δ_i for authenticated shares and the global difference used in free-XOR*. Indeed, looking ahead to the main protocol in Figure 2 step 4 (d), the content of the garbled circuit can be viewed as some authentication information plus shares of the garbled output labels for each garbler.

High level picture of the protocol. Given the above discussion, we can now picture the high level idea of our protocol. Our idea, from a high level view, is to let $n - 1$ parties be garblers, each maintaining a set of garbled labels, and to let the remaining party be the evaluator. Each garbler knows its own set of garbled labels. However, for each gate and for each garbler, the *permuted garbled output labels* are secretly shared to all parties, and thus no party knows how these labels are permuted. For each garbler P_i and a garbler row, P_i has shares of permuted garbled output labels for all garblers. In the garbled table, P_i encrypts all these shares using its own set of garbled input labels. Further, shares of the mask value are authenticated similarly. The evaluator decrypts the same row of garbled tables from all garblers in order to recompute the garbled output labels for each garbler. Intuitively, this ensures that for any set of $n - 1$ parties, they cannot garble or evaluate any gate.

4 THE MAIN SCHEME

Since we have discussed the main intuition of our protocol, we will proceed to the details directly. The proof of the main protocol can be found in Section A.1. In Figure 2 and Figure 3, we present the

complete MPC protocol in the \mathcal{F}_{Pre} -hybrid model. The protocol can be divided into five phases:

1. **Circuit Pre-scan.** (Step 1-3) In this phase, each party obtains their own private global MAC keys (Δ_i) from \mathcal{F}_{Pre} , and generate authenticated shares on wire masks for all wires.
2. **Circuit Garbling.** In this phase, each party compute shares of garbled tables for each garbler (Step 4 (a) - 4 (c)). Garblers then compute the distributed garbled circuits based on these shares (Step 4 (d)).
3. **Circuit Input Processing.** For each input wire that corresponds to P_i 's input, all other parties reveal their share of the wire mask to P_i . Party[i] then broadcasts the masked input values. All garblers, upon receiving this masked input value, send the corresponding key to the evaluator.
4. **Circuit Evaluation.** The evaluator evaluate the circuit following the topological order. In detail, the garbled wire labels from each garbler is used to obtain a set of shares of the wire labels for the output of the gate. The wire labels can then be constructed from the shares.
5. **Circuit Output Processing.** Now the evaluator holds masked output. All garblers reveal their shares of the output wire masks for the circuit to let evaluator unmask the value.

5 EFFICIENTLY REALIZING \mathcal{F}_{PRE}

In this section, we describe an efficient instantiation of \mathcal{F}_{Pre} , which is a multi-party version of TinyOT protocol. All previous related protocols [12, 20, 30] for multi-party TinyOT rely on cut-and-choose to ensure correctness and another bucketing to ensure privacy, resulting in a communication/computation complexity at least $\Omega(B^2 n^2)$ per AND triple, where bucket size $B = \rho / \log |C|$ (See Table 1 for more detail). Furthermore, these protocols output SPDZ-style shares that are not compatible with our main protocol. Our new protocol introduced in this section works with BDOZ-style shares; furthermore the complexity per AND triple is $O(Bn^2)$ with a very small constant. The new protocol features a new distributed AND triple checking protocol that checks the correctness of an AND triple *without* cut-and-choose. The adversary is still able to perform selective failure attacks on a triple with probability of being caught at least one-half. Such leakage can be easily eliminated using bucketing.

In the following, we will build our protocol from the bottom up. In Section 5.1 and Section 5.2, we discuss the multi-party authenticated bits and authenticated shares that we also introduced in Section 2; in Section 5.4, we discuss an AND triple generation protocol that allows an adversary to perform selective failure attacks; the final protocol that eliminates such an attack follows the bucketing protocols used in previous works [35, 37], and is detailed in Section 5.5.

Note that similar to the prior work [37], our protocol also relies on two-party $\mathcal{F}_{\text{aBit}}$ functionality, which has a random global key for

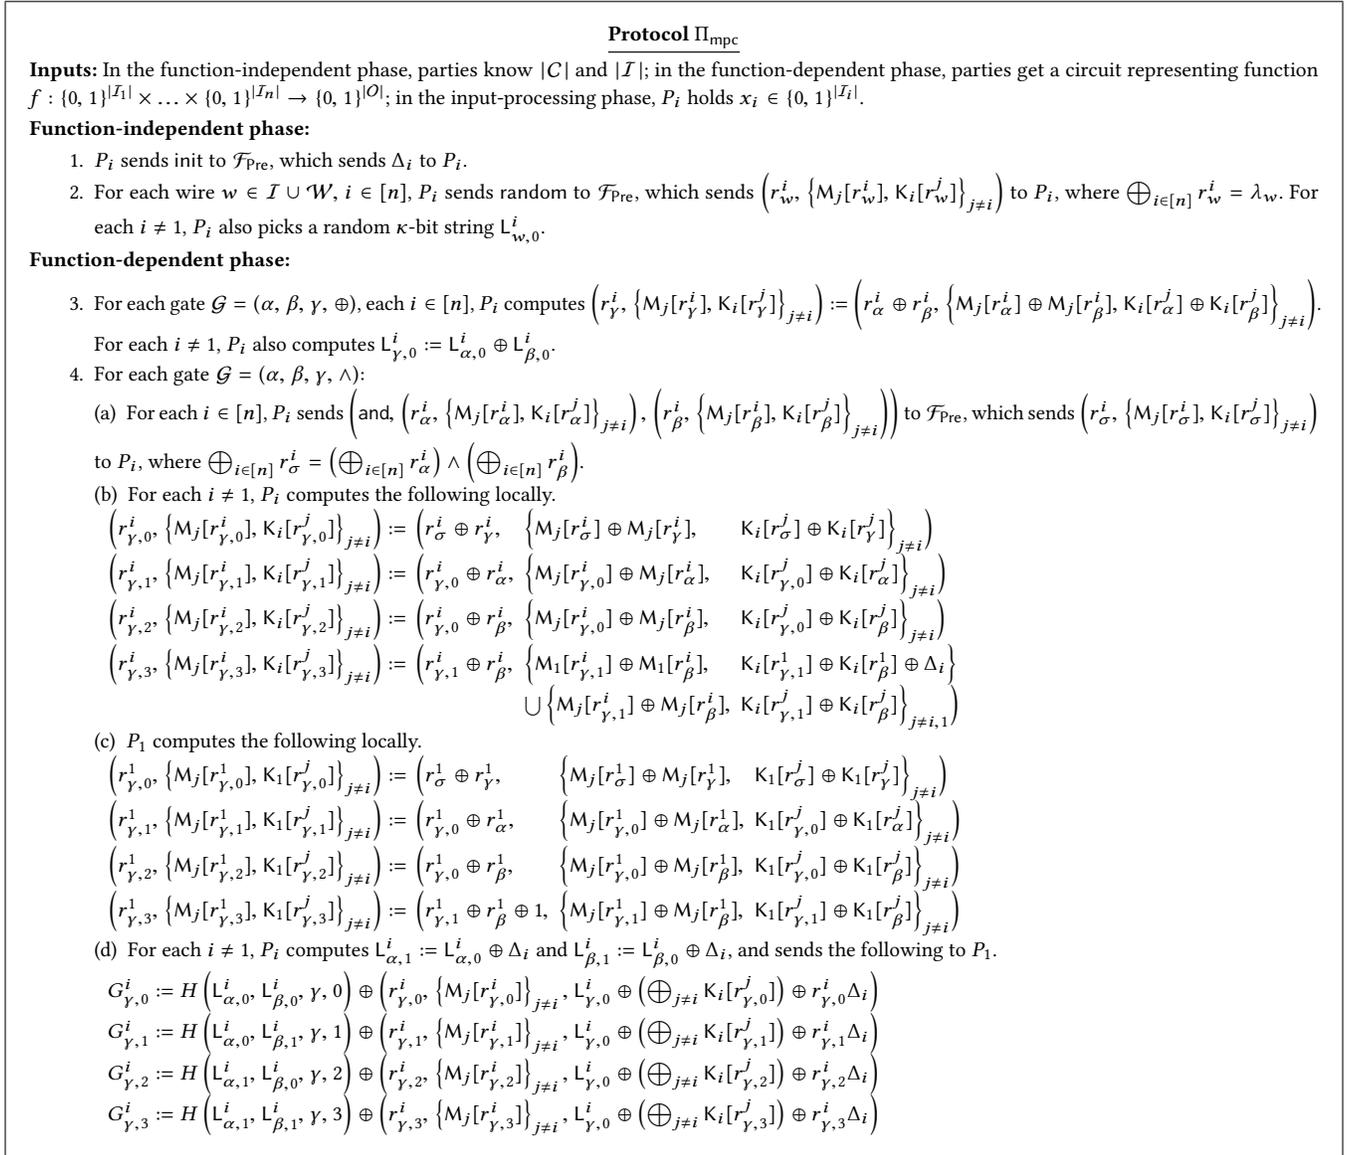


Figure 2: Our main protocol. Here ρ is set to κ for clarity, but this is not necessary.

each party preserved across all executions and allows an adversary to make “global-key queries” to honest parties’ global keys. Both these features are preserved in all our ideal functionalities, but we suppress explicit mention of them in our descriptions. Global-key queries have little effect on security, since the probability that the attacker can correctly guess the honest party’s value of Δ using polynomially many queries is negligible.

5.1 Multi-Party Authenticated Bit

The first step of our protocol is to generate multi-party authenticated bit. The functionality $\mathcal{F}_{\text{aBit}}^n$, also discussed in Section 2, is shown in Figure 4. Notice that if we set $n = 2$, then $\mathcal{F}_{\text{aBit}}^2$ is the original two-party authenticated bit functionality [35]. One naive

solution to realize $\mathcal{F}_{\text{aBit}}^n$ is to let P_i run the two-party authenticated bit protocol with every other party using the same bit x . This solution is not secure, since a malicious P_i can potentially use inconsistent values when running $\mathcal{F}_{\text{aBit}}^2$ with other parties. In our protocol, we use this general idea and we also perform additional checks to ensure that P_i uses consistent values. The check is similar to the recent malicious OT extension protocol by Keller et al. [27], where parties perform checks based on random linear combination: a malicious P_i who uses inconsistent values is able to pass ρ checks with probability at most $2^{-\rho}$. Note that these checks also reveal some linear relationship of x ’s. To eliminate this leakage, a small number of random authenticated bits are computed and checked together. They are later discarded to break the linear relationships. The protocol is described in Figure 5 with proof in Section A.2.

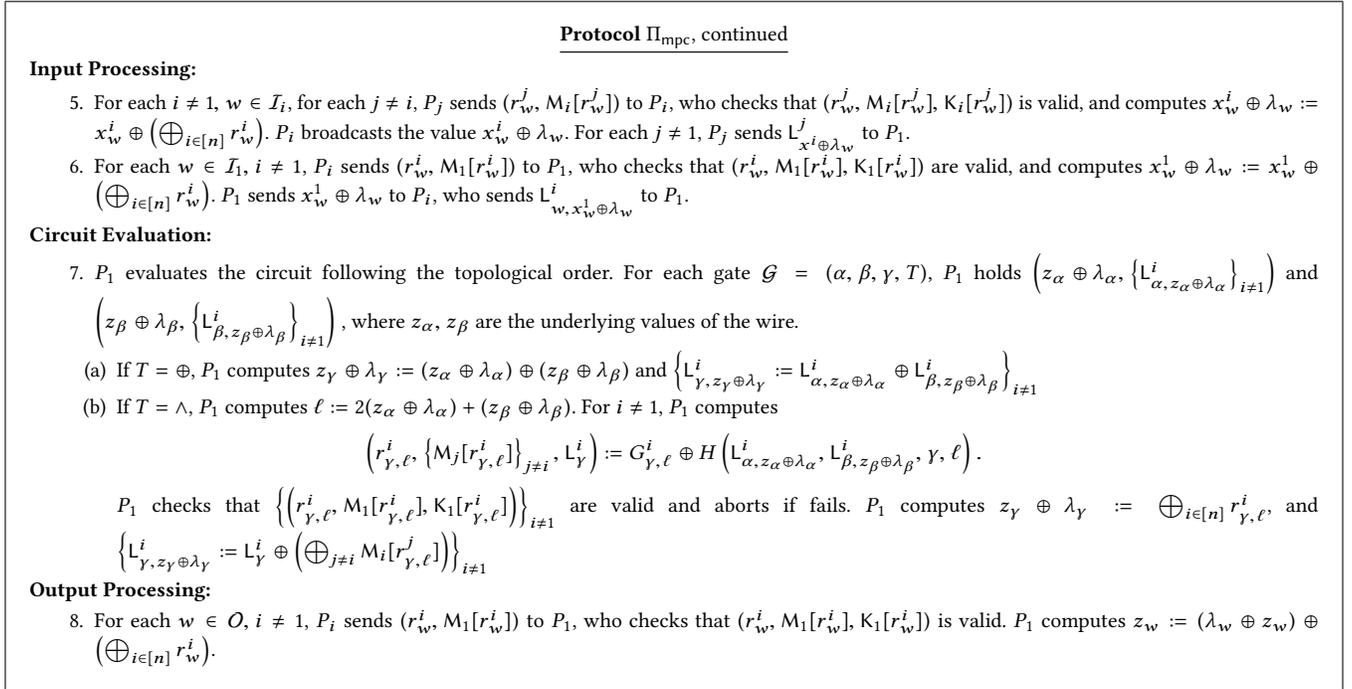


Figure 3: Our main protocol, continued. Here ρ is set to κ for clarity, but this is not necessary.

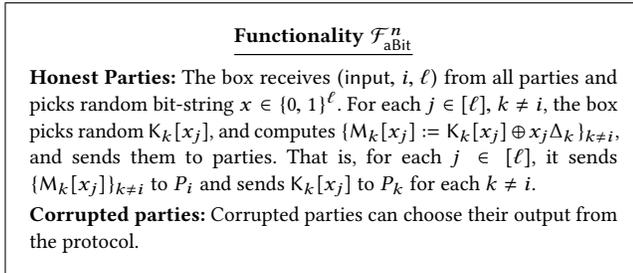


Figure 4: Functionality for multi-party authenticated bit.

5.2 Multi-Party Authenticated Shares

In this section, we aim to construct a protocol that allows multiple parties to obtain authenticated shares of a secret bit, as shown in Figure 6. One straightforward idea is to call $\mathcal{F}_{\text{aBit}}^n$ n times, where in the i -th execution, they compute $[x^i]^i$ for some random x^i known only to P_i . However, the adversary is still able to perform an attack: a malicious P_i can potentially use different global MAC keys (Δ_i) in different executions of $\mathcal{F}_{\text{aBit}}^n$. The result is that $[x^j]^j$ is authenticated with a global MAC key Δ_i , while some other $[x^k]^k$ is authenticated with a different global MAC key Δ'_i . This attack does not happen in the two-party setting, because each party is authenticated to only one party.

Our key idea is based on the observation that the two-party authenticated bit protocol already ensured that, when P_i and P_j compute multiple authenticated bits, P_i uses the same Δ_i across different authenticated bits. Therefore, in the above insecure attempt, if one authenticated share has consistent global MAC keys,

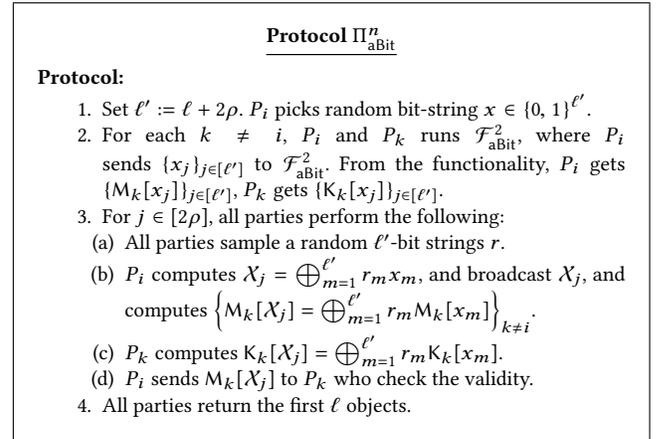


Figure 5: The protocol Π_{aBit}^n instantiating $\mathcal{F}_{\text{aBit}}^n$.

then all authenticated shares have consistent global MAC keys, and vice versa. In our secure construction, we first let all parties compute $\ell + \rho$ number of multi-party authenticated shares as described above, which may not be secure. Then we partially open the last ρ tuples to check the consistency of global MAC keys. A malicious party who uses inconsistent Δ_i 's will get caught with probability one-half for each partially opened shares.

In more detail, each player P_i will take the role of a prover once to prove that he uses a consistent Δ_i and the remaining players will take the role of verifier for the given prover. The basic idea is that if the prover used a consistent Δ_i , then these authenticated bits across different parties are XOR homomorphic. Taking a three-party setting as an example. Say P_1 has $K_1[x]$, $K_2[y]$ with global

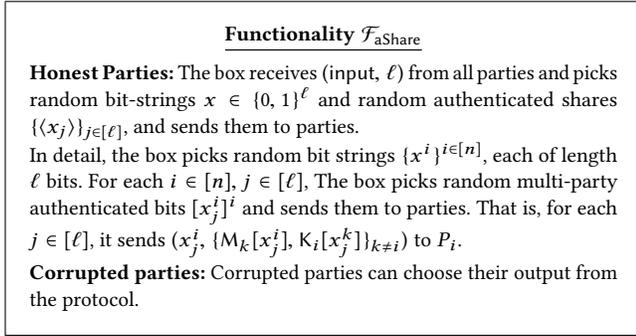
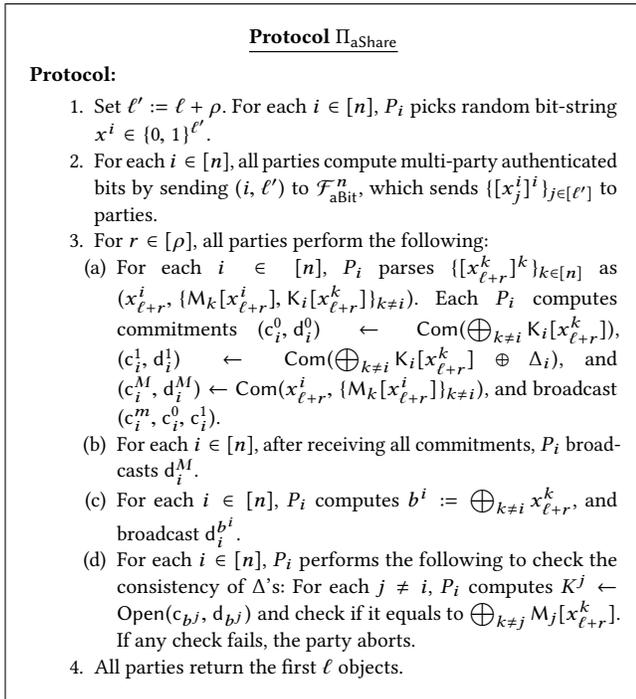


Figure 6: Functionality for multi-party authenticated share.

Figure 7: The protocol Π_{aShare} instantiating $\mathcal{F}_{\text{aShare}}$.

keys Δ_1^x and Δ_1^y which are potentially different; P_2 has $(x, M_1[x])$; P_3 has $(y, M_1[y])$. In our checking protocol, we let P_1 commit to values $K_1[x] \oplus K_1[y]$ and $K_1[x] \oplus K_1[y] \oplus \Delta_1$. For an adversary who uses inconsistent global keys, it needs to choose two values out of the following four values to commit.

$x = 0$	$y = 0$	$K_1[x] \oplus K_2[y]$
$x = 0$	$y = 1$	$K_1[x] \oplus K_2[y] \oplus \Delta_1^y$
$x = 1$	$y = 0$	$K_1[x] \oplus K_2[y] \oplus \Delta_1^x$
$x = 1$	$y = 1$	$K_1[x] \oplus K_2[y] \oplus \Delta_1^x \oplus \Delta_1^y$

Later in the protocol, the adversary is asked to open the MAC for $x \oplus y$. If inconsistent global keys are used, this value can be any of the four values each with one-fourth probability, therefore, the adversary can win with probability at most $2/4 = 1/2$. The details of the protocol are shown in Figure 7 with proof in Section A.3.

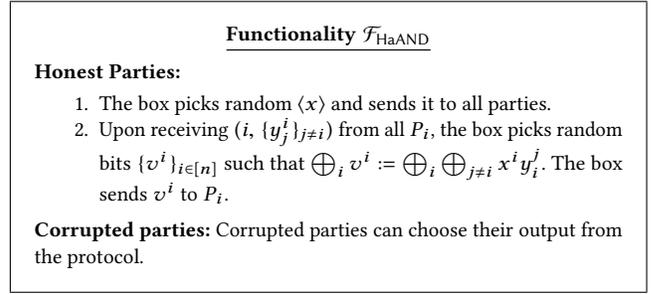
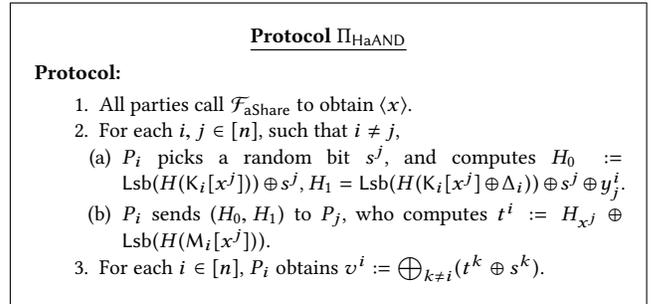


Figure 8: The Half Authenticated AND Functionality

Figure 9: Protocol Π_{HaAND} instantiating $\mathcal{F}_{\text{HaAND}}$.

5.3 Half-Authenticated AND Triple

Before introducing the protocol for leaky authenticated AND triples, there is yet another tool that we need. As described in Figure 8, the functionality $\mathcal{F}_{\text{HaAND}}$ is introduced to compute cross terms in AND triples. It takes unauthenticated and potentially inconsistent y 's and outputs authenticated share $\langle x \rangle$ as well as unauthenticated shares of cross product terms. Details about the protocol is in Figure 9 with proof included in Section A.3.

5.4 Multi-Party Leaky Authenticated AND Triple

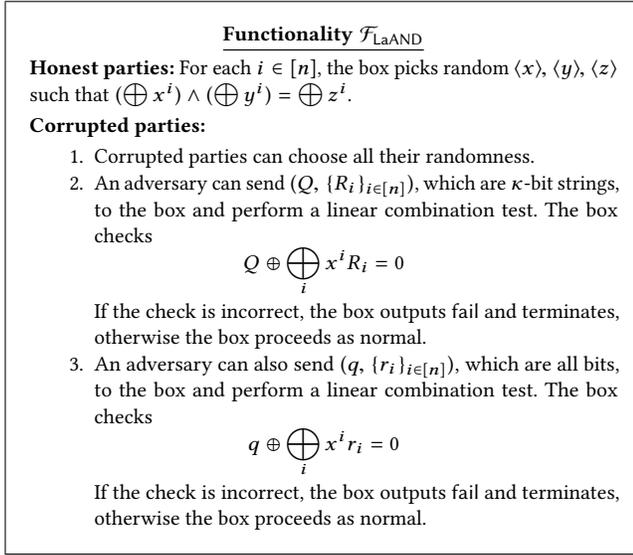
Now we are ready to discuss the protocol for leaky authenticated AND triples. It can be divided into following steps:

1. Call $\mathcal{F}_{\text{aShare}}$ to obtain some random $\langle y \rangle$ and $\langle r \rangle$.
2. Call $\mathcal{F}_{\text{HaAND}}$ with y to obtain a random $\langle x \rangle$ and compute shares $\{z^i\}$, such that $(\bigoplus_i x^i) \wedge (\bigoplus_i y^i) = \bigoplus_i z^i$.
3. Reveal $d = z \oplus r$ and computes $\langle z \rangle := \langle r \rangle \oplus d$.
4. Perform additional check to ensure the correctness of the AND relationship.

In the above steps, the adversary is able to cheat by using inconsistent values of y and z between step 1 and 2. However, this only allows the adversary to perform selective failure attack on x^i 's. For example, the AND relationship checked is

$$\left(\bigoplus_i x^i \right) \wedge \left(\bigoplus_i y^i \right) = \left(\bigoplus_i z^i \right).$$

The adversary can guess that the value of $\bigoplus_i x^i = 0$ and flip y^j for some $j \in \mathcal{M}$. If the guess is correct, than the check will go through

Figure 10: Functionality $\mathcal{F}_{\text{LaAND}}$ for leaky AND triple.

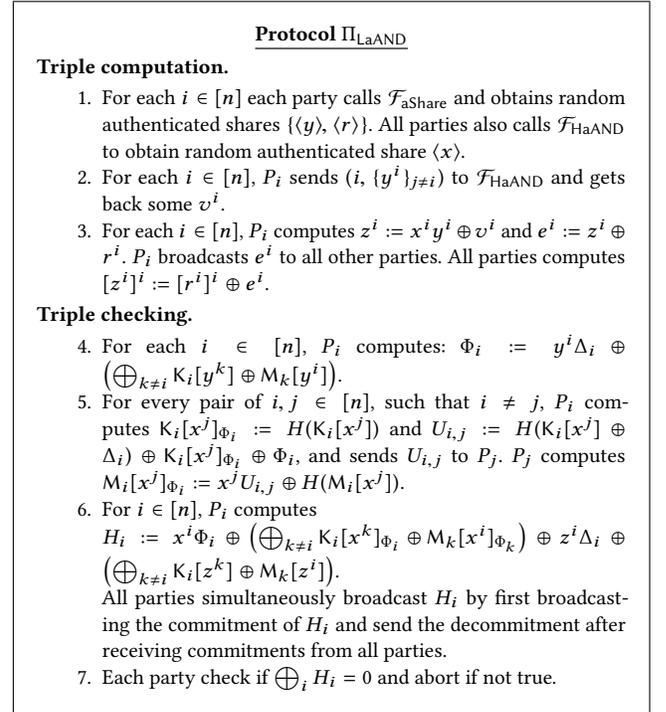
and the protocol will proceed as normal. However, if the guess is wrong, then the checking will abort and the adversary is caught.

One main challenge is what leakage we should aim for in this functionality. We would like to limit the leakage to be possible only on x^i 's, otherwise we would need more bucketing for each possible leakage, as also noted by Nielsen et al. [35]. On the other hand, the adversary can do more attacks than the one mentioned above: it is also possible to, for example, learn $\bigoplus_{i \in S} x^i$ for some set $S \subset [n]$. We find that the best way to abstract such attack is to allow the adversary to perform a linear check on the value of x^i 's. As shown in Figure 10, the adversary is allowed to send a list of coefficients and check if the inner product between the coefficients and x values is zero or not.

Our checking phase differs substantially from existing works. We design an efficient checking protocol, that ensures the correctness of the triple (if no party aborts) which allows malicious parties to learn k bits of some specific information with probability at most 2^{-k} . In the two-party protocol, one party constructs "checking tables" and lets the other party to evaluate/check. In the multi-party protocol here, we instead let all parties distributively construct the "checking tables". Interestingly, distributively constructing these checks is inspired by the main protocol where parties distributively construct garbled tables. As noted before, this protocol is vulnerable to selective failure attacks. The full description of this protocol is presented in Figure 11.

In the following, we will show the correctness and unforgeability of the protocol, which are crucial to the security proof of the protocol.

5.4.1 Correctness of the protocol. We want to show that the protocol will compute a correct triple and will not abort if all parties are honest. Notice that the value we are checking can be written as:

Figure 11: The protocol Π_{LaAND} .

$$\begin{aligned}
 & \bigoplus_i H_i \\
 &= \bigoplus_i \left(x^i \Phi_i \oplus \left(\bigoplus_{k \neq i} K_i[x^k]_{\Phi_i} \oplus M_k[x^i]_{\Phi_k} \right) \oplus z^i \Delta_i \oplus \left(\bigoplus_{k \neq i} K_i[z^k] \oplus M_k[z^i] \right) \right) \\
 &= \bigoplus_i \left(x^i \Phi_i \oplus \left(\bigoplus_{k \neq i} K_i[x^k]_{\Phi_i} \oplus M_k[x^i]_{\Phi_k} \right) \right) \oplus \bigoplus_i \left(z^i \Delta_i \oplus \left(\bigoplus_{k \neq i} K_i[z^k] \oplus M_k[z^i] \right) \right) \\
 &= \bigoplus_i \left(x^i \Phi_i \oplus \left(\bigoplus_{k \neq i} K_i[x^k]_{\Phi_i} \oplus M_i[x^k]_{\Phi_k} \right) \right) \oplus \bigoplus_i \left(z^i \Delta_i \oplus \left(\bigoplus_{k \neq i} K_i[z^k] \oplus M_i[z^k] \right) \right) \\
 &= \bigoplus_i \left(x^i \Phi_i \oplus \left(\bigoplus_{k \neq i} x^k \Phi_i \right) \right) \oplus \bigoplus_i \left(z^i \Delta_i \oplus \left(\bigoplus_{k \neq i} z^k \Delta_i \right) \right) \\
 &= \left(\bigoplus_i x^i \right) \cdot \left(\bigoplus_i \Phi_i \right) \oplus \left(\bigoplus_i z^i \right) \cdot \left(\bigoplus_i \Delta_i \right)
 \end{aligned}$$

Notice further that

$$\bigoplus_i \Phi_i = \bigoplus_i \left(y^i \Delta_i \oplus \left(\bigoplus_{k \neq i} K_i[y^k] \oplus M_k[y^i] \right) \right) = \left(\bigoplus_i y^i \right) \cdot \left(\bigoplus_i \Delta_i \right)$$

Therefore we know that

$$\begin{aligned}
 \bigoplus_i H_i &= \left(\bigoplus_i x^i \right) \cdot \left(\bigoplus_i \Phi_i \right) \oplus \left(\bigoplus_i z^i \right) \cdot \left(\bigoplus_i \Delta_i \right) \\
 &= \left(\bigoplus_i x^i \right) \cdot \left(\bigoplus_i y^i \right) \cdot \left(\bigoplus_i \Delta_i \right) \oplus \left(\bigoplus_i z^i \right) \cdot \left(\bigoplus_i \Delta_i \right) \\
 &= \left(\left(\bigoplus_i x^i \right) \cdot \left(\bigoplus_i y^i \right) \oplus \left(\bigoplus_i z^i \right) \right) \cdot \left(\bigoplus_i \Delta_i \right)
 \end{aligned}$$

Since $\bigoplus_i \Delta_i$ is non-zero, $\bigoplus_i H_i = 0$ if and only if the logic of this AND is correct.

5.4.2 *Unforgeability.* Now we want to show that any incorrect AND triple cannot pass the check.

LEMMA 5.1. Define x^i, y^i from $\langle x \rangle, \langle y \rangle$ which are outputs from $\mathcal{F}_{\text{aShare}}$ and $\mathcal{F}_{\text{HaAND}}$; define $z^i := r^i \oplus e^i$, where $\langle r \rangle$ is output from $\mathcal{F}_{\text{aShare}}$, e^i is the value broadcast from P_i . If $(\bigoplus_i x^i) \wedge (\bigoplus_i y^i) \neq (\bigoplus_i z^i)$ then the protocol results in an abort except with negligible probability.

We use $U_{i,j}^*$ and H_i^* to denote the values that an honest party would have computed, and define $Q_{i,j} = U_{i,j}^* \oplus U_{i,j}$, $Q_i = H_i^* \oplus H_i$. In the following, we will assume that the logic of the AND does not hold while at the same time that the check passes, and we will derive a contradiction from it.

First note that if P_i uses some $Q_{i,j}$, then P_j will obtain $M_i[x^j]_{\Phi_i}$ with an additive error of $x^j Q_{i,j}$. Note that

$$\bigoplus_i H_i^* = \left(\left(\bigoplus_i x^i \right) \cdot \left(\bigoplus_i y^i \right) \oplus \left(\bigoplus_i z^i \right) \right) \cdot \left(\bigoplus_i \Delta_i \right) = \bigoplus_i \Delta_i$$

Therefore, we know that

$$\begin{aligned} \bigoplus_i H_i &= \bigoplus_{i \in \mathcal{M}} H_i \oplus \bigoplus_{i \in \mathcal{H}} H_i \\ &= \bigoplus_{i \in \mathcal{M}} (H_i^* \oplus Q_i) \oplus \bigoplus_{i \in \mathcal{H}} \left(H_i^* \oplus \left(\bigoplus_{k \neq i} x^k Q_{k,i} \right) \right) \\ &= \bigoplus_i H_i^* \oplus \bigoplus_{i \in \mathcal{M}} Q_i \oplus \bigoplus_{i \in \mathcal{H}} \left(\bigoplus_{k \neq i} x^k Q_{k,i} \right) \\ &= \bigoplus_i \Delta_i \oplus \bigoplus_{i \in \mathcal{M}} Q_i \oplus \bigoplus_{i \in \mathcal{H}} \left(\bigoplus_{k \neq i} x^k Q_{k,i} \right) \end{aligned}$$

In order to make $\bigoplus_i H_i$ to be 0, the adversary needs to find paddings such that

$$\bigoplus_{i \in \mathcal{M}} Q_i \oplus \bigoplus_{i \in \mathcal{H}} \left(\bigoplus_{k \neq i} x^k Q_{k,i} \right) = \bigoplus_i \Delta_i$$

The above happens with at most negligible probability.

THEOREM 5.2. Assuming an adversary corrupting up to $n - 1$ parties, the protocol in Figure 11, where H is modeled as a random oracle, securely instantiates $\mathcal{F}_{\text{LaAND}}$ functionality in the $(\mathcal{F}_{\text{aShare}}, \mathcal{F}_{\text{HaAND}})$ -hybrid model.

Note that since no party has private input, the simulation proof is straightforward given the lemmas above. We provide full details of the proof in Section A.5.

5.5 Multi-Party Authenticated AND Triple

Once we have a protocol for leaky authenticated AND triple, it is straightforward to obtain a non-leaky authenticated AND triple, using the combine protocol in [37]. We show the details of the protocol in Figure 13.

6 EVALUATION

6.1 Implementation Details

We implemented our protocol in the EMP-toolkit [36] framework and will be made publicly available as a part of it. To fully explore performance characteristics of our protocol, we evaluate our implementation in three different settings:

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

Honest parties: For each $i \in [n]$, the box picks random $\langle x \rangle, \langle y \rangle, \langle z \rangle$ such that $(\bigoplus x^i) \wedge (\bigoplus y^i) = \bigoplus z^i$.

Corrupted parties: Corrupted parties get to choose all of their randomness.

Figure 12: Functionality $\mathcal{F}_{\text{aAND}}$ for generating AND triples

Protocol Π_{aAND}

Protocol:

1. P_i call $\mathcal{F}_{\text{LaAND}} \ell' = \ell B$ times and obtains $\{\langle x_j \rangle, \langle y_j \rangle, \langle z_j \rangle\}_{j \in [\ell']}$.
2. All parties randomly partition all objects into ℓ buckets, each with B objects.
3. For each bucket, parties combine B leaky ANDs into one non-leaky AND. To combine two leaky ANDs, namely $(\langle x_1 \rangle, \langle y_1 \rangle, \langle z_1 \rangle)$ and $(\langle x_2 \rangle, \langle y_2 \rangle, \langle z_2 \rangle)$:
 - (a) Parties reveal $d := y_1 \oplus y_2$ with its MACs checked.
 - (b) Each party P_i sets $\langle x \rangle := \langle x_1 \rangle \oplus \langle x_2 \rangle$, $\langle y \rangle := \langle y_1 \rangle$, $\langle z \rangle := \langle z_1 \rangle \oplus \langle z_2 \rangle \oplus d \langle x_2 \rangle$.
 Parties iterate all B leaky objects, by taking the resulted object and combine with the next element.

Figure 13: Protocol Π_{aAND} instantiating $\mathcal{F}_{\text{aAND}}$.



Figure 14: Amazon EC2 regions used in the WAN experiment. Details see Table 9.

Circuit	n_1	n_2	n_3	$ C $
AES	128	128	128	6800
SHA-128	256	256	160	37300
SHA-256	256	256	256	90825

Table 4: Circuits used in our evaluation.

- **LAN setting.** Machines are located in the same Amazon EC2 region. Experiments are performed for up to 14 parties.
- **WAN setting.** Each Machine is located in a *different* Amazon EC2 region (locations shown in Figure 14). For a k -party computation experiment, a prefix subset of the machines in Table 9 are selected. For example, parties in 3PC experiments are located in North Virginia, Ohio, and North California

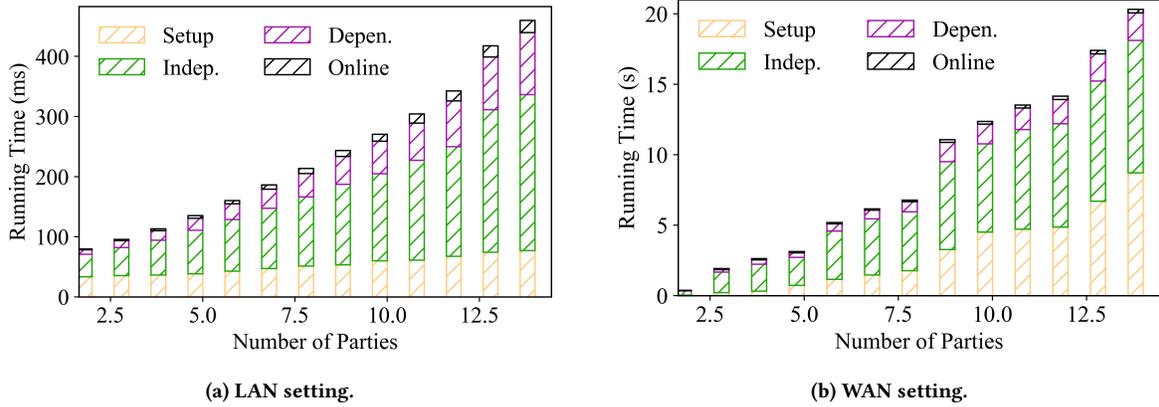


Figure 15: Running time breakdown for evaluating AES. In the LAN setting, all parties are located in the same region; In the WAN setting, all parties are located in different regions worldwide, for example, **2PC**: within US-east; **5PC**: within North America; **8PC**: within North America and Europe; **12PC**: within North America, Europe and Asia; **13PC**: further adds Sydney; **14PC**: all parties in Figure 14.

respectively. Experiments are run for up to 14 parties, which is the number of different Amazon EC2 regions available.

- **Crowd setting.** In the LAN case, we evaluate up to 128 parties all located in the same Amazon EC2 region. In the global-scale case, we choose 8 different cities across 5 continents and open up to 16 parties in each city (totally 128 parties).

All machines are of type c4.8xlarge, with 36 cores and 60 GB RAM. Network bandwidth within the same region is about 10Gbps. The bandwidth across different regions depends on the location of the machines. All experiments are based on $\rho = 40, \kappa = 128$. We extend justGarble [4] to support garbling of longer tables in a straightforward manner. In the implementation, we used the “broadcast with abort” protocol by Goldwasser and Lindell [24] and achieves the notion of secure computation with abort. We observe small variance when running in the LAN setting and slightly higher variance in the WAN setting. All numbers reported in LAN setting are on average of 10 runs, ones in WAN setting are on average of 20 runs, and ones in the Crowd setting are on average of 5 runs due to the lengthy experiments.

6.2 Performance on Basic Circuits

We evaluate commonly used benchmark circuits on our protocol, including AES, SHA-1, and SHA-256. Information about these circuits can be found in Table 4.

We plot in Figure 15 the results for AES in different network setting and performance breakdown as described above. Detailed timings and more results for all three circuits can be found in Table 10 in the Appendix. First, the performance of three-party computation is extremely efficient: it takes 95 ms to evaluate a circuit for AES, with 2 ms online time. We also find that in the LAN setting, the slowdown from 2PC to 3PC is roughly 1.5 \times ; the slowdown in the WAN setting is larger. This is caused by the network latency: the first two parties are both located in the U.S. east coast, while the third party is located in the U.S. west coast with much higher latency.

We also find that the cost of the one-time setup is almost independent of number of parties for small number of parties. This is mainly due to the parallelization in the implementation that allows all base-OT to run at the same time.

World-wide MPC experiment. We would like to emphasize that in the case of WAN setting with 14 parties, it is a “world-wide” MPC experiment over 5 continents. To the best of our knowledge, we are the first to conduct MPC over such large range even considering semi-honest MPC protocols.

We also notice a big “jump” in running time from 8 parties to 9 parties in the WAN setting. We believe this is because of the network condition: for experiments up to 8 parties, it is within the US/Europe area; the ninth party is located in asia, where the communication to US/Europe is much slower. More details see Figure 15b.

6.3 Evaluation in a Global Setting

In this section, we focus on the performance of our protocol with a large number of parties. We summarize our results in Figure 5. We notice that our protocol scales very well with increasing number of parties. Even in a setting with 128 parties located in the same LAN, where up to 127 of them can be corrupted, it takes less than 17 seconds end-to-end running time to compute AES. Note that the performance of a 64-party computation on AES is comparable to the performance of what used to be the state-of-the-art malicious 2-party computation three years ago [1], and we believe further optimizations and improvements based on our work will flourish too.

When comparing the running time of 128 parties to the one of 8 parties, we find that the cost of function-dependent phase increases much faster than the cost of function-independent phase. This is because our function-independent phase is symmetric and all communication loads are evenly distributed among all parties; while in the function-dependent phase, $n - 1$ garblers send the garbled circuit to the evaluator, and the bandwidth of the evaluator becomes the bottleneck. Therefore in the case where there are a lot

n	LAN setting					WAN setting				
	Setup	Indep.	Depen.	Online	Total	Setup	Indep.	Depen.	Online	Total
8	49.4	122.2	35.7	7.8	215.1	16736.0	30647.4	5905.2	783.3	54071.9
16	78.8	227.8	121.1	29.4	457.0	18708.1	21699.6	12243.5	598.8	53250.0
32	129.9	627.0	446.2	112.3	1315.5	35838.8	19038.4	8242.3	716.6	63836.1
64	212.9	1182.2	2630.1	476.5	4501.7	71913.8	25280.7	29416.6	1564.0	128175.2
128	383.0	2727.4	11669.6	1870.2	16650.2	88055.9	30795.9	22659.1	2316.2	143827.0

Table 5: Detailed experiment results for the crowd setting. Timings are measured in terms of milliseconds. In the LAN setting, all parties are located in the same region. In the Worldwide setting, 8PC is performed with each party located in a different region; 16PC is performed with 2 parties located in each region; others can be interpreted similarly.

of parties, when we double the number of parties, the running time of the function-dependent phase almost doubles.

We also run the same experiment in the worldwide range. We choose 8 most separate regions out of 14 and open up to 16 machines in each region (thus totally 128 machines). The performance is also shown in Table 5: it takes slightly more than a minute for 64 parties to compute AES and about 2.5 minutes for 128 parties located all around the world. We also observe that setup takes much more time; we believe it is due to the high latency.

6.4 Comparison to Other Work

Malicious MPC on AES. Evaluating AES with malicious security against $n - 1$ corruption was studied by Damgård et al. [17]. They reported 240 ms *online time* for 3 parties and 340 ms online time for 10 parties. The offline time for 3 and 10 parties are around 4200 seconds and 15000 seconds respectively. Our protocol takes 95 ms total time to evaluate AES for 3 parties with online time as small as 2 ms; and 268 ms total time with online time 12 ms. The improvement for online phase ranges from $28\times$ to $120\times$; and the improvement for total time ranges from $44000\times$ to $56000\times$. This is a huge improvement even considering hardware differences.

BMR-style protocols. Lindell et al. [31, 32] studied how to use SPDZ and SHE to construct a BMR-style protocol. Since their protocol is not implemented, we compare the communication complexity. After incorporating various optimizations, every AND gate still need $3n + 1$ SPDZ multiplication triples. Together with the most recent advance in SPDZ triple generation by Keller et al. [28], generating one SPDZ triple with n parties requires communication about $180(n - 1)$ kilobits per party. Therefore the communication cost per AND gate per party is about $540n(n - 1)$ kilobits. In our protocol, each AND gate only needs one AND triple from \mathcal{F}_{Pre} , which, using our new protocol in Section 5, requires communication roughly $2.28(n - 1)$ kilobits per party. Therefore, the improvement of our protocol compared to the best-optimized BMR protocol based on Lindell et al. is about $237n\times$ with n parties. For a three-party setting, it is an improvement of $711\times$; for the 128-party computation that we perform, the improvement is as high as $30,000\times$!

Ben-Efraim et al. [6] presented a protocol secure in the *semi-honest* model based on BMR. Surprisingly, given the fact that our protocol is maliciously secure, while theirs only has a semi-honest security, our implementation has roughly the same performance as theirs. In Table 6, we compare the running time of our protocol with the running time of theirs based on the same hardware. We

		3	8	16	32
[6]	Online	80	150	400	1500
	Total	228	2000	5900	-
This work	Online	24	95	370	1632
	Total	618	1945	6711	18828

Table 6: Compare with Ben-Efraim et al. [6] Timings are in terms of milliseconds. Their protocol works in the semi-honest setting; ours is maliciously secure. Comparison based on SHA-256, with the same hardware configuration.

		ρ	3	8	16
[25]	40	14	49	105	
	80	55	193	413	
This work	40	4.8	16.9	36.4	
	80	8.6	30	64.5	

Table 7: Compare bandwidth consumption with Hazay et al. [25]. All numbers are the maximum amount of data one party needs to send in the function-independent phase, measured in terms of megabytes (MB). Numbers for $\rho = 80$ are calculated based on the complexity of both protocols.

notice that for both online time and total time, the performance of the two protocols are roughly the same.

Malicious 3PC protocols with honest majority. Mohassel et al. [33] proposed an efficient protocol for malicious 3PC with honest majority. Their protocol requires only one garbled circuit to be sent and therefore has a smaller communication complexity than us. We estimate that our protocol requires about $14\times$ more communication than theirs. However interestingly we also find that the online time of two protocols are roughly the same: their protocol requires 31 ms evaluation time, while ours needs 23.4 ms evaluation time. We believe this is due to the fact that their protocol needs to check that the garbled circuit received from two garblers are the same, while it is not needed in our protocol.

Furukawa et al. [21] also presented a malicious 3PC protocol with honest majority. Their protocol has a smaller communication overhead compared to the protocol above by Mohassel et al. but requires at least one round of communication per level of the circuit. In addition to a stronger security guarantee that we support dishonest majority, our protocol has a better latency, especially for deep circuit (e.g. SHA-256 has a depth of 4000), while their protocol has a better throughput.

n	Setup	Indep.	Depen.	Online	Total
AES					
3	57.1 KB	4.8 MB	1.3 MB	4.5 KB	6.2 MB
4	85.7 KB	7.2 MB	1.8 MB	4.5 KB	9.1 MB
5	114.2 KB	9.7 MB	2.2 MB	4.5 KB	12.0 MB
6	142.8 KB	12.1 MB	2.7 MB	4.5 KB	14.9 MB
7	171.4 KB	14.5 MB	3.1 MB	4.5 KB	17.8 MB
8	199.9 KB	16.9 MB	3.5 MB	4.5 KB	20.7 MB
16	428.4 KB	36.4 MB	7.1 MB	4.5 KB	44.0 MB
SHA-256					
3	57.1 KB	63.3 MB	17.4 MB	9.0 KB	80.8 MB
4	85.7 KB	95.0 MB	23.4 MB	9.0 KB	118.5 MB
5	114.2 KB	126.6 MB	29.4 MB	9.0 KB	156.2 MB
6	142.8 KB	158.3 MB	35.4 MB	9.0 KB	193.9 MB
7	171.4 KB	190.0 MB	41.4 MB	9.0 KB	231.6 MB
8	199.9 KB	221.7 MB	47.4 MB	9.0 KB	269.3 MB
16	428.4 KB	475.1 MB	95.4 MB	9.0 KB	570.9 MB

Table 8: Communication complexity of our protocol. Bandwidth are measured for evaluating AES and SHA-256. All numbers are the maximum amount of data one party needs to send.

Compare with Hazay et al. [25]. We also compare with the concurrent work by Hazay et al.. As their protocol is benchmarked on different hardware and network configurations, we only compare the bandwidth usage. We find that both protocols has similar function-dependent cost and online cost. However, due to our improved preprocessing protocol, our function-independent cost is much smaller than theirs. In Table 7, we compare the function-independent cost of for AES evaluation with different value of ρ . Our protocol uses $3\times$ to $6.5\times$ less communication compared to theirs. Note that the cost of function-independent phase dominates the overall cost, therefore the speed up here also translates to the speed up to the whole computation.

6.5 Communication Complexity

In this section, we evaluate the communication complexity of our protocol. All numbers reported here are the maximum amount of data sent from one party. All numbers are obtained by running our implementation, which are slightly higher than calculated values due to implementation details. In Table 8, we summarize the bandwidth use for AES and SHA-256 for up to 16 parties.

We can see from the figure that the communication required per party grows linearly with the number of parties. In addition, the communication cost of the Setup phase and the Online phase are very small. The total communication cost is dominated by the function-independent phase.

ACKNOWLEDGMENTS

The authors would like to thank Roberto Trifiletti, Yan Huang, and Ruiyu Zhu for their helpful comments.

This material is based on work supported by NSF awards #1111599, #1563722, and #1564088. 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 author(s) and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

REFERENCES

- [1] Arash Afshar, Payman Mohassel, Benny Pinkas, and Ben Riva. 2014. Non-Interactive Secure Computation Based on Cut-and-Choose. In *Eurocrypt 2014 (LNCS)*, Vol. 8441. 387–404.
- [2] Toshinori Araki, Jun Furukawa, Yehuda Lindell, Ariel Nof, and Kazuma Ohara. 2016. High-Throughput Semi-Honest Secure Three-Party Computation with an Honest Majority. In *ACM CCS 2016*. 805–817.
- [3] Donald Beaver, Silvio Micali, and Phillip Rogaway. 1990. The round complexity of secure protocols. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing*. ACM, 503–513.
- [4] Mihir Bellare, Viet Tung Hoang, Sriram Keelveedhi, and Phillip Rogaway. 2013. Efficient Garbling from a Fixed-Key Blockcipher. In *IEEE Symposium on Security & Privacy*. 478–492.
- [5] Assaf Ben-David, Noam Nisan, and Benny Pinkas. 2008. FairplayMP: a system for secure multi-party computation. In *ACM CCS 2008*. 257–266.
- [6] Aner Ben-Efraim, Yehuda Lindell, and Eran Omri. 2016. Optimizing Semi-Honest Secure Multiparty Computation for the Internet. In *ACM CCS 2016*. 578–590.
- [7] Rikke Bendlin, Ivan Damgård, Claudio Orlandi, and Sarah Zakarias. 2011. Semi-homomorphic Encryption and Multiparty Computation. In *Eurocrypt 2011 (LNCS)*, Vol. 6632. 169–188.
- [8] Dan Bogdanov, Liina Kamm, Baldur Kubo, Reimo Rebane, Ville Sokk, and Riivo Talviste. 2016. Students and Taxes: A Privacy-Preserving Social Study Using Secure Computation. In *Privacy Enhancing Technologies Symposium (PETS)*.
- [9] Dan Bogdanov, Sven Laur, and Jan Willemson. 2008. Sharemind: A Framework for Fast Privacy-Preserving Computations. In *ESORICS 2008 (LNCS)*, Vol. 5283. 192–206.
- [10] Peter Bogetoft, Dan Lund Christensen, Ivan Damgård, Martin Geisler, Thomas Jakobsen, Mikkel Kroigaard, Janus Dam Nielsen, Jesper Buus Nielsen, Kurt Nielsen, Jakob Pagter, Michael I. Schwartzbach, and Tomas Toft. 2009. Secure Multiparty Computation Goes Live. In *FC 2009 (LNCS)*, Vol. 5628. 325–343.
- [11] Martin Burkhart, Mario Strasser, Dilip Many, and Xenofontas Dimitropoulos. 2010. SEPIA: Privacy-preserving aggregation of multi-domain network events and statistics. In *19th USENIX Security Symposium*, Ian Goldberg (Ed.). USENIX Association, Washington, D.C., USA.
- [12] Sai Sheshank Burra, Enrique Larraia, Jesper Buus Nielsen, Peter Sebastian Nordholt, Claudio Orlandi, Emmanuela Orsini, Peter Scholl, and Nigel P. Smart. 2015. High Performance Multi-Party Computation for Binary Circuits Based on Oblivious Transfer. *Cryptology ePrint Archive*, Report 2015/472. (2015). <http://eprint.iacr.org/2015/472>.
- [13] Seung Geol Choi, Kyung-Wook Hwang, Jonathan Katz, Tal Malkin, and Dan Rubenstein. 2012. Secure Multi-Party Computation of Boolean Circuits with Applications to Privacy in On-Line Marketplaces. In *CT-RSA 2012 (LNCS)*, Vol. 7178. 416–432.
- [14] Seung Geol Choi, Jonathan Katz, Alex J. Malozemoff, and Vassilis Zikas. 2014. Efficient Three-Party Computation from Cut-and-Choose. In *Crypto 2014, Part II (LNCS)*, Vol. 8617. 513–530.
- [15] Ivan Damgård, Martin Geisler, Mikkel Kroigaard, and Jesper Buus Nielsen. 2009. Asynchronous Multiparty Computation: Theory and Implementation. In *PKC 2009 (LNCS)*, Vol. 5443. 160–179.
- [16] Ivan Damgård and Yuval Ishai. 2005. Constant-Round Multiparty Computation Using a Black-Box Pseudorandom Generator. In *Crypto 2005 (LNCS)*, Vol. 3621. 378–394.
- [17] Ivan Damgård, Marcel Keller, Enrique Larraia, Christian Miles, and Nigel P. Smart. 2012. Implementing AES via an Actively/Covertly Secure Dishonest-Majority MPC Protocol. In *SCN 12 (LNCS)*, Vol. 7485. 241–263.
- [18] Ivan Damgård, Marcel Keller, Enrique Larraia, Valerio Pastro, Peter Scholl, and Nigel P. Smart. 2013. Practical Covertly Secure MPC for Dishonest Majority - Or: Breaking the SPDZ Limits. In *ESORICS 2013 (LNCS)*, Vol. 8134. 1–18.
- [19] Ivan Damgård, Valerio Pastro, Nigel P. Smart, and Sarah Zakarias. 2012. Multiparty Computation from Somewhat Homomorphic Encryption. In *Crypto 2012 (LNCS)*, Vol. 7417. 643–662.
- [20] Tore Kasper Frederiksen, Marcel Keller, Emmanuela Orsini, and Peter Scholl. 2015. A Unified Approach to MPC with Preprocessing Using OT. In *ASIACRYPT 2015, Part I (LNCS)*, Vol. 9452. 711–735.
- [21] Jun Furukawa, Yehuda Lindell, Ariel Nof, and Or Weinstein. 2017. High-Throughput Secure Three-Party Computation for Malicious Adversaries and an Honest Majority. In *Eurocrypt 2017, Part II (LNCS)*, Vol. 10211. 225–255.

- [22] Oded Goldreich. 2009. *Foundations of Cryptography: Volume 2, Basic Applications*. Vol. 2. Cambridge University Press.
- [23] Oded Goldreich, Silvio Micali, and Avi Wigderson. 1987. How to Play any Mental Game, or A Completeness Theorem for Protocols with Honest Majority. In *19th ACM STOC*. 218–229.
- [24] Shafi Goldwasser and Yehuda Lindell. 2005. Secure Multi-Party Computation without Agreement. *Journal of Cryptology* 18, 3 (July 2005), 247–287.
- [25] Carmit Hazay, Peter Scholl, and Eduardo Soria-Vazquez. 2017. Low Cost Constant Round MPC Combining BMR and Oblivious Transfer. Cryptology ePrint Archive, Report 2017/214. (2017). To appear in *Asiacrypt 2017*.
- [26] Thomas P. Jakobsen, Marc X. Makkes, and Janus Dam Nielsen. 2010. Efficient Implementation of the Orlandi Protocol. In *ACNS 10 (LNCS)*, Vol. 6123. 255–272.
- [27] Marcel Keller, Emmanuela Orsini, and Peter Scholl. 2015. Actively Secure OT Extension with Optimal Overhead. In *Crypto 2015, Part I (LNCS)*, Vol. 9215. 724–741.
- [28] Marcel Keller, Emmanuela Orsini, and Peter Scholl. 2016. MASCOT: Faster Malicious Arithmetic Secure Computation with Oblivious Transfer. In *ACM CCS 2016*. 830–842.
- [29] Marcel Keller, Peter Scholl, and Nigel P. Smart. 2013. An architecture for practical actively secure MPC with dishonest majority. In *ACM CCS 2013*. 549–560.
- [30] Enrique Larraia, Emmanuela Orsini, and Nigel P. Smart. 2014. Dishonest Majority Multi-Party Computation for Binary Circuits. In *Crypto 2014, Part II (LNCS)*, Vol. 8617. 495–512.
- [31] Yehuda Lindell, Benny Pinkas, Nigel P. Smart, and Avishay Yanai. 2015. Efficient Constant Round Multi-party Computation Combining BMR and SPDZ. In *Crypto 2015, Part II (LNCS)*, Vol. 9216. 319–338.
- [32] Yehuda Lindell, Nigel P. Smart, and Eduardo Soria-Vazquez. 2016. More Efficient Constant-Round Multi-party Computation from BMR and SHE. In *TCC 2016-B, Part I (LNCS)*, Vol. 9985. 554–581.
- [33] Payman Mohassel, Mike Rosulek, and Ye Zhang. 2015. Fast and Secure Three-party Computation: The Garbled Circuit Approach. In *ACM CCS 2015*. 591–602.
- [34] Jesper Nielsen, Thomas Schneider, and Roberto Trifiletti. 2017. Constant-Round Maliciously Secure 2PC with Function-Independent Preprocessing Using LEGO. In *Network and Distributed System Security Symposium (NDSS)*.
- [35] Jesper Buus Nielsen, Peter Sebastian Nordholt, Claudio Orlandi, and Sai Sheshank Burra. 2012. A New Approach to Practical Active-Secure Two-Party Computation. In *Crypto 2012 (LNCS)*, Vol. 7417. 681–700.
- [36] Xiao Wang, Alex J. Malozemoff, and Jonathan Katz. 2016. EMP-Toolkit: Efficient Multiparty Computation Toolkit. <https://github.com/emp-toolkit>. (2016).
- [37] Xiao Wang, Samuel Ranellucci, and Jonathan Katz. 2017. Authenticated Garbling and Efficient Maliciously Secure Two-Party Computation. In *ACM CCS 2017*.
- [38] Andrew Chi-Chih Yao. 1986. How to Generate and Exchange Secrets. In *IEEE FOCS*. 162–167.

A SECURITY PROOFS

A.1 Proof of the Main Protocol

THEOREM A.1. *If H is modeled as a random oracle, the protocol in Figures 2 and 3 securely realizes \mathcal{F}_{mpc} in the \mathcal{F}_{pre} -hybrid model with security $\text{negl}(\kappa)$ against an adversary corrupting up to $n - 1$ parties.*

PROOF. We consider separately the case where $P_1 \in \mathcal{H}$ and where $P_1 \in \mathcal{M}$ and $P_2 \in \mathcal{H}$. The case where $P_1 \in \mathcal{M}$ and $P_i \in \mathcal{H}$ for some $i \geq 3$ is similar to the second case.

Honest P_1 . Let \mathcal{A} be an adversary corrupting $\{P_i\}_{i \in \mathcal{M}}$. We construct a simulator \mathcal{S} that runs \mathcal{A} as a subroutine and plays the role of $\{P_i\}_{i \in \mathcal{M}}$ in the ideal world involving an ideal functionality \mathcal{F}_{mpc} evaluating f . \mathcal{S} is defined as follows.

- 1-4 \mathcal{S} acts as honest $\{P_i\}_{i \in \mathcal{H}}$ and plays the functionality of \mathcal{F}_{pre} , recording all outputs. If any honest party or \mathcal{F}_{pre} would abort, \mathcal{S} outputs whatever \mathcal{A} outputs and then aborts.
- 5 \mathcal{S} interacts with \mathcal{A} acting as an honest $\{P_i\}_{i \in \mathcal{H}}$, using input $\{x^i := 0\}_{i \in \mathcal{H}}$. For each $i \in \mathcal{M}$, $w \in \mathcal{I}_i$, \mathcal{S} receives \hat{x}_w^i and computes $x_w^i := \hat{x}_w^i \oplus \bigoplus_{i \in [n]} r_w^i$. If any honest party would abort, \mathcal{S} outputs whatever \mathcal{A} outputs and aborts.
- 6 \mathcal{S} interacts with \mathcal{A} acting as honest $\{P_i\}_{i \in \mathcal{H}}$, using input $x^1 := 0$.

7-8 \mathcal{S} interacts with \mathcal{A} acting as honest $\{P_i\}_{i \in \mathcal{H}}$. If an honest P_1 would abort, \mathcal{S} outputs whatever \mathcal{A} outputs and aborts; otherwise for each $i \in \mathcal{M}$, \mathcal{S} sends (input, x^i) on behalf of P_i to \mathcal{F}_{mpc} .

At any time, \mathcal{S} will answer \mathcal{A} 's global key query honestly, since \mathcal{S} knows the global keys of all parties.

Note that since the global keys are randomly selected from $\{0, 1\}^\kappa$, \mathcal{A} cannot guess any global key with more than negligible probability. Therefore, in the following, we will assume that it does not happen.

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

Hybrid₁. Same as the hybrid-world protocol, where \mathcal{S} plays the role of honest $\{P_i\}_{i \in \mathcal{H}}$, using the actual inputs $\{x^i\}_{i \in \mathcal{H}}$.

Hybrid₂. Same as **Hybrid₁**, except that in step 5, for each $i \in \mathcal{M}$, $w \in \mathcal{I}_i$, \mathcal{S} receives \hat{x}_w^i and computes $x_w^i := \hat{x}_w^i \oplus \bigoplus_{i \in [n]} r_w^i$. If any honest party would abort, \mathcal{S} outputs whatever \mathcal{A} outputs; otherwise for each $i \in \mathcal{M}$, \mathcal{S} sends (input, x^i) on behalf of P_i to \mathcal{F}_{mpc} .

The views produced by the two hybrids are exactly the same. According to Lemma A.2, P_1 will learn the same output in both hybrids with all but negligible probability.

Hybrid₃. Same as **Hybrid₂**, except that, for each $i \in \mathcal{H}$, \mathcal{S} computes $\{r_w^i\}_{w \in \mathcal{I}_i}$ as follows: \mathcal{S} first randomly pick $\{u_w^i\}_{w \in \mathcal{I}_i}$, and then computes $r_w^i := u_w^i \oplus x_w^i$.

The two Hybrids produce exactly the same view.

Hybrid₄. Same as **Hybrid₃**, except that \mathcal{S} uses $\{x^i = 0\}_{i \in \mathcal{H}}$ as input in step 5 and step 6.

Note that although the distribution of $\{x^i\}_{i \in \mathcal{H}}$ in **Hybrid₃** and **Hybrid₄** are different, the distribution of $\{x_w^i \oplus r_w^i\}_{i \in \mathcal{H}}$ are exactly the same. The views produced by the two Hybrids are therefore the same, we will show that P_1 aborts with the same probability in both Hybrids.

Observe that the only place where P_1 's abort can possibly depend on $\{x^i\}_{i \in \mathcal{H}}$ 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 when P_1 is honest.

Malicious P_1 and honest P_2 . Let \mathcal{A} be an adversary corrupting $\{P_i\}_{i \in \mathcal{M}}$. We construct a simulator \mathcal{S} that runs \mathcal{A} as a subroutine and plays the role of $\{P_i\}_{i \in \mathcal{M}}$ in the ideal world involving an ideal functionality \mathcal{F}_{mpc} evaluating f . \mathcal{S} is defined as follows.

- 1-4 \mathcal{S} acts as honest $\{P_i\}_{i \in \mathcal{H}}$ and plays the functionality of \mathcal{F}_{pre} , recording all outputs. If any honest party would abort, \mathcal{S} output whatever \mathcal{A} outputs and aborts.
- 5-6 \mathcal{S} interacts with \mathcal{A} acting as honest $\{P_i\}_{i \in \mathcal{H}}$, using input $\{x^i := 0\}_{i \in \mathcal{H}}$. For each $i \in \mathcal{M}$, $w \in \mathcal{I}_i$, \mathcal{S} receives \hat{x}_w^i and computes $x_w^i := \hat{x}_w^i \oplus \bigoplus_{i \in [n]} r_w^i$. If any honest party would abort, \mathcal{S} output whatever \mathcal{A} outputs and aborts.

8 For each $i \in \mathcal{M}$, \mathcal{S} sends (input, x^i) on behalf of P_i to \mathcal{F}_{mpc} . If \mathcal{F}_{mpc} aborts, \mathcal{S} aborts, outputting whatever \mathcal{A} outputs. Otherwise, if \mathcal{S} receives z as the output, \mathcal{S} computes $z' := f(y^1, \dots, y^n)$, where $\{y^i := 0\}^{i \in \mathcal{H}}$, and $\{y^i := x^i\}^{i \in \mathcal{M}}$. For each $i \in \mathcal{H}$, $w \in \mathcal{O}$, if $z'_w = z_w$, \mathcal{S} sends $(r_w^i, M_1[r_w^i])$ on behalf of P_i to \mathcal{A} ; otherwise, \mathcal{S} sends $(r_w^i \oplus 1, M_1[r_w^i] \oplus \Delta_1)$.

At any time, \mathcal{S} will answer \mathcal{A} 's global key query honestly, since \mathcal{S} knows the global keys of all parties.

Note that since the global keys are randomly selected from $\{0, 1\}^k$, \mathcal{A} cannot guess any global key with more than negligible probability. Therefore, in the following, we will assume it does not happen.

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

Hybrid₁. Same as the hybrid-world protocol, where \mathcal{S} plays the role of honest $\{P_i\}_{i \in \mathcal{H}}$ using the actual inputs $\{x^i\}_{i \in \mathcal{H}}$.

Hybrid₂. Same as **Hybrid₁**, except that in step 5 and step 6, for each $i \in \mathcal{M}$, $w \in \mathcal{I}_i$, \mathcal{S} receives \hat{x}_w^i and computes $x_w^i := \hat{x}_w^i \oplus \bigoplus_{i \in [n]} r_w^i$. If any honest party would abort, \mathcal{S} outputs whatever \mathcal{A} outputs; otherwise for each $i \in \mathcal{M}$, \mathcal{S} sends (input, x^i) on behalf of P_i to \mathcal{F}_{mpc} .

P_1 does not have output; furthermore the view of \mathcal{A} does not change between the two Hybrids.

Hybrid₃. Same as **Hybrid₂**, except that in step 5 and step 6, \mathcal{S} uses $\{x^i := 0\}_{i \in \mathcal{H}}$ as input and in step 8, \mathcal{S} computes z' as defined. For each $w \in \mathcal{O}$, if $z'_w = z_w$, \mathcal{S} sends $(r_w^i, M_1[r_w^i])$; otherwise, \mathcal{S} sends $(r_w^i \oplus 1, M_1[r_w^i] \oplus \Delta_1)$. \mathcal{A} has no knowledge of r_w^i , therefore r_w^i and $r_w^i \oplus 1$ are indistinguishable.

Note that since \mathcal{S} uses different values for x between the two Hybrids, we also need to show that the distribution of garbled rows opened by P_1 are indistinguishable for the two Hybrids. According to the security of garbled circuits, P_1 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_1 , masked values and garbled keys are indistinguishable between two Hybrids.

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

LEMMA A.2. Consider an \mathcal{A} corrupting parties $\{P_i\}_{i \in \mathcal{M}}$ such that $P_1 \in \mathcal{H}$, and denote $x_w^i := \hat{x}_w^i \oplus \bigoplus_{i=1}^n r_w^i$, where \hat{x}_w^i is the value \mathcal{A} sent, r_w^i are the values from \mathcal{F}_{Pre} . With all but negligible probability, P_1 either aborts or learns $z = f(x^1, \dots, x^n)$.

PROOF. Define z_w^* as the correct wire values computed using x defined above and y , z_w as the actually wire values P_1 holds in the evaluation.

We will first show that P_1 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_1 has $\{z_w^* \oplus \lambda_w = z_w \oplus \lambda_w\}_{w \in \{\alpha, \beta\}}$, then P_1 also obtains $z_w^* \oplus \lambda_w = z_w \oplus \lambda_w$.

- $T = \oplus$: It is true according to the following: $z_\gamma^* \oplus \lambda_\gamma = (z_\alpha^* \oplus \lambda_\alpha) \oplus (z_\beta^* \oplus \lambda_\beta) = (z_\alpha \oplus \lambda_\alpha) \oplus (z_\beta \oplus \lambda_\beta) z_\gamma \oplus \lambda_\gamma$
- $T = \wedge$: According to the protocol, P_1 will open the garbled row defined by $i := 2(z_\alpha \oplus \lambda_\alpha) + (z_\beta \oplus \lambda_\beta)$. If P_1 learns $z_\gamma \oplus \lambda_\gamma \neq z_\gamma^* \oplus \lambda_\gamma$, then it means that P_1 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_1 learns the correct masked output. P_1 can therefore learn the correct output $f(x, y)$ unless \mathcal{A} is able to flip $\{r_w\}_{w \in \mathcal{O}}$, which, again, only happens with negligible probability. \square

A.2 Multi-Party Authenticated Bits

THEOREM A.3. The protocol in Figure 5 securely realizes $\mathcal{F}_{\text{aBit}}^n$ with statistical security $2^{-\rho}$ in the $\mathcal{F}_{\text{aBit}}^2$ -hybrid model.

PROOF. We consider two cases.

Case 1: $P_i \in \mathcal{H}$. Note that in this case, the only way malicious parties can break the protocol is by learning some information about $\{x_i\}_{i \in [\ell]}$ in the checking step. However, we will show that, because we “throw out” the last 2ρ authenticated bits, the adversary can learn nothing about x 's.

We use s_j to denote the last 2ρ bits of r in the j -th check. According to Lemma A.4 and the parameters we chose, the probability that any subset of $\{s_j\}_{j \in [2\rho]}$ is linearly independent is $1 - 2^{-\rho}$. Now we will show that if linear independence holds then the adversary cannot learn anything.

For the j -checking, $X = \left(\bigoplus_{m=1}^{\ell} r_m x_m\right) \oplus \left(\bigoplus_{m=1}^{2\rho} s_m x_{\ell+m}\right)$. Note that $\bigoplus_{m=1}^{2\rho} s_m x_{\ell+m}$ from each checking are independent random bits, where $\{x_m\}_{m=\ell}^{\ell'}$ is random. This is true because the s_i 's are linearly independent. Therefore, $\bigoplus_{m=1}^{2\rho} s_m x_{\ell+m}$ acts as one-time pad to $\bigoplus_{m=1}^{\ell} r_m x_m$. Given the above, the simulation is straightforward. Note that for all global key queries, \mathcal{S} can answer them honestly, since \mathcal{S} knows the global key for both parties.

Case 2: $P_i \in \mathcal{M}$. The simulation is straightforward if we could show that for any \mathcal{A} who uses inconsistent x 's can pass all 2ρ checks with at most negligible probability. This is what we will proceed to show.

Suppose that \mathcal{A} sends x^1 to $\mathcal{F}_{\text{aBit}}^2$ when interacting with one honest party, and uses a different x^2 with another honest party, where $x^1 \neq x^2$. We also assume that \mathcal{A} passes all checks. Note that for the j -th checking, if \mathcal{A} is not able to forge a MAC, then the probability that the checking passes is the probability that $X_j = \bigoplus_m r_m x_m^1$ and that $X_j = \bigoplus_m r_m x_m^2$.

$$\begin{aligned} & \Pr \left\{ \bigoplus_m r_m x_m^1 = \bigoplus_m r_m x_m^2 \right\} \\ &= \Pr \left\{ \bigoplus_m r_m (x_m^1 \oplus x_m^2) = 0 \right\} \\ &= \Pr \left\{ \bigoplus_{m \in I} r_m = 0 : I \text{ is the set of indices where } x_m^1 \neq x_m^2 \right\} \\ &= 1/2 \end{aligned}$$

Each checking is independent as long as r is selected independently. Therefore, \mathcal{A} can pass all checks with probability at most $2^{-2\rho}$. \square

LEMMA A.4. *Let r_1, \dots, r_ℓ be random bit vectors of length k . With probability at most $2^{\ell-k}$, there exists some subset $I \subset [\ell]$, such that*

$$\bigoplus_{i \in I} r_i = 0$$

PROOF. Note that given a fixed interval $I \subset [\ell]$, the probability that $\bigoplus_{i \in I} r_i = 0$ is 2^{-k} . According to the union bound, the probability that any subset $I \subset [\ell]$ has $\bigoplus_{i \in I} r_i = 0$ is $2^{-k} \times 2^\ell = 2^{\ell-k}$. \square

A.3 Multi-Party Authenticated Bits

THEOREM A.5. *The protocol in Figure 7 securely realizes $\mathcal{F}_{\text{aShare}}$ with statistical security $2^{-\rho}$ in the $\mathcal{F}_{\text{aBit}}^n$ -hybrid model.*

PROOF. Without loss of generality, we consider the case when P_1 is honest. The simulator plays the role of $\mathcal{F}_{\text{aBit}}^n$ honestly, recording all values it sends to \mathcal{A} and values \mathcal{A} sent to $\mathcal{F}_{\text{aBit}}^n$. \mathcal{S} acts as honest parties and check for each $i \in \mathcal{M}$, if P_i sent consistent Δ_i in all instructions to $\mathcal{F}_{\text{aBit}}^n$. If not, \mathcal{S} aborts outputting whatever \mathcal{A} outputs.

Note that this simulator has a $2^{-\rho}$ statistical difference to the real world execution given Lemma A.6 \square

LEMMA A.6. *When a malicious P_i computes MACs with P_j , denote Δ_i^j as the value P_i sent to $\mathcal{F}_{\text{aBit}}^n$. If for some $\Delta_i^{j1} \neq \Delta_i^{j2}$, the honest parties abort with probability at least $1 - 2^{-\rho}$.*

PROOF. In the following, we will prove that malicious party passes each single test with probability $1/2$, independently. Since malicious parties are ensured to use the same Δ for each party, it can either cheat for all bits or be honest for all bits. Therefore cheating malicious parties cannot pass all checks with more than $2^{-\rho}$ probability.

We will prove by contradiction. Suppose at least one party uses inconsistent value and the check passes. We use K^i to denote the value P_i opened in step 3 (d), and use $\{M_k^i\}_{k \neq i}$ to denote the value committed in step 3 (c). First compute $Q^i := K^i \oplus \bigoplus_{k \neq i} K_i[x_{\ell+r}^k]$ and $Q_k^i = M_k^i \oplus M_k[x_{\ell+r}^i]$. Since the check for P_i passes, we know that the following is zero.

$$\begin{aligned} K^i \oplus \bigoplus_{k \neq i} M_k^i &= \left(\bigoplus_{k \neq i} K_i[x_{\ell+r}^k] \right) \oplus Q^i \oplus \left(\bigoplus_{k \neq i} M_i[x_{\ell+r}^k] \oplus Q_k^i \right) \\ &= \left(\bigoplus_{k \neq i} K_i[x_{\ell+r}^k] \right) \oplus Q^i \\ &\quad \oplus \left(\bigoplus_{k \neq i} K_i[x_{\ell+r}^k] \oplus x_{\ell+r}^k \Delta_i^k \right) \oplus \bigoplus_{k \neq i} Q_k^i \\ &= \left(Q^i \oplus \bigoplus_{k \neq i} Q_k^i \right) \oplus \left(\bigoplus_{k \neq i} x_{\ell+r}^k \Delta_i^k \right) \end{aligned}$$

If P_i uses $\Delta_{\ell+r}^{k1} \neq \Delta_{\ell+r}^{k2}$ for some $k1 \neq k2, k1, k2 \in \mathcal{H}$, then the value of $(Q^i \oplus \bigoplus_{k \neq i} Q_k^i)$ that makes the equation as 0 is different

depending on the value of $x_{\ell+r}^{k1}$ and $x_{\ell+r}^{k2}$. This means that P_i needs to guess at least one of them to pass the check. \square

A.4 Half-Authenticated AND Triple

LEMMA A.7. *If H is modeled as a random oracle, the protocol in Figure 9 securely realizes $\mathcal{F}_{\text{HaAND}}$ in the $\mathcal{F}_{\text{aShare}}$ -hybrid model.*

PROOF. Note that for each $i \in [n]$, P_i has values $\{s^j, t^j\}_{j \neq i}$. We denote the value s^j, t^j held by P_i as s_i^j, t_i^j .

Correctness. First we will show the correctness of the protocol. We further first show that for any $i \neq j$, $s_i^j \oplus t_j^i = x^j y_i^i$. We will discuss in two cases:

- $x^j = 0$. In this case, P_j obtains $t_j^i = s^j$.
- $x^j = 1$. In this case, P_j obtains $t_j^i = s^j \oplus y_j^i$.

In any case, the above equation holds. Now the correctness of the protocol can be seen given the following equation.

$$\begin{aligned} \bigoplus_i \bigoplus_{j \neq i} x^j y_i^i &= \bigoplus_i \bigoplus_{j \neq i} (s_i^j \oplus t_j^i) \\ &= \bigoplus_i \bigoplus_{j \neq i} s_i^j \oplus \bigoplus_i \bigoplus_{j \neq i} t_j^i \\ &= \bigoplus_i \bigoplus_{j \neq i} s_j^i \oplus \bigoplus_i \bigoplus_{j \neq i} t_j^i \\ &= \bigoplus_i \left(\bigoplus_{j \neq i} s_j^i \oplus t_j^i \right) \\ &= \bigoplus_i v^i \end{aligned}$$

Simulation proof. We will prove the security assuming P_1 is honest, that is, $P_1 \in \mathcal{H}$. The simulation is as follows:

1. \mathcal{S} plays the role of $\mathcal{F}_{\text{aShare}}$ storing all values used.
2. For each pair $i \neq j$, such that $i \in \mathcal{M}$, \mathcal{S} obtains (H_0, H_1) sent by malicious P_i . \mathcal{S} computes $s^j := H_0 \oplus \text{Lsb}(H(K_i[x^j]))$ and $y_j^i := H_1 \oplus \text{Lsb}(H(K_i[x^j] \oplus \Delta_i)) \oplus s^j$. For each $i \in \mathcal{M}$, \mathcal{S} sends $(i, \{y_j^i\}_{j \neq i})$ to $\mathcal{F}_{\text{HaAND}}$, which sends back $\{v^i\}_{i \in \mathcal{M}}$.
3. For each $i \in \mathcal{M}$, \mathcal{S} picks random $\{t^k\}_{k \in \mathcal{H}}$, such that $\bigoplus_{k \neq i} s^i \oplus \bigoplus_{k \neq i, k \in \mathcal{M}} t_k^i \oplus \bigoplus_{k \neq i, k \in \mathcal{H}} t^i = v^i$. For each $j \in \mathcal{H}$, \mathcal{S} computes $H_{x^i} = \text{Lsb}(H(K[x_i] \oplus x_i \Delta_j)) \oplus t^j$, and picks a random $H_{1 \oplus x^i}$. \mathcal{S} sends (H_0, H_1) to P_i on behalf of an honest P_j .

First of all, the first two steps are perfect simulation. For the last step, it is also a perfect simulation: first the one that is not opened is random since H is a random oracle. The other value is also random, depending on the value of s^j . However, in order to make the joint distribution of the value \mathcal{A} learns here and the output of an honest P_2 indistinguishable between ideal and real world protocol, t^k are tweaked such that \mathcal{A} will learn the same value in both Hybrids. \square

A.5 Multi-Party Leaky Authenticated AND Triple

We have described the protocol and the key ideas of the proof in the main body. Here we will directly proceed to the proof.

THEOREM A.8. *If H is modeled as a random oracle, the protocol in Figure 11 securely realizes $\mathcal{F}_{\text{LaAND}}$ in the $(\mathcal{F}_{\text{aShare}}, \mathcal{F}_{\text{HaAND}})$ -hybrid model.*

PROOF. We construct a simulator in the following. For all global key queries, \mathcal{S} redirect them to $\mathcal{F}_{\text{aShare}}$ and redirect the answer to \mathcal{A} .

1. \mathcal{S} plays the role of $\mathcal{F}_{\text{aShare}}$ storing all information sent to parties.
- 2-3 \mathcal{S} obtains $(i, \{y_j^i\}_{j \neq i})$ for each $P_i \in \mathcal{M}$. \mathcal{S} also obtains $\{e^i\}_{i \in \mathcal{M}}$. \mathcal{A} broadcasts. \mathcal{S} first computes e^{*i} , which are what an honest P_i would have broadcast and compute $q_i := e^i \oplus e^{*i}$. \mathcal{S} further computes $r_{i,j} := y_j^i \oplus y^i$, where y^i is the value \mathcal{S} used when playing the role of $\mathcal{F}_{\text{aShare}}$. \mathcal{S} computes $r_i := \bigoplus_{j \in \mathcal{M}, j \neq i} r_{j,i}$, $q := \bigoplus_{i \in \mathcal{M}} q_i$, and sends $(q, \{r_i\}_n)$ to $\mathcal{F}_{\text{LaAND}}$. If $\mathcal{F}_{\text{LaAND}}$ terminates, \mathcal{S} follows the protocol as honest parties and abort in step 7.

- 4-5. For each $i \in \mathcal{M}$, \mathcal{S} receives $\{U_{i,j}\}_{j \in \mathcal{H}}$ from P_i . \mathcal{S} picks random $\{U_{j,i}\}_{j \in \mathcal{H}}$ and sends them to P_i playing the role of P_j for each $j \in \mathcal{H}$.

- 6-7 If $\mathcal{F}_{\text{LaAND}}$ terminates in step 2, then \mathcal{S} follows the protocol as honest parties and abort in step 7. If the equation hold, \mathcal{S} will extract another selective failure attack query.

Similar to the unforgeability proof, we use $U_{i,j}^*$ and H_i^* to denote the values that an honest party would have compute, and define $Q_{i,j} := U_{i,j}^* \oplus U_{i,j}$, $Q_i = H_i^* \oplus H_i$. This means that a malicious P_i uses some $Q_{i,j}$, then P_j will obtain some $M_i[x^j]_{\Phi_i}$ with an additive error of $x^j Q_{i,j}$. \mathcal{S} defines $R_k = \bigoplus_{i \neq k, i \in \mathcal{H}} Q_{k,i}$. \mathcal{S} sends $(\bigoplus_{i \in \mathcal{M}} Q_i, \{R_i\}_{i \in [n]})$ to $\mathcal{F}_{\text{LaAND}}$. If $\mathcal{F}_{\text{LaAND}}$ terminates, \mathcal{S} aborts outputting whatever \mathcal{A} outputs; otherwise, \mathcal{S} obtains $\{H_i\}_{i \in \mathcal{M}}$ and picks random $\{H_i\}_{i \in \mathcal{H}}$ such that $\bigoplus_i H_i = 0$.

Note that the first five steps are perfectly indistinguishable given that H is a random oracle, except that \mathcal{A} can perform a selective failure attack. We will show that the probability of abort due to this attack is the same between real-world protocol and ideal-world protocol. The probability that the value \mathcal{A} sent in step 2 and 3 cause an abort is the same as \mathcal{S} 's query to $\mathcal{F}_{\text{LaAND}}$, noticing that the following is true.

$$\begin{aligned}
& \bigoplus_i \bigoplus_{j \neq i} x_i y_j^i \oplus \bigoplus_i x^i y^i \oplus \bigoplus_i (e^i \oplus r^i) \\
&= \bigoplus_i \bigoplus_{j \in \mathcal{M}, j \neq i} x_i y_j^i \oplus \bigoplus_i x_i y_i^i \oplus \bigoplus_{j \in \mathcal{H}, j \neq i} x_i y_j^i \oplus \bigoplus_i x^i y^i \oplus \bigoplus_i z^i \oplus \bigoplus_{i \in \mathcal{M}} q_i \\
&= \bigoplus_i \bigoplus_{j \in \mathcal{M}, j \neq i} x_i r_{j,i} \oplus \bigoplus_i \bigoplus_{j \neq i} x_i y_j^i \oplus \bigoplus_i z^i \oplus \bigoplus_{i \in \mathcal{M}} q_i \\
&= \bigoplus_i \bigoplus_{j \in \mathcal{M}, j \neq i} x_i r_{j,i} \oplus \bigoplus_{i \in \mathcal{M}} q_i \\
&= \left(\bigoplus_i x_i \oplus \bigoplus_{j \in \mathcal{M}, j \neq i} r_{j,i} \right) \oplus \bigoplus_{i \in \mathcal{M}} q_i
\end{aligned}$$

We will focus on the last step. If in step 6, it is the case that $(\bigoplus_i x^i) (\bigoplus_i y^i) \neq (\bigoplus_i z^i)$, then it is easy to see that the views are indistinguishable: all parties behave the same between hybrids. According to the unforgeability lemma, the protocol will abort with all but negligible probability. In the following, we will further focus on the case when the equation holds.

Note that in the idea world protocol, all H_i from \mathcal{H} are picked randomly. We need to show that in the real world protocol all H_i 's

Continent	Region	
North America	North Virginia	Ohio
	North California	Oregon
		Toronto
Europe	Ireland	London
		Frankfurt
Asia	Mumbai	Tokyo
	Seoul	Singapore
Australia	Sydney	
South America	São Paulo	

Table 9: List of all Amazon EC2 regions used in the WAN experiment.

is also a random share of 0. In particular, we define

$$F_i^* = \left(\bigoplus_{k \neq i} K_i[x^k]_{\Phi_i} \oplus M_k[x^i]_{\Phi_k} \right)$$

and will first show that for any proper subset $S \subset \mathcal{H}$, $\bigoplus_{i \in S} F_i$ is indistinguishable from random to the \mathcal{A} . We use e to denote an honest party such that $e \in \mathcal{H}$, $e \notin S$. Such e always exists, since S is a proper subset of \mathcal{H} .

$$\begin{aligned}
\bigoplus_{i \in S} F_i^* &= \bigoplus_{i \in S} \bigoplus_{k \neq i} (K_i[x^k]_{\Phi_i} \oplus M_k[x^i]_{\Phi_k}) \\
&= \bigoplus_{i \in S} \bigoplus_{k \neq i} (K_i[x^k]_{\Phi_i}) \oplus \bigoplus_{i \in S} \bigoplus_{k \neq i} (M_k[x^i]_{\Phi_k}) \\
&= \bigoplus_{i \in S} \bigoplus_{k \neq i} (K_i[x^k]_{\Phi_i}) \oplus \bigoplus_{k \in S} \bigoplus_{i \neq k} (M_i[x^k]_{\Phi_i}) \\
&= \bigoplus_{i \in S} \bigoplus_{k \neq i} (K_i[x^k]_{\Phi_i}) \oplus \bigoplus_{i \in [n]} \bigoplus_{k \in S, k \neq i} (M_i[x^k]_{\Phi_i})
\end{aligned}$$

From the equation, it is clear that for $i \in S$, $K_e[x^i]$ is not in the computation, while $M_e[x^i]$ is. Since $K_e[x^i]$ is randomly picked by P_e , we know $\bigoplus_{i \in S} F_i^*$ is random. Therefore we can see that for any proper subset $S \subset \mathcal{H}$, $\bigoplus_{i \in S} H_i$ is indistinguishable from random.

Finally we need to show that the probability of abort due to selective failure attack is also the same. This is straightforward given the equation used in the unforgeability proof:

$$\begin{aligned}
\bigoplus_i H_i &= \bigoplus_{i \in \mathcal{M}} H_i \oplus \bigoplus_{i \in \mathcal{H}} H_i \\
&= \bigoplus_{i \in \mathcal{M}} (H_i^* \oplus Q_i) \oplus \bigoplus_{i \in \mathcal{H}} \left(H_i^* \oplus \left(\bigoplus_{k \neq i} x^k Q_{k,i} \right) \right) \\
&= \bigoplus_i H_i^* \oplus \bigoplus_{i \in \mathcal{M}} Q_i \oplus \bigoplus_{i \in \mathcal{H}} \left(\bigoplus_{k \neq i} x^k Q_{k,i} \right) \\
&= \bigoplus_{i \in \mathcal{M}} Q_i \oplus \bigoplus_{i \in \mathcal{H}} \left(\bigoplus_{k \neq i} x^k Q_{k,i} \right) \\
&= \bigoplus_{i \in \mathcal{M}} Q_i \oplus \bigoplus_k \left(x^k \bigoplus_{i \in \mathcal{H}, i \neq k} Q_{k,i} \right) \\
&= \bigoplus_{i \in \mathcal{M}} Q_i \oplus \bigoplus_k x^k R_k
\end{aligned}$$

□

n		LAN					WAN				
		Setup	Indep.	Depen.	Online	Total	Setup	Indep.	Depen.	Online	Total
2	AES	33.5	37.4	7.6	1.3	79.8	43.7	286.7	32.5	12.1	375.0
	SHA-1	33.8	148.2	25.2	5.0	212.2	43.9	587.0	101.7	15.9	748.5
	SHA-256	34.3	329.3	58.7	11.8	434.1	44.3	1097.0	205.2	23.1	1369.6
3	AES	35.2	46.7	11.5	2.0	95.4	213.6	1455.1	183.8	66.0	1918.5
	SHA-1	35.5	203.6	36.2	9.3	284.6	212.5	2038.4	270.7	74.7	2596.3
	SHA-256	36.7	470.1	88.0	23.4	618.2	213.7	2939.7	511.0	87.8	3752.2
4	AES	36.5	57.7	15.7	3.2	113.1	321.2	1903.3	319.2	85.4	2629.1
	SHA-1	36.8	265.5	50.3	13.7	366.3	322.7	2665.4	358.7	96.6	3443.4
	SHA-256	37.9	632.2	123.7	31.6	825.4	328.2	3669.0	786.5	113.1	4896.8
5	AES	38.4	72.4	19.9	4.3	135.0	727.8	1988.6	327.6	82.7	3126.7
	SHA-1	39.2	326.2	73.9	20.8	460.1	719.0	2828.3	442.5	106.0	4095.8
	SHA-256	40.0	792.3	182.7	48.9	1063.9	728.4	3953.5	1101.1	131.2	5914.2
6	AES	42.8	85.3	26.1	5.4	159.6	1151.8	3432.1	511.2	87.8	5182.9
	SHA-1	44.5	410.8	110.3	26.1	591.7	1131.1	4921.8	1159.2	118.6	7330.7
	SHA-256	46.0	968.2	282.0	63.8	1360.0	1112.3	6933.2	1617.7	149.8	9813.0
7	AES	47.0	99.8	31.8	7.3	185.9	1459.5	3986.6	620.3	90.8	6157.2
	SHA-1	49.7	493.1	141.8	31.2	715.8	1480.4	5690.3	936.0	115.4	8222.1
	SHA-256	49.4	1176.3	397.0	74.5	1697.2	1455.3	7862.8	1765.8	155.7	11239.6
8	AES	50.9	114.7	38.7	8.5	212.8	1767.7	4189.7	710.6	100.4	6768.4
	SHA-1	52.4	551.2	180.1	39.9	823.6	1791.9	5920.0	1065.8	132.7	8910.4
	SHA-256	51.4	1328.0	470.7	94.5	1944.6	1807.1	8504.8	2129.4	185.0	12626.3
9	AES	53.4	134.7	45.2	9.7	243.0	3279.2	6237.3	1356.2	198.3	11071.0
	SHA-1	54.4	613.6	236.7	47.2	951.9	3263.7	8843.3	2769.5	246.9	15123.4
	SHA-256	54.7	1505.4	573.5	114.9	2248.5	3274.6	12712.0	4705.4	315.2	21007.2
10	AES	59.3	144.2	53.0	11.7	268.2	4502.8	6279.1	1388.2	200.7	12370.8
	SHA-1	60.3	721.0	285.2	54.9	1121.4	4479.5	8962.9	2671.9	252.4	16366.7
	SHA-256	60.2	1694.9	658.3	128.4	2541.8	4524.3	12465.7	6037.4	328.9	23356.3
11	AES	61.3	165.8	60.9	15.0	303.0	4709.8	7083.2	1534.9	205.0	13532.9
	SHA-1	63.0	798.3	318.1	69.8	1249.2	4823.6	12112.6	3539.2	273.7	20749.1
	SHA-256	62.1	1904.9	881.1	173.2	3021.3	4752.5	13432.2	5501.9	378.4	24065.0
12	AES	68.0	181.6	75.8	16.7	342.1	4860.2	7349.1	1718.9	232.1	14160.3
	SHA-1	71.4	872.8	382.4	79.3	1405.9	4826.6	10596.2	3519.9	302.8	19245.5
	SHA-256	69.7	2045.5	1033.9	200.5	3349.6	4814.5	16738.7	6743.5	437.9	28734.6
13	AES	73.5	237.2	85.6	18.4	414.7	6703.6	8531.2	1932.3	248.4	17415.5
	SHA-1	74.0	1281.8	465.1	88.8	1909.7	6743.2	12151.9	4051.9	324.5	23271.5
	SHA-256	75.5	2953.4	1277.7	216.5	4523.1	6682.9	18355.0	8879.7	490.8	34408.4
14	AES	76.8	258.2	102.5	20.3	457.8	8710.8	9412.2	1947.0	250.2	20320.2
	SHA-1	77.0	1375.3	546.7	91.2	2090.2	8654.9	13512.7	4118.9	338.7	26625.2
	SHA-256	79.7	3283.2	1573.0	238.8	5174.7	8714.4	18104.8	8236.1	480.2	35535.5

Table 10: Detailed numbers for experiments of basic circuits. Timings are measured in terms of milliseconds.