

MAESTRO: Multi-Party AES Using Lookup Tables

Hiraku Morita
Aarhus University
University of Copenhagen

Erik Pohle
COSIC, KU Leuven

Kunihiko Sadakane
The University of Tokyo

Peter Scholl
Aarhus University

Kazunari Tozawa
The University of Tokyo

Daniel Tschudi
Concordium
Eastern Switzerland University of Applied Sciences (OST)

Abstract

Secure multi-party computation (MPC) enables multiple distrusting parties to jointly compute a function while keeping their inputs private. Computing the AES block cipher in MPC, where the key and/or the input are secret-shared among the parties is important for various applications, particularly threshold cryptography.

In this work, we propose a family of dedicated, high-performance MPC protocols to compute the non-linear S-box part of AES in the honest majority setting. Our protocols come in both semi-honest and maliciously secure variants. The core technique is a combination of lookup table protocols based on random one-hot vectors and the decomposition of finite field inversion in $GF(2^8)$ into multiplications and inversion in the smaller field $GF(2^4)$, taking inspiration from ideas used for hardware implementations of AES. We also apply and improve the analysis of a batch verification technique for checking inner products with logarithmic communication. This allows us to obtain malicious security with almost no communication overhead, and we use it to obtain new, secure table lookup protocols with only $O(\sqrt{N})$ communication for a table of size N , which may be useful in other applications.

Our protocols have different trade-offs, such as having a similar round complexity as previous state-of-the-art by Chida et al. [WAHC'18] but 37% lower bandwidth costs, or having 27% fewer rounds and 16% lower bandwidth costs. An experimental evaluation in various network conditions using three party replicated secret sharing shows improvements in throughput between 28% and 71% in the semi-honest setting. For malicious security, we improve throughput by 319% to 384% in LAN and by 717% in WAN due to sublinear batch verification.

1 Introduction

Secure multi-party computation (MPC) has become a practical component to realize privacy-preserving computation, improving both privacy and security of existing processes

and data flows. Using MPC, a set of distrusting parties can jointly evaluate a function while keeping their own inputs private. The security relies on distributed trust that no adversary corrupts more parties than the allowed corruption threshold.

MPC protocols are often designed to support generic computation with good performance over commonly used functions. Unfortunately, complex functions such as block ciphers will not always yield optimal performance when using a generic MPC protocol. In these cases, the effort of designing highly-specialized and high-performance MPC protocols for essential building blocks is worthwhile to improve the overall performance of privacy-preserving systems and applications.

In this work, we will focus on such specialized protocols for AES to evaluate the AES block cipher with a secret-shared key and input, which has various applications. Such multi-party AES can be used when clients communicate and exchange data with a cluster of MPC engines in a secure manner using, e.g., oblivious TLS [2] or clients secret-share a secret key to enable MPC parties to distributively decrypt [10, 51] and arbitrarily process their data on their behalf [49] in a secure way, e.g., for secure IoT data collection and processing [1]. Furthermore, it enables secure database joins [43], keyword search [30], private set intersection [36] or allows to increase the trust in systems that rely on centralized secrets, like the Key Distribution Center in the Kerberos authentication protocol or brokered identification systems [17]. Distributed variants [6] and similar distributed authentication protocols [9] rely on AES evaluations with a secret-shared key or can use AES as an oblivious PRF with high confidence in its security. The general case of this class of applications is *threshold cryptography*, which uses MPC to protect cryptographic keys while they are used, adding a layer of distributed trust to a secure system. NIST identified this use-case and initiated the multi-party threshold cryptography project¹ to study, among others, symmetric-key functions like AES-based enciphering, CMAC and HMAC in a *threshold* way.

Since cryptographic primitives tend to be fairly complex to

¹<https://csrc.nist.gov/projects/threshold-cryptography>

evaluate inside an MPC protocol, much effort has been put into designing *MPC-friendly* variants of standard primitives such as hash functions [3], block ciphers [4] and pseudorandom functions [34]. This approach comes with two major drawbacks, however. Firstly, since these primitives are relatively new, they have not been subjected to the same level of scrutiny from cryptanalysts as established, longstanding primitives like AES or SHA-256, so there is less confidence in their security. Secondly, none of these primitives are standardized or even in widespread use. This rules out deploying these types of constructions in applications where MPC must be integrated into an existing system, which only uses standard cryptographic primitives. Aside from these adoption and integration challenges, some MPC-friendly primitives also require many more rounds, e.g., MiMC [3] needs ≈ 8 times more rounds for the same security level compared to AES. Furthermore, large (prime) field arithmetic or extensive use of bit-bit multiplication (e.g., in the linear layers of LowMC [4]) makes them relatively slow to evaluate in plain software.

1.1 Contribution

We present a family of MPC protocols to evaluate the AES block cipher in the multi-party, honest majority setting with both semi-honest and maliciously secure variants. Our contribution is three-fold.

- (i) We securely compute the inversion in $GF(2^8)$ (where 0 maps to 0) with novel techniques. The proposed protocols either rely on lookup table protocols based on preprocessed random one-hot vectors, or on the isomorphism between $GF(2^8)$ and $GF((2^4)^2)$ or on a combination of both techniques. This results in a range of different trade-offs in computational complexity, bandwidth costs and round complexity.
- (ii) We present several lookup table protocols for (t, n) replicated and (n, n) additive secret-sharing, which may be of independent interest. In particular, our protocol based on replicated secret sharing is maliciously secure and only needs $O(\sqrt{N})$ communication for a table of size N . All prior malicious protocols based on information-theoretic primitives require $\Omega(N)$ communication.
- (iii) We implemented the described protocols in the 3-party setting, and experimentally verified their performance in multiple network settings. We obtain improvements between 28% and 71% for online phase and total throughput in the semi-honest setting, compared with the best prior work by Chida et al. [21]. For malicious security, we improve the total throughput by 319% in LAN and by 717% in WAN compared to a version of [21] using cut-and-choose.

1.2 Technical Overview

We now give a brief overview of our main protocols. Their performance characteristics in the 3-party setting, together with

those of the most competitive related work, are summarised in Table 1.

Here, the communication volume represents the number of bits sent by each party, excluding the cost of key expansion steps. As shown in Table 1, the execution of a 10-round oblivious AES protocol involves a communication volume of 4320 bits in 22 communication rounds. We have reduced the round complexity by 27% and communication by 16% compared to the state-of-the-art [21], whilst also adding malicious security with almost no communication overhead.

Warm up: $GF(2^8)$ Inversion from [21]. We start with the inversion protocol [21] that computes $z^{-1} \in GF(2^8)$ as $z^{-1} = z^{254} = (z^{15})^{16} \cdot z^{14}$, where $z^{15} = (z^3)^4 \cdot (z^3)$, $z^{14} = (z^3)^4 \cdot z^2$ and $z^3 = z \cdot z^2$. This can be done with four multiplications in three rounds since squaring is $GF(2)$ -linear.

S-box via $GF(2^4)$ Inversion. Our main approach to evaluating the S-box in MPC views the inversion in $GF(2^8)$ as an extension of $GF(2^4)$ at the cost of one inversion and 3 multiplications in $GF(2^4)$. We observe that this representation considerably simplifies the complexity of the S-box and identify suitable MPC sub-protocols for the $GF(2^4)$ inversion. Although methods for finite field inversion via a tower of field extensions are well-known [39] and have been applied to AES [12, 20], as far as we are aware, they have not been explicitly considered in an MPC context.

$GF(2^4)$ -circuit: Inversion as an Arithmetic Circuit. First, we consider the simple approach of computing $x^{-1} = x^{14} = x^2 \cdot x^4 \cdot x^8$ in $GF(2^4)$. This can be done with just 2 multiplications, since squaring is $GF(2)$ -linear. This already reduces the communication by more than one third, compared with a $GF(2^8)$ circuit-based approach [21], while preserving the same round complexity.

$GF(2^4)$ -LUT-16: Inversion with a Lookup Table. Our next variant uses lookup table techniques to evaluate the inverse. Here, we adapt techniques from the dishonest majority MPC literature [18, 25, 40], which allow to offload the work of computing a lookup table to a preprocessing phase. The high-level idea is to use the preprocessing phase to compute a secret-shared, random one-hot binary vector (that is, all-zero except for a single position) of length equal to the table size, which we do based on a protocol from [40, 42]. Using this, in the online phase the parties can open the masked input and then compute the table lookup with a single linear combination. This reduces communication in the online phase by a further 20%, and reduces round complexity, but adds some cost in an input-independent preprocessing phase.

LUT-256: S-box via a Single Table Lookup. Our second class of protocols treats the S-box as a single lookup table of 256 elements. The main advantage of this approach is that it gets the best round complexity since the S-box can be evaluated in a single round in the online phase (or 10 rounds for 1 AES block). Here, we present two different variants, based on either additive secret sharing ($\langle\langle\cdot\rangle\rangle$) or replicated secret

Table 1: Performance comparison of our multi-party AES protocols and other approaches.

[†] the protocol communicates $O(\kappa)$ during the input phase but nothing is sent during the computation phase.

* these protocols incur an additional $O(\log N)$ communication and rounds where N is the total number of multiplications verified.

Protocol	Preprocessing Phase		Online Phase		Total Comm.	Malicious
	Comm. Bit	Rounds	Comm.Bit	Rounds		
Obliv. select [43]	–	–	286720	30	286720	✗
$GF(2)$ -circuit [6, 43]	–	–	5120	60	5120	✗
$GF(2)$ -circuit [5]	–	–	35840	60	35840	✓
$GF(2)$ -circuit [44]	–	–	5120*	60*	5120	✓
$GF(2^8)$ -circuit [21]	–	–	5120	30	5120	✗
$GF(2^8)$ -circuit [21, 31]	≈ 15000	$O(1)$	6144	30	≈ 21144	✓
$GF(2^8)$ -circuit [15, 21]	–	–	5120*	30*	5120	✓
Garbled circuit [45, 50]	614400	1	0 [†]	1	≈ 614400	✓
$GF(2^4)$ -circuit	–	–	3200*	30*	3200	✓
$GF(2^4)$ -LUT-16	1760	2	2560*	20*	4320	✓
$\langle\langle\cdot\rangle\rangle$ -LUT-256	3520	2	2560*	10*	6080	✓
$[\![\cdot]\!]$ -LUT-256	39520	6	1280*	10*	39800	✓

sharing ($[\![\cdot]\!]$). Replicated secret sharing has the lowest online communication cost, but has very expensive preprocessing. With additive secret sharing, through a novel approach, we are able to reduce the preprocessing cost by more than 10x, whilst only doubling the communication in the online phase. Based on our implementation, it seems that the LUT-256 protocols are best suited to a WAN setting, where round complexity is more critical, especially since the local computation cost of the table lookups is larger.

Achieving Malicious Security. One of the main challenges in our protocols is to achieve malicious security with a low overhead. One reason this is difficult is that our protocols rely on additive secret sharing at various points, instead of purely robust schemes like replicated or Shamir secret sharing. This makes it hard to apply standard verification techniques, such as the batch multiplication procedure of Boneh et al. [11], which is often used for distributed zero-knowledge proofs and MPC protocols [15, 32, 44]. To overcome this, we carefully design our protocols such that the necessary replicated shares can be extracted from our additively shared table lookup protocols. This allows the result of a table lookup on additively shared inputs to be cheaply verified by checking the *previous* multiplication gate. Additionally, for our additively shared LUT-256 protocol, we rely on the algebraic structure of the S-box to reduce the cost of the correctness check of an S-box computation, after a potentially faulty table lookup, inspired by recent work in zero-knowledge proofs [8]. This allows all of our maliciously secure protocols to have the same amortized communication cost as their semi-honest counterparts.

As a stepping stone in one of our protocols, we also obtain a general-purpose, malicious protocol for table lookups with $O(\sqrt{N})$ communication complexity for a table of size N . To

the best of our knowledge, all prior practical approaches with malicious security require a cost of $\Omega(N)$. Our protocol relies on the observation that a table lookup can be verified using a single inner product check. By applying the generalised version of the batch multiplication verification from [11], which allows for checking inner product relations, we show how to verify a $O(\sqrt{N})$ complexity lookup table with almost no communication overhead.

Interestingly, for our additively shared AES protocol based on a specialization of this technique, we observe that it’s not sufficient to directly use the verification protocol from [11, 15, 32], which inherently leak any errors in multiplication (or inner product) triples being verified to the adversary. While this leakage would typically be harmless, since the errors are already known to the adversary, this turns out not to be the case for us (for further discussion, see Section 3.5.2). We therefore show how to modify the verification procedure to remove any leakage.

Overall, as can be seen in Table 1, our maliciously secure protocols lead to a large reduction in communication costs compared with prior approaches. Our implementation results show that this approach comes with a slight increase in computational costs, but is still highly practical.

1.3 Related Work

Background. Known oblivious AES protocols are classified into two primary categories: those utilizing garbled circuits and those employing secret sharing schemes. Garbled circuit-based approaches [35, 37, 40, 50, 54] have the advantage of fewer communication rounds, but have high bandwidth costs due to the extensive size of garbled circuits. On the other hand, secret sharing schemes require more communication

rounds, but less communication. In recent years, efficient methods for performing secure Boolean operations have been proposed [6, 21] in the honest-majority setting, improving the performance of AES in MPC. The main challenge in constructing oblivious AES protocols lies in computing S-boxes, which requires non-linear operations [6] and thus communication. Previous research addressing this challenge can be broadly categorized into two types: methods using secure table lookup protocols [42, 43], and secure computation of algorithms that focus on the specific structure of S-boxes [23, 43]. Our proposed protocol integrates good aspects from both of these methods.

Oblivious AES using Lookup Tables. Oblivious AES computation using secure table lookups was proposed by Launchbury et al. [42] and Laur et al. [43], by converting the secret lookup index x into a one-hot vector encoding, with a one in position x and zeroes elsewhere. This approach performs the encoding entirely in the online phase and requires 304 secure multiplications and 3 rounds to process one S-box as a size-256 table. Later works have used this technique in both the dishonest majority and honest majority settings [7, 18, 25, 40, 48], with the main improvement being to offload the computation of the one-hot vector to a preprocessing phase, leading to a very lightweight online phase. For example, the protocols from [18, 40] require $2^k - k - 1$ secure AND gates to preprocess a table of size $N = 2^k$. Our $\llbracket \cdot \rrbracket$ -LUT-256 protocol is based on these ideas applied to the setting of replicated secret sharing. Other approaches using distributed point functions can reduce the bandwidth cost to $O(k\lambda)$, for security parameter λ , but this comes with computational security and an expensive setup phase [13, 14].

Structure of the S-box. Protocols exploiting the structure of the AES S-box have proposed various ways to improve efficiency. Laur et al. [43] securely compute S-boxes with the optimized Boolean circuit of Boyar et al. [12], obtaining a protocol in 6 rounds and with 32 AND gates. Subsequent works [5, 6] employ the same technique in the replicated secret-sharing setting for semi-honest and malicious adversaries. By focusing on the algebraic structure of the S-box, methods like the one in [21, 23, 24, 40] securely compute the S-box as a multiplicative inverse $x^{-1} = x^{254}$ in $GF(2^8)$. The previous most efficient AES protocols compute the inversion as a circuit in $GF(2^8)$ [21, 40]. In Chida et al.’s protocol [21], this inversion can be performed with only four multiplications in three rounds. While it is also possible to interpolate the AES S-box as a sparse polynomial in $GF(2^8)$ [46], as explored in [24], the cost of 18 multiplications in 12 rounds is prohibitively high compared to other techniques.

Technique for Multiplicative Inverse. The technique we use in this paper to calculate the multiplicative inverse using an extension of $GF(2^4)$ has mainly been used in hardware implementations of AES [52, 53]. Garbled circuit-based methods for oblivious AES have used optimized circuits based on

this approach [37].

Maliciously Secure Protocols. There are also oblivious AES protocols that are secure against malicious adversaries [5, 24, 26, 27]. Our maliciously secure protocols target the three-party honest-majority setting. In the same setting, [5], improving over [31], use bucket cut-and-choose techniques to compute Boolean circuits with a total communication cost of 7 bits per AND gate. Afterwards, batched multiplication checks with logarithmic overhead have been employed to reduce communication [15, 44] where each party proves correct behaviour separately. Extending this approach, in all of our protocols, the parties *jointly* prove correctness of the multiplications. This leads to a factor 3 improvement in local computation which is the bottleneck when employed over medium to high bandwidth networks. Furthermore, we leverage the AES circuit-specific relations between multiplications and encode multiple multiplication triples to be checked into a single triple in the larger field which is needed for soundness. The three-party garbling framework by Mohassel et al. [45] lifts any semi-honest two-party garbling scheme into a malicious three-party protocol in the honest-majority setting. Instantiated with the ThreeHalves scheme [50], the communication cost per AND gate is $1.5\kappa \approx 120$ bits for 80-bit security.

In other settings, maliciously secure two-party protocols [26] and multi-party protocols [23, 24, 27] for dishonest majority implement the AES function. These protocols require more than five times the communication cost compared to semi-honest secure protocols.

2 Preliminaries

We begin by outlining some notation. We write \vec{x} to denote vectors and index them as \vec{x}_i . We write $\vec{x} \cdot \vec{y}$ to describe the inner product. We use $\vec{z} = \vec{x} \parallel \vec{y}$ to denote concatenation, i.e., \vec{z} first contains elements of \vec{x} , then of \vec{y} . One-hot vectors are written as $e^{(r)}$ where r is the index of the single one in the vector, i.e., $e_r^{(r)} = 1$ and $e_i^{(r)} = 0$ for $i \neq r$. Public truth table vectors are denoted with T , omitting the $\vec{\cdot}$.

2.1 Finite Fields and Field Inversion

Let \mathbb{F}_2 be the field of order 2. We define the finite fields $GF(2^8)$ and $GF(2^4)$ as follows:

$$\begin{aligned} GF(2^8) &:= \mathbb{F}_2[X]/(X^8 + X^4 + X^3 + X + 1), \\ GF(2^4) &:= \mathbb{F}_2[X]/(X^4 + X + 1). \end{aligned}$$

Note that the irreducible polynomial used in the definition of $GF(2^8)$ comes from the AES specification [28]. We consider elements of $GF(2^8)$ as bit-strings of length 8 corresponding to the coefficients of a degree-7 polynomial over \mathbb{F}_2 . In particular, we write $\sum_{i=0}^7 a_i x^i \in GF(2^8)$ as the string

$\{a_7a_6 \dots a_1a_0\}_2$ or the corresponding 2-digit hexadecimal notation. For example, $\{6E\}_{16}$ represents the equivalence class of $X^6 + X^5 + X^3 + X^2 + X$. Similarly, elements of $GF(2^4)$ are represented as 4-bit sequences or 1-digit hexadecimal notation. For example, $\{E\}_{16} \in GF(2^4)$ represents the equivalence class of $X^3 + X^2 + X$.

We define the finite field (extension) $GF((2^4)^2)$ over $GF(2^4)$ as $GF((2^4)^2) := GF(2^4)[X]/(X^2 + X + \{E\}_{16})$. Elements of $GF((2^4)^2)$ are represented as a degree-1 polynomial $a_hX + a_\ell$ over $GF(2^4)$. Each coefficient's binary representation is $\{a_{h3}a_{h2}a_{h1}a_{h0}\}_2$ and $\{a_{\ell3}a_{\ell2}a_{\ell1}a_{\ell0}\}_2$.

Multiplicative Inversion in $GF((2^4)^2)$. A formula for the inverse of $a = a_hX + a_\ell \in GF((2^4)^2)$ can be calculated by solving a system of equations derived from $aa^{-1} = 1$ (see, for instance, [20, 29]), giving

$$(a_hX + a_\ell)^{-1} = (a_h \otimes v^{-1})X + (a_h \oplus a_\ell) \otimes v^{-1}. \quad (1)$$

Here, \oplus, \otimes represent addition and multiplication in $GF(2^4)$ respectively, and $v \in GF(2^4)$ is defined as follows

$$v := (a_h^2 \otimes \{E\}_{16}) \oplus (a_h \otimes a_\ell) \oplus a_\ell^2. \quad (2)$$

This shows that the inverse in $GF((2^4)^2)$ can be obtained through the calculation of one inverse v^{-1} in $GF(2^4)$, plus three multiplications and two squarings in $GF(2^4)$.

Isomorphism Between $GF(2^8)$ and $GF((2^4)^2)$. The finite fields $GF(2^8)$ and $GF((2^4)^2)$ are isomorphic. We use the explicit isomorphism and its inverse, described by Wolkerstorfer, Oswald and Lamberger [53], given by the following maps.

$$\begin{aligned} \Phi: GF(2^8) &\xrightarrow{\sim} GF((2^4)^2) : \{a_7a_6 \dots a_0\}_2 \mapsto (a_h, a_\ell), \\ a_{h0} &:= a_4 \oplus a_5 \oplus a_6, & a_{\ell0} &:= a_0 \oplus a_4 \oplus a_5 \oplus a_6, \\ a_{h1} &:= a_1 \oplus a_4 \oplus a_6 \oplus a_7, & a_{\ell1} &:= a_1 \oplus a_2, \\ a_{h2} &:= a_2 \oplus a_3 \oplus a_5 \oplus a_7, & a_{\ell2} &:= a_1 \oplus a_7, \\ a_{h3} &:= a_5 \oplus a_7, & a_{\ell3} &:= a_2 \oplus a_4. \end{aligned}$$

and

$$\Phi^{-1}: GF((2^4)^2) \xrightarrow{\sim} GF(2^8) : (a_h, a_\ell) \mapsto \{a_7a_6 \dots a_0\}_2.$$

$$\begin{aligned} a_0 &:= a_{\ell0} \oplus a_{h0}, & a_4 &:= a_{\ell1} \oplus a_{\ell3} \oplus a_{h0} \oplus a_{h1} \oplus a_{h3}, \\ a_1 &:= a_{h0} \oplus a_{h1} \oplus a_{h3}, & a_5 &:= a_{\ell2} \oplus a_{h0} \oplus a_{h1}, \\ a_2 &:= a_{\ell1} \oplus a_{h0} \oplus a_{h1} \oplus a_{h3}, & a_6 &:= a_{\ell1} \oplus a_{\ell2} \oplus a_{\ell3} \oplus a_{h0} \oplus a_{h3}, \\ a_3 &:= a_{h0} \oplus a_{h1} \oplus a_{h2} \oplus a_{\ell1}, & a_7 &:= a_{\ell2} \oplus a_{h0} \oplus a_{h1} \oplus a_{h3}. \end{aligned}$$

2.2 Advanced Encryption Standard (AES)

AES is a block cipher standardized by NIST [28], with a 128-bit block size. We focus on AES-128, with a key length of 128 bits. For a detailed overview of the algorithm, see Appendix C.

High-Level Structure and S-box. AES follows a substitution-permutation network design, with alternating layers of non-linear S-boxes and linear permutations. In addition, there is a key schedule that expands the 128-bit key into a set of round keys to be used in each round. The linear components of a round are ShiftRows, MixColumns and AddRoundKey, which XORs the state with the round key. These can be expressed as linear operations in $GF(2^8)$.

The S-box of AES — also called SubBytes — operates on one byte of the state at a time, and is the only non-linear part of AES. It can be expressed as the mapping over the finite field $GF(2^8)$ that sends $x \mapsto \text{Affine}(x^{254})$, where Affine is an invertible affine transformation. Since the multiplicative group of $GF(2^8)$ has order 255, the computation of x^{254} maps every non-zero x to x^{-1} and 0 to 0. We will often abuse terminology slightly and refer to this as an inversion in $GF(2^8)$.

2.3 Security Model

We consider protocols secure against up to $t - 1$ out of n corrupted parties, where $t - 1 < n/2$. We will often focus on the case of 1-out-of-3 corruptions, where $n = 3, t = 2$. All of our protocols are presented and analyzed in the malicious model with abort, where a corrupt party may deviate from the protocol specification and honest parties are not guaranteed to receive output. We also consider relaxations in the semi-honest model, where each party is assumed to follow the protocol, which are obtained by omitting any verification steps in the protocol.

We model ideal functionalities and give security proofs in the Universal Composability (UC) framework [19], which gives strong composition guarantees.

2.4 Secure Multi-Party Computation

We build upon MPC protocols based on replicated secret sharing [38]. This approach is well-suited for secure computations involving $GF(2^k)$ values such as those occurring in oblivious AES. Additionally, it offers the advantage of a lightweight protocol for multiplications of secret-shared values.

Replicated Secret Sharing and Additive Secret Sharing. We use both t -out-of- n replicated secret sharing (t, n) -RSS and n -out-of- n additive secret sharing (n, n) -SS.

In (t, n) -RSS, we write $\llbracket a \rrbracket$. For a secret a in a finite field \mathbb{F} , a is split into $\binom{n}{t-1}$ random shares $a^{(T)} \in \mathbb{F}$, for each subset $T \subset [n]$ of size $t - 1$, with the shares sampled so that $a = \sum_T a^{(T)}$. Party i gets the $\binom{n-1}{t-1}$ shares $\llbracket a \rrbracket_i = \{a^{(T)}\}_{T, i \notin T}$. The overall sharing of a is $\llbracket a \rrbracket = (\llbracket a \rrbracket_1, \dots, \llbracket a \rrbracket_n)$. In case of three parties, each party holds exactly two of the three shares.

In (n, n) -SS, we write $\langle\langle a \rangle\rangle$. For a secret $a \in \mathbb{F}$, each party i holds $\langle\langle a \rangle\rangle_i = a^{(i)} \in \mathbb{F}$ s.t. $a = a^{(1)} + \dots + a^{(n)}$. A share of a value a is $\langle\langle a \rangle\rangle = (\langle\langle a \rangle\rangle_1, \dots, \langle\langle a \rangle\rangle_n)$.

In this paper, the field \mathbb{F} is always characteristic 2, and typically either $GF(2)$, $GF(2^4)$, or $GF(2^8)$. With both replicated

and additive sharing, a sharing in $GF(2^k)$ can be decomposed into an array of k shares over $GF(2)$ without communication by doing binary decomposition.

We denote the set of corrupted parties by \mathcal{C} , and the set of honest parties by \mathcal{H} . For a sharing $\llbracket x \rrbracket$, we let $\llbracket x \rrbracket^{\mathcal{C}} = \{\llbracket x \rrbracket_i\}_{i \in \mathcal{C}}$ be the set of all corrupt parties' shares, and $\llbracket x \rrbracket^{\mathcal{H}}$ the set of all honest shares (and similarly define $\langle\langle \cdot \rangle\rangle^{\mathcal{C}}$, $\langle\langle \cdot \rangle\rangle^{\mathcal{H}}$ for additive sharings).

Definition 2.1 (Consistent shares) Let $S \subset [n]$ and consider a set of replicated shares $\{\llbracket x \rrbracket_i\}_{i \in S}$, where $\llbracket x \rrbracket_i = \{x_i^{(T)}\}_{T \neq i}$. We say that the set of shares is consistent if $x_i^{(T)} = x_{i'}^{(T)}$ for every i, i' and T where $i, i' \notin T$.

When modelling security, we often define ideal functionalities where the adversary provides as input a set of shares $\llbracket x \rrbracket^{\mathcal{C}}$. In this case, we implicitly require that the functionality only accepts a set of consistent shares. Our functionalities also often rely on the following straightforward fact.

Proposition 2.2 Given a secret x and a consistent set of corrupted parties' shares $\llbracket x \rrbracket^{\mathcal{C}}$, a consistent set of honest shares $\llbracket x \rrbracket^{\mathcal{H}}$ can always be defined.

Note that if exactly t parties are corrupted, then the honest parties' shares are defined uniquely; otherwise, they can be sampled at random.

We consider reconstruction as an interactive protocol, where parties exchange shares and reconstruct a secret. We have the following two protocols.

Opening Additive Shares: $x = \text{Reconst}(\langle\langle x \rangle\rangle)$. Each party sends its share $x^{(i)}$ to all other parties, and reconstructs $x = \sum x^{(i)}$. This requires $n - 1$ field elements of communication per party in one round. Alternatively, one can use the "king" approach, where the parties send their shares to a designated party, who reconstructs and sends back x . This takes on average only $2(n - 1)/n$ field elements per party but two rounds of interaction. In our implementation, since we focus on the 3-party setting, we chose to take the one-round approach with two elements of communication. Note that in the malicious setting, a corrupted party can easily change the result of reconstruction by lying about their share.

Opening Replicated Shares: $x = \text{Reconst}(\llbracket x \rrbracket)$. With replicated secret sharing, the parties can *robustly* open a secret, guaranteeing that each party either outputs the correct value or aborts. The simplest protocol is for the parties to exchange all of their shares, and check whether the resulting sharing is consistent before reconstructing x . When reconstructing many values, this protocol can be optimized by optimistically sending the minimal number of shares needed to reconstruct x , and later verifying all openings in a batch by exchanging and comparing hashes of the remaining shares. This was demonstrated for the 3-party setting in [31] and later extended to the multi-party setting [41].

<p>Functionality $\mathcal{F}_{\text{rand}}(\mathbb{F})$</p> <ol style="list-style-type: none"> 1. $\mathcal{F}_{\text{rand}}$ receives from the adversary the shares $\llbracket r \rrbracket^{\mathcal{C}}$. 2. $\mathcal{F}_{\text{rand}}$ samples $r \leftarrow \mathbb{F}$, and uses $(r, \llbracket r \rrbracket^{\mathcal{C}})$ to define $\llbracket r \rrbracket^{\mathcal{H}}$. 3. $\mathcal{F}_{\text{rand}}$ distributes the shares $\llbracket r \rrbracket^{\mathcal{H}}$ to the honest parties.
<p>Functionality $\mathcal{F}_{\text{zero}}(\mathbb{F})$</p> <ol style="list-style-type: none"> 1. $\mathcal{F}_{\text{zero}}$ receives from the adversary the shares $\langle\langle z \rangle\rangle^{\mathcal{C}}$, and samples random shares $\langle\langle z \rangle\rangle^{\mathcal{H}}$, such that $z = 0 \in \mathbb{F}$. 2. $\mathcal{F}_{\text{zero}}$ distributes the shares $\langle\langle z \rangle\rangle^{\mathcal{H}}$ to the honest parties.
<p>Functionality $\mathcal{F}_{\text{coin}}(k)$</p> <ol style="list-style-type: none"> 1. $\mathcal{F}_{\text{coin}}$ samples $(b_1, \dots, b_k) \leftarrow \{0, 1\}^k$ and sends this to the adversary. 2. On receiving <i>OK</i> from the adversary, it delivers (b_1, \dots, b_k) to the honest parties.

Figure 1: The functionalities $\mathcal{F}_{\text{rand}}$, $\mathcal{F}_{\text{zero}}$ and $\mathcal{F}_{\text{coin}}$.

2.5 Correlated Randomness and Coin Tossing

We require the parties to have access to different forms of correlated randomness, for obtaining replicated sharings of random values, and additive sharings of zero (see Fig. 1). The $\mathcal{F}_{\text{rand}}$ functionality can be implemented using pseudorandom secret sharing [22], where after a one-time setup to distribute pseudorandom function keys among the parties, the correlated randomness can be generated non-interactively. Obtaining random additive sharings of zero, modelled in $\mathcal{F}_{\text{zero}}$, can similarly be done using pseudorandom secret sharing. One simple approach is to use $\mathcal{F}_{\text{rand}}$ to obtain a random $\llbracket r \rrbracket$, and then each party XORs together $n - 1$ of its share elements, appropriately selected such that all shares cancel out and sum to zero.

Finally, we also rely on the coin-tossing functionality, $\mathcal{F}_{\text{coin}}$, which can be realized by running Reconst on a random sharing from $\mathcal{F}_{\text{rand}}$.

2.6 Computations on Shares

Linear Operations. Any $GF(2)$ -linear operation can be performed locally on replicated or additively shared values, by simply applying the operation to each element of each party's share. Similarly, addition by a constant can be performed by adding it to a fixed subset of the shared elements. Given sharings $\llbracket x \rrbracket$, $\llbracket y \rrbracket$ and public values a, c , we denote these operations by $\llbracket ax + y + c \rrbracket := a\llbracket x \rrbracket + \llbracket y \rrbracket + c$ and similarly $\langle\langle ax + y + c \rangle\rangle := a\langle\langle x \rangle\rangle + \langle\langle y \rangle\rangle + c$.

Free Squaring. Since squaring in $GF(2^k)$ is linear over $GF(2)$, it can be performed locally in both (t, n) -RSS and (n, n) -SS: each party simply squares their corresponding shares. We denote $\llbracket x^2 \rrbracket := \text{Square}(\llbracket x \rrbracket)$.

Local Multiplication. We rely on the multiplicative property of replicated secret sharing: given sharings $\llbracket x \rrbracket$ and $\llbracket y \rrbracket$, the parties can obtain an additive sharing $\langle\langle xy \rangle\rangle$ with-

Functionality $\mathcal{F}_{\text{weakMult}} / \mathcal{F}_{\text{weakDotProduct}}$
Input: $[\![\vec{x}]\!], [\![\vec{y}]\!]: (t, n)$ -RSS share of \vec{x} and \vec{y} .
Output: $[\![\vec{x} \cdot \vec{y}]\!]: (t, n)$ -RSS share of the product $\vec{x} \cdot \vec{y}$.
1. $\mathcal{F}_{\text{weakMult}}$ receives the shares $[\![\vec{x}]\!]^{\mathcal{H}}, [\![\vec{y}]\!]^{\mathcal{H}}$ from honest parties and reconstructs the secrets \vec{x}, \vec{y} .
2. $\mathcal{F}_{\text{weakMult}}$ computes the corrupted parties' shares $[\![\vec{x}]\!]^{\mathcal{C}}, [\![\vec{y}]\!]^{\mathcal{C}}$ and sends these shares to the adversary.
3. $\mathcal{F}_{\text{weakMult}}$ receives an error d and a set of shares $[\![z]\!]^{\mathcal{C}}$ from the adversary.
4. $\mathcal{F}_{\text{weakMult}}$ computes $z := \vec{x} \cdot \vec{y} + d$ and samples honest parties' shares $[\![z]\!]^{\mathcal{H}}$ by using d and $[\![z]\!]^{\mathcal{C}}$.
5. $\mathcal{F}_{\text{weakMult}}$ sends the shares $[\![z]\!]^{\mathcal{H}}$ to honest parties.

Figure 2: Functionality for multiplication with additive error. We refer to the functionality as $\mathcal{F}_{\text{weakMult}}$ if \vec{x}, \vec{y} have length 1, and $\mathcal{F}_{\text{weakDotProduct}}$ otherwise.

out any interaction. This holds because when $t - 1 < n/2$, any pair of share elements $x^{(T)}, x^{(T')}$, for size- $(t - 1)$ subsets T, T' , is held by at least one party. For example, when $n = 3$, party i holds $(x^{(i)}, x^{(i+1)}), (y^{(i)}, y^{(i+1)})$ and can compute $\langle\langle xy \rangle\rangle_i := x^{(i)}y^{(i)} + (x^{(i)} + x^{(i+1)})(y^{(i)} + y^{(i+1)})$. We write this multiplication as $\langle\langle xy \rangle\rangle := \text{MultLocal}([\![x]\!], [\![y]\!])$.

Share Conversion from $\langle\langle x \rangle\rangle$ to $[\![x]\!]$. This can be achieved with a single communication round. Essentially, each party creates a replicated secret sharing $[\![x^{(i)}]\!]$ of its additive share, and distributes the resulting shares to the remaining parties. With $n = 3$, this can be achieved by having the parties first re-randomize their shares by adding a share of zero from $\mathcal{F}_{\text{zero}}$. Then, party i sends its share $x^{(i)}$ to party $i + 1$. The cost is sending 1 field element per party in one round [6]. We write this as $[\![x]\!] = \text{Reshare}(\langle\langle x \rangle\rangle)$.

Multiplication with Additive Errors. Combining the local multiplication and resharing protocols above, we obtain a multiplication protocol on $[\![\cdot]\!]$ -shared values. In the malicious setting, a corrupted party may cheat during the resharing step, introducing an error into the output. We model this in the $\mathcal{F}_{\text{weakMult}}$ functionality (Fig. 2), adapted from [33], which allows the adversary to choose an error d that is added into the output shares. We additionally extend this to $\mathcal{F}_{\text{weakDotProduct}}$, for computing an inner product, where the inner product is computed locally followed by a single resharing.

2.7 Verifying Multiplications and Dot Products with Malicious Security

To ensure correct multiplications with malicious security, we use a batch verification procedure. The idea is that during the main MPC execution, the parties can use $\mathcal{F}_{\text{weakMult}}$, and later verify that all multiplications were correct. This batch verification can be safely postponed until the end of the protocol, as long as any outputs of the computation are only revealed after the multiplications have been verified.

Protocol 1 CheckTriple($[\![x]\!], [\![y]\!], [\![z]\!]$) $\rightarrow 0/1$

Input: $[\![x]\!], [\![y]\!], [\![z]\!]$.

Output: Verify that $xy = z$.

- 1: $[\![x']]\!, [\![r]\!] \leftarrow \mathcal{F}_{\text{rand}}(\mathbb{F})$
 - 2: $[\![z']]\! \leftarrow \mathcal{F}_{\text{weakMult}}([\![x']]\!, [\![y]\!])$
 - 3: $t \leftarrow \mathcal{F}_{\text{coin}}(\mathbb{F})$
 - 4: $\rho := \text{Reconst}([\![x]\!] + t[\![x']]\!)$
 - 5: $[\![\sigma]\!] := \mathcal{F}_{\text{weakMult}}([\![z]\!] + t[\![z']]\! - \rho[\![y]\!], [\![r]\!])$
 - 6: **if** $\text{Reconst}([\![\sigma]\!]) = 0$ **then**
 - 7: **return** 1
 - 8: **else**
 - 9: **return** 0
 - 10: **end if**
-

Some prior works implementing MPC for Boolean circuits, such as [5, 31], use cut-and-choose techniques to verify multiplications. These have a large communication overhead, and require running in very large batches to obtain reasonable parameters. Instead, we verify multiplication triples by adapting the protocol of [32, 33], originally presented for Shamir-based MPC, and based on similar ideas used previously for distributed zero-knowledge proofs [11] and MPC [15, 16, 47]. The work of Li et al. [44] uses similar techniques to check multiplications in Boolean circuits but their encoding of bit triples into prime field triples does not efficiently generalize for binary extension fields that our protocols require.

Our protocol realizes the functionality $\mathcal{F}_{\text{verify}}$ (Fig. 3). The main protocol uses a recursive inner product check, which is essentially that of [16, 32] adapted to more general secret-sharing schemes. However, we make two key changes that are needed in some of our applications. Firstly, we extend the protocol to verifying not just multiplications, but also inner product relations using the approach from [11]. Secondly, prior works [16, 32] only realized a functionality that leaks the errors $z_i - x_i y_i$ in all multiplication triples to the adversary. Instead, we modify the base case of the protocol to realize a stronger functionality, which *only* leaks the result of the verification check and no additional information. In typical usage, the errors in multiplications are chosen by the adversary so leaking them in $\mathcal{F}_{\text{verify}}$ would not be an issue. However, specifically for our AES protocol in Section 3.5.3, it turns out that leaking these errors would compromise security.

The main protocol is shown in Protocol 2. To ensure correct triples with high probability, we need to work over an exponentially large finite field \mathbb{F} . We therefore use the protocol by first taking our Boolean (or small field) triples and lifting the shares into a large extension field. Then, the protocol begins by randomizing the batch of inner product triples, converting it into a single, large inner product of length N . To verify the inner product, the protocol proceeds in $\log N$ rounds, where in each round the dimension is halved, by first viewing the inner product as an inner product on length- $N/2$ vectors of suitably defined degree-1 polynomials, followed

Protocol 2 Verifying a batch of inner products

Functionality: $0/1 \leftarrow \mathcal{F}_{\text{verify}}(\llbracket \cdot \cdot \cdot \rrbracket)$.

Input: Shared triples $\{\llbracket \vec{x}_i \rrbracket, \llbracket \vec{y}_i \rrbracket, \llbracket z_i \rrbracket\}_{i=0}^{m-1}$, where $\vec{x}_i, \vec{y}_i \in \mathbb{F}^{n_i}$ and $z_i \in \mathbb{F}$ for each i .

Output: Verify that $\vec{x}_i \cdot \vec{y}_i = z_i$, for all i .

```

1:  $r \leftarrow \mathcal{F}_{\text{coin}}(\mathbb{F})$ 
2: return VerifyDotProduct( $\llbracket \vec{x}_0 \rrbracket r \vec{x}_1 \parallel \dots \parallel r^{m-1} \vec{x}_{m-1} \rrbracket$ ,
 $\llbracket \vec{y}_0 \rrbracket \dots \parallel \vec{y}_{m-1} \rrbracket, \sum_{i=0}^{m-1} r^i \llbracket z_i \rrbracket$ )

3: procedure VerifyDotProduct( $(\llbracket x_0 \rrbracket, \dots, \llbracket x_{N-1} \rrbracket), (\llbracket y_0 \rrbracket, \dots, \llbracket y_{N-1} \rrbracket), \llbracket z \rrbracket$ )
4:   if  $N = 1$  then
5:     return CheckTriple( $\llbracket x_0 \rrbracket, \llbracket y_0 \rrbracket, \llbracket z \rrbracket$ )
6:   end if
7:    $\triangleright$ polynomials in  $X$ :  $f_i(b) = x_{2i+b}, g_i(b) = y_{2i+b}$ , for  $b \in \{0, 1\}$ 
8:   for  $i = 0$  to  $N/2 - 1$  do
9:      $\llbracket f_i(X) \rrbracket := \llbracket x_{2i} \rrbracket + (\llbracket x_{2i} \rrbracket + \llbracket x_{2i+1} \rrbracket)X$ 
10:     $\llbracket g_i(X) \rrbracket := \llbracket y_{2i} \rrbracket + (\llbracket y_{2i} \rrbracket + \llbracket y_{2i+1} \rrbracket)X$ 
11:   end for
12:    $\llbracket h(1) \rrbracket \leftarrow \mathcal{F}_{\text{weakDotProduct}}(\llbracket \vec{f}(1) \rrbracket, \llbracket \vec{g}(1) \rrbracket)$ 
13:    $\llbracket h(2) \rrbracket \leftarrow \mathcal{F}_{\text{weakDotProduct}}(\llbracket \vec{f}(2) \rrbracket, \llbracket \vec{g}(2) \rrbracket)$ 
14:    $\llbracket h(0) \rrbracket := \llbracket z \rrbracket - \llbracket h(1) \rrbracket$   $\triangleright$ defines degree-2  $h(X)$ 
15:    $r \leftarrow \mathcal{F}_{\text{coin}}(\mathbb{F})$ 
16:   Locally compute  $\llbracket h(r) \rrbracket$  via Lagrange interpolation
17:   return VerifyDotProduct( $\llbracket \vec{f}(r) \rrbracket, \llbracket \vec{g}(r) \rrbracket, \llbracket h(r) \rrbracket$ )
18: end procedure

```

by evaluating the polynomials at a random challenge to compress this to an inner product of vectors of field elements. Eventually, it reaches a base case where $N = 1$, and performs a naive triple check (Protocol 1) that uses one extra, random multiplication and a random challenge to check the remaining one. Importantly, step 5 of the protocol computes the value $z + tz' - ry = (z - xy) + t(z' - x'y)$, which should equal 0 if the triple is correct. However, this cannot be revealed directly, as it would leak information on $z - xy$ to the adversary. To prevent this, we re-randomize it via the additional multiplication with $\llbracket r \rrbracket$; this change allows us to realize the stronger $\mathcal{F}_{\text{verify}}$ functionality.

We prove the following in Appendix D.

Theorem 2.3 *Protocol 2 (Verify) securely realizes the functionality $\mathcal{F}_{\text{verify}}$ (see Fig. 3), in the $(\mathcal{F}_{\text{weakMult}}, \mathcal{F}_{\text{coin}}, \mathcal{F}_{\text{rand}})$ -hybrid model. The failure probability in the simulation is at most $(m + 2 \log N)/|\mathbb{F}|$.*

In the above bound, m denotes the number of inner product triples of length n_i , $1 \leq i \leq m$, and $N = \sum_{i=1}^m n_i$. In order to obtain statistical security ρ , the size of the field in Protocol 2 must be $\geq \rho$ bit. Thus the computation phase of all our MPC protocols uses the small fields $GF(2^4)$, $GF(2^8)$ but the *multiplication check phase* embeds the observed small field triples into triples in $GF(2^{64})$ to achieve the required soundness level.

Table 2: Embeddings used to convert smaller field triples into $GF(2^{64})$ elements for $\mathcal{F}_{\text{verify}}$.

Field	
$GF(2^4)$	$\psi(X) = \{\text{a181e7d66f5ff794}\}_{16}$
$GF(2^8)$	$\psi(X) = \{\text{033ce8beddc8a656}\}_{16}$
$(GF(2^4))^4$	$\psi(\alpha) = \{\text{14f1968d182dd50f}\}_{16}$

Details on Embeddings. Since the inputs to $\mathcal{F}_{\text{verify}}$ are $GF(2^{64})$ elements, the parties locally convert their shares of each inner product triple $\llbracket \vec{x} \rrbracket, \llbracket \vec{y} \rrbracket, \llbracket z \rrbracket$ with $x_i, y_i, z \in GF(2^k)$ to an inner product triple in $GF(2^{64})$ by computing the isomorphism $\psi: GF(2^k) \xrightarrow{\sim} G$ where G is a subfield of $GF(2^{64})$ of size 2^k (where $k = 2^4$ or $k = 2^8$). Due to the linearity of ψ , $\psi(\llbracket x_i \rrbracket) = \llbracket \psi(x_i) \rrbracket$, all this computation is local and secure due to the $\llbracket \cdot \rrbracket$ -sharing. For correctness, note that the check $\sum \psi(x_i) \psi(y_i) = \psi(z + \delta) = \psi(z) + \psi(\delta) \in GF(2^{64})$ with the error $\delta \in GF(2^k)$ (in the small field) holds iff. $\delta = 0$ since the bijective ψ only maps $0 \in GF(2^k)$ to $0 \in GF(2^{64})$.

In Table 2 we detail the concrete isomorphisms used in our protocols. We let $GF(2^{64}) := \mathbb{F}_2[Y]/(Y^{64} + Y^4 + Y^3 + Y + 1)$ and write its elements as 64-bit integers in hexadecimal notation in little endian order, i.e., $\{141\}_{16}$ denotes $Y^9 + Y^7 + 1$. Further, we define the degree- k extension of $GF(2^4)$ (see Protocol 3) for the check with a reduced number of multiplications as $(GF(2^4))^4 := GF(2^4)[\alpha]/\{2\}_{16}\alpha^4 + \{2\}_{16}\alpha^2 + \{4\}_{16}\alpha + \{8\}_{16}$. Finally, the embedding used for the improved multiplication check in Protocol 8 embeds $(b_0, \dots, b_7) \in \mathbb{F}_2^8$ into $GF(2^{64})$ as $b_0 + b_1Y + \dots + b_7Y^7$. Note that this is not an isomorphism to a subfield of $GF(2^{64})$ of size 2^8 , however the check Eq. 5 and 6 only requires a single multiplication where this naive embedding is still correct. Importantly, the adversary can only add bit errors on the right hand side of the equations, thus cannot cancel other errors by triggering a reduction modulo $(Y^{64} + Y^4 + Y^3 + Y + 1)$.

2.8 Functionality for Table Lookup

In Fig. 4, we present a functionality for performing a secret-shared table lookup. The functionality is parameterized by the secret-sharing schemes used for the inputs and outputs, denoted ss_1 and ss_2 respectively, which can be any combination of replicated shares ($\llbracket \cdot \rrbracket$) and additive shares ($\langle \langle \cdot \rangle \rangle$). We shorten the description to $\mathcal{F}_{\text{LUT}}^{ss}$ if $ss_1 = ss_2$. Note that the variant $\llbracket \cdot \rrbracket \mapsto \llbracket \cdot \rrbracket$ is fully maliciously secure, while variants where the input or output is $\langle \langle \cdot \rangle \rangle$ -shared inherently allow a corrupt party to change the input or output. Importantly, if the input is $\langle \langle \cdot \rangle \rangle$ -shared then the functionality additionally outputs a $\llbracket \cdot \rrbracket$ sharing of the input that was used. We use this in our maliciously secure protocols for verifying the correct inputs were used after the protocol execution.

In Section 3.3, we will describe a protocol for realizing

Functionality $\mathcal{F}_{\text{verify}}$
Input: $(\llbracket \vec{x}^{(1)} \rrbracket, \llbracket \vec{y}^{(1)} \rrbracket, \llbracket \vec{z}^{(1)} \rrbracket), \dots, (\llbracket \vec{x}^{(m)} \rrbracket, \llbracket \vec{y}^{(m)} \rrbracket, \llbracket \vec{z}^{(m)} \rrbracket)$, where m is the number of inner product triples to be verified. Output: $b \in \{\text{accept}(1), \text{abort}(0)\}$ to honest parties.
<ol style="list-style-type: none"> 1. $\mathcal{F}_{\text{verify}}$ receives from honest parties their shares of $(\llbracket \vec{x}^{(i)} \rrbracket, \llbracket \vec{y}^{(i)} \rrbracket, \llbracket \vec{z}^{(i)} \rrbracket)$ to reconstruct $(\vec{x}^{(i)}, \vec{y}^{(i)}, \vec{z}^{(i)})$ for all $i \in [m]$. 2. $\mathcal{F}_{\text{verify}}$ computes the rest of the corrupted parties' shares of $(\llbracket \vec{x}^{(i)} \rrbracket, \llbracket \vec{y}^{(i)} \rrbracket, \llbracket \vec{z}^{(i)} \rrbracket)$ for all $i \in [m]$ and sends these shares to the adversary. 3. $\mathcal{F}_{\text{verify}}$ sets $b := \text{abort}$ if there exists $i \in [m]$ such that $\vec{z}^{(i)} \neq \vec{x}^{(i)} \cdot \vec{y}^{(i)}$ and sets $b := \text{accept}$ otherwise. 4. $\mathcal{F}_{\text{verify}}$ sends b to the adversary and proceeds as follows: <ul style="list-style-type: none"> • If the adversary replies continue, send b to honest parties. • If the adversary replies abort, send abort to honest parties.

Figure 3: The functionality $\mathcal{F}_{\text{verify}}$.

$\mathcal{F}_{\text{LUT}}^{\langle \cdot \rangle \rightarrow \llbracket \cdot \rrbracket}$. Later, in Section 3.5.3, we give protocols for the other variants.

3 Protocols for Multi-Party AES

In this section, we construct the MPC protocols that compute AES. We define the ideal functionality \mathcal{F}_{AES} as the functionality that computes the AES block encryption taking secret-shared inputs $\{\llbracket x_i \rrbracket\}_{i=0,\dots,127}$ and a secret-shared encryption key $\{\llbracket k_i \rrbracket\}_{i=0,\dots,127}$, and returns shared outputs of $\text{Enc}(k, x)$, $\{\llbracket z_i \rrbracket\}_{i=0,\dots,127}$. Note that in our setting, i.e., replicated secret sharing with honest majority, we can transparently switch between shares of the binary extension $GF(2^k)$ and k many bit shares $GF(2)$, as these two have the same size. All proposed protocols securely compute the ideal functionality \mathcal{F}_{AES} .

3.1 Overview of the Proposed Protocols

To compute the AES algorithm in MPC, all steps need to be computed on secret-shared data. However, the linear layers of AES (ShiftRows, MixColumns, AddRoundKey) can be computed locally on the shares, as detailed in Sect. 2.6. Thus, we focus on constructing the non-linear operations within the KeyExpansion and SubBytes steps, that is, multiplicative inversions in $GF(2^8)$ within the S-box. This is (essentially) the only place where our variants differ. For completeness, we give the full oblivious AES algorithm in Protocol 9 in Appendix C. We briefly summarize the following approaches to compute multiplicative inversions in $GF(2^8)$; their costs are also summarized in Table 3.

The straightforward approach is to compute the inverse directly using a 256-elements lookup table. We call this proto-

Functionality $\mathcal{F}_{\text{LUT}}^{ss_1 \mapsto ss_2} / \mathcal{F}_{\text{LUT}}^{ss}$
Input: v shared under scheme $ss_1 \in \{\langle \cdot \rangle, \llbracket \cdot \rrbracket\}$, and public vector T . Output: T 's v -th element T_v shared under scheme $ss_2 \in \{\langle \cdot \rangle, \llbracket \cdot \rrbracket\}$. If $ss_1 = \langle \cdot \rangle$, also output $\llbracket v \rrbracket$.
<ol style="list-style-type: none"> 1. To define the input sharings, if $ss_1 = \langle \cdot \rangle$: <ul style="list-style-type: none"> • \mathcal{F}_{LUT} receives from each party a share $\langle v \rangle^i$, and reconstructs $v = \sum_i \langle v \rangle^i$. 2. Otherwise, if $ss_1 = \llbracket \cdot \rrbracket$: <ul style="list-style-type: none"> • \mathcal{F}_{LUT} receives from honest parties their shares $\llbracket v \rrbracket^{\mathcal{H}}$, and reconstructs v. • \mathcal{F}_{LUT} computes the corrupted parties' shares $\llbracket v \rrbracket^{\mathcal{C}}$ and sends these to the adversary. 3. \mathcal{F}_{LUT} looks up T_v. 4. \mathcal{F}_{LUT} receives from the adversary a set of corrupted shares $\text{sh}(T_v)^{\mathcal{C}}$, under scheme ss_2. 5. \mathcal{F}_{LUT} samples consistent honest shares $\text{sh}(T_v)^{\mathcal{H}}$ using $(T_v, \text{sh}(T_v)^{\mathcal{C}})$, under scheme ss_2. 6. If $ss_1 = \langle \cdot \rangle$: \mathcal{F}_{LUT} additionally receives consistent shares $\llbracket v \rrbracket^{\mathcal{C}}$ from the adversary, and defines the honest shares $\llbracket v \rrbracket^{\mathcal{H}}$ using $(v, \llbracket v \rrbracket^{\mathcal{C}})$. 7. \mathcal{F}_{LUT} outputs the shares $\text{sh}(T_v)^{\mathcal{H}}$, and optionally $\llbracket v \rrbracket^{\mathcal{H}}$, to the honest parties.

Figure 4: Ideal functionality for secret-shared table lookup.

col the LUT-256 (see Sect. 3.5). Here, parties prepare a public 256-elements lookup table for inversion and convert an input share $\llbracket x \rrbracket$ for $x \in GF(2^8)$ into a one-hot vector $e^{(x)} \in \{0, 1\}^{256}$ that has 1 at the position x and 0 otherwise. Then, they compute an inner product of the lookup table and the one-hot vector to obtain $\llbracket x^{-1} \rrbracket$. While this approach is optimal in round complexity, it requires heavy offline communication to generate random one-hot vectors.

The main protocol (Protocol 3 in Sect. 3.2), LUT-16, reduces such offline cost by using a smaller lookup table. The idea is to reduce the computation of multiplicative inversion over $GF(2^8)$ into the one over $GF(2^4)$ as shown in Eq. (1), (2) by using the isomorphism between $GF(2^8)$ and $GF((2^4)^2)$ from Sect. 2.1. Computing the inverse in $GF(2^4)$ only requires a lookup table of size 16.

Our lookup table protocol (Protocol 4) in Sect. 3.3 takes $\langle v \rangle, v \in GF(2^4)$ as well as a public table T and outputs the corresponding shared value $\llbracket T_v \rrbracket$ using the table. When the table is the inversion table, the obtained output $T_v = v^{-1}$. The key subfunctionality is $\mathcal{F}_{\text{RandOHV}}$ by which we can obtain a shared randomness $\llbracket r \rrbracket$ and the corresponding shared one-hot vector $\llbracket e^{(r)} \rrbracket$. This allows the table lookup to be performed as a linear function on $e^{(r)}$, after reconstructing a masked value $c = v \oplus r$ in the online phase, for a lookup table input v .

Table 3: Comparison of maliciously secure S-Box evaluation methods. # Mult is the number of multiplication triples checked by $\mathcal{F}_{\text{verify}}$.

	Offline Comm.	Online Comm.	Rounds	# Mult
Boolean circuit	–	32	6	32
$GF(2^8)$ circuit	–	32	4	4
$GF(2^4)$ circuit	–	20	4	3
$GF(2^4)$ /LUT-16	11	16	2	3
(3,3) LUT-256	22	16	1	6
(2,3) LUT-256	247	8	1	7

Functionality \mathcal{F}_{Inv}	
Input: $\llbracket x \rrbracket$	
Output: $\llbracket x^{-1} \rrbracket$	
1. \mathcal{F}_{Inv} receives from honest parties their shares of $\llbracket x \rrbracket^{\mathcal{H}}$ to reconstruct x .	
2. \mathcal{F}_{Inv} computes the corrupted parties' shares $\llbracket x \rrbracket^C$ and sends these shares to the adversary.	
3. \mathcal{F}_{Inv} obtains x^{-1} from x .	
4. \mathcal{F}_{Inv} receives a set of shares $\llbracket x^{-1} \rrbracket^C$ from the adversary.	
5. \mathcal{F}_{Inv} computes the honest shares $\llbracket x^{-1} \rrbracket^{\mathcal{H}}$ by using the set of shares $\llbracket x^{-1} \rrbracket^C$ and x^{-1} .	
6. \mathcal{F}_{Inv} outputs the shares $\llbracket x^{-1} \rrbracket^{\mathcal{H}}$ to the honest parties.	

Figure 5: The functionality \mathcal{F}_{Inv} for $GF(2^8)$.

3.2 Secure Protocol for Multiplicative Inverse

We propose a protocol for securely computing the multiplicative inverse in $GF(2^8)$ in Protocol 3. The ideal functionality \mathcal{F}_{Inv} is described in Fig. 5. By applying Φ from Sect. 2.1, the secure computation of a multiplicative inverse in $GF(2^8)$ can be reduced to three secure multiplications in $GF(2^4)$ and one secure multiplicative inverse in $GF(2^4)$.

To achieve malicious security, we need to ensure that corrupt parties use the correct additively shared input $\langle\langle v \rangle\rangle$ to $\mathcal{F}_{\text{LUT}}^{\langle\langle \cdot \rangle\rangle \rightarrow \llbracket \cdot \rrbracket}$. To do this, we rely on the fact that $\mathcal{F}_{\text{LUT}}^{\langle\langle \cdot \rangle\rangle \rightarrow \llbracket \cdot \rrbracket}$ also outputs the replicated sharing $\llbracket v \rrbracket$, giving a commitment to what value was actually used as input. We then check that v was correct by working backwards until the previous multiplication (step 6), computing replicated shares of the multiplication input. Then, if we verify this multiplication using $\mathcal{F}_{\text{verify}}$, this guarantees that the correct value was input to $\mathcal{F}_{\text{LUT}}^{\langle\langle \cdot \rangle\rangle \rightarrow \llbracket \cdot \rrbracket}$.

We prove the following in Appendix B.1.

Lemma 3.1 *The protocol Inv in Protocol 3 securely computes \mathcal{F}_{Inv} with abort in the $\{\mathcal{F}_{\text{LUT}}^{\langle\langle \cdot \rangle\rangle \rightarrow \llbracket \cdot \rrbracket}, \mathcal{F}_{\text{weakMult}}, \mathcal{F}_{\text{verify}}\}$ -hybrid model in the presence of a malicious adversary under the honest majority setting.*

Protocol 3 Multiplicative Inversion over $GF(2^8)$ against Malicious Adversary

Functionality: $\llbracket x^{-1} \rrbracket \leftarrow \mathcal{F}_{\text{Inv}}(\llbracket x \rrbracket)$

Input: Share $\llbracket x \rrbracket$ of $x \in GF(2^8)$

Output: Share $\llbracket x^{-1} \rrbracket$ of the inverse of $x \in GF(2^8)$

Subfunctionality: \mathcal{F}_{LUT}

- 1: $(\llbracket a_h \rrbracket, \llbracket a_\ell \rrbracket) \leftarrow \Phi(\llbracket x \rrbracket) \in GF((2^4)^2) \quad \triangleright \Phi \text{ in Sect. 2.1}$
- 2: $\llbracket a_h^2 \rrbracket := \text{Square}(\llbracket a_h \rrbracket)$
- 3: $\langle\langle a_h^2 \rangle\rangle \leftarrow \text{ToAdditive}(\llbracket a_h^2 \rrbracket)$
- 4: $\llbracket a_\ell^2 \rrbracket := \text{Square}(\llbracket a_\ell \rrbracket)$
- 5: $\langle\langle a_\ell^2 \rangle\rangle := \text{ToAdditive}(\llbracket a_\ell^2 \rrbracket)$
- 6: $\langle\langle a_h \times a_\ell \rangle\rangle := \text{MultLocal}(\llbracket a_h \rrbracket, \llbracket a_\ell \rrbracket)$
- 7: $\langle\langle v \rangle\rangle := (\{E\}_{16} \times \langle\langle a_h^2 \rangle\rangle) \oplus \langle\langle a_h \times a_\ell \rangle\rangle \oplus \langle\langle a_\ell^2 \rangle\rangle$
- 8: $(\llbracket v^{-1} \rrbracket, \llbracket v \rrbracket) \leftarrow \mathcal{F}_{\text{LUT}}^{\langle\langle \cdot \rangle\rangle \rightarrow \llbracket \cdot \rrbracket}(\langle\langle v \rangle\rangle, T^{\text{inv}}) \quad \triangleright 1 \text{ round, 8 bits}$
- 9: $\llbracket a_h \times a_\ell \rrbracket := \llbracket v \rrbracket \oplus (\{E\}_{16} \times \llbracket a_h^2 \rrbracket) \oplus \llbracket a_\ell^2 \rrbracket$
- 10: $\llbracket a_h' \rrbracket \leftarrow \mathcal{F}_{\text{weakMult}}(\llbracket a_h \rrbracket, \llbracket v^{-1} \rrbracket)$
- 11: $\llbracket a_\ell' \rrbracket \leftarrow \mathcal{F}_{\text{weakMult}}(\llbracket a_h \rrbracket \oplus \llbracket a_\ell \rrbracket, \llbracket v^{-1} \rrbracket) \quad \triangleright 1 \text{ round, 8 bits}$
- 12: $\llbracket y \rrbracket \leftarrow \Phi^{-1}(\llbracket a_h' \rrbracket, \llbracket a_\ell' \rrbracket) \in GF(2^8)$
- 13: Execute $\mathcal{F}_{\text{verify}}$ for the following multiplication triples:
 1. $(\llbracket a_h \rrbracket, \llbracket a_\ell \rrbracket, \llbracket a_h \times a_\ell \rrbracket)$
 2. $(\llbracket a_h \rrbracket, \llbracket v^{-1} \rrbracket, \llbracket a_h' \rrbracket)$
 3. $(\llbracket a_h \oplus a_\ell \rrbracket, \llbracket v^{-1} \rrbracket, \llbracket a_\ell' \rrbracket)$
- 14: **return** $\llbracket y \rrbracket$

Reducing the number of multiplication checks. Instead of checking 3 multiplications, we observe that the protocol can be optimized by embedding all checks into one multiplication. For some $k \geq 3$ and degree- k , irreducible polynomial $f(\alpha)$ over $GF(2^4)$, define the extension field $K = GF((2^4)^k) = GF(2^4)[\alpha]/f(\alpha)$. An element of K can be expressed as a polynomial $a_{k-1}\alpha^{k-1} + \dots + a_1\alpha + a_0$, for $a_i \in GF(2^4)$. Then, we can check the following equation over K :

$$(\llbracket a_h \rrbracket + \alpha \llbracket a_\ell \rrbracket) \cdot (\llbracket a_\ell \rrbracket + \alpha \llbracket v^{-1} \rrbracket) = \llbracket a_h \times a_\ell \rrbracket + \alpha \llbracket a_h' \oplus a_\ell^2 \rrbracket + \alpha^2 \llbracket a_\ell' \oplus a_h' \rrbracket. \quad (3)$$

Note that shares of a_ℓ^2 can be computed locally.

Lemma 3.2 *If the checking equation Eq. (3) holds, the multiplication triples (1), (2) and (3) in Protocol 3 are all correct.*

Proof: Let $d_v \in GF(2^4)$ be the error introduced by \mathcal{A} for the input $\langle\langle v \rangle\rangle$ of $\mathcal{F}_{\text{LUT}}^{\langle\langle \cdot \rangle\rangle \rightarrow \llbracket \cdot \rrbracket}$, thus the functionality returns $(\llbracket (v \oplus d_v)^{-1} \rrbracket, \llbracket v \oplus d_v \rrbracket)$. Then further, let $d_1, d_2 \in GF(2^4)$ denote the additive errors introduced in $\llbracket a_h' \oplus d_1 \rrbracket \leftarrow \mathcal{F}_{\text{weakMult}}(\llbracket a_h \rrbracket, \llbracket (v \oplus d_v)^{-1} \rrbracket)$ and $\llbracket a_\ell' \oplus d_2 \rrbracket \leftarrow \mathcal{F}_{\text{weakMult}}(\llbracket a_h \rrbracket \oplus \llbracket a_\ell \rrbracket, \llbracket (v \oplus d_v)^{-1} \rrbracket)$.

The parties call $\mathcal{F}_{\text{verify}}$ with the tuple

$$(\llbracket a_h \rrbracket + \alpha \llbracket a_\ell \rrbracket, \llbracket a_\ell \rrbracket + \alpha \llbracket (v \oplus d_v)^{-1} \rrbracket, \llbracket a_h \times a_\ell \oplus d_v \rrbracket + \alpha \llbracket a_h' \oplus d_1 \oplus a_\ell^2 \rrbracket + \alpha^2 \llbracket a_\ell' \oplus d_2 \oplus a_h' \rrbracket),$$

corresponding to Eq. (3) where $a'_h = a_h(v \oplus d_v)^{-1}$ and $a'_\ell = (a_h \oplus a_\ell)(v \oplus d_v)^{-1}$. Note that $\mathcal{F}_{\text{verify}}$ computes the check and obtains

$$0 = d_v + \alpha d_1 + \alpha^2(d_1 \oplus d_2),$$

where both sides are elements in the extension field K . The check is zero if and only if d_v , d_1 and d_2 are all zero. \square

3.3 Secure Table Lookup Protocol

We present Protocol 4 that securely computes the ideal functionality $\mathcal{F}_{\text{LUT}}^{\langle\cdot\rangle \rightarrow \llbracket \cdot \rrbracket}$ from Fig. 4. Step 1 executes the ideal functionality $\mathcal{F}_{\text{RandOHV}}$, which performs secure random one-hot vector encoding that takes $N = 2^k$ as input and outputs the shared randomness $\llbracket r \rrbracket$ and its corresponding shared random one-hot vector $\llbracket e^{(r)} \rrbracket$ (see Fig. 6). Using the randomness, the protocol masks the input and reveals $c = v + r$. The rest of the computation can be done locally as in Step 4-5.

We show that the proposed approach correctly computes the desired value. It is sufficient to show that for each i , the value t_i computed in Step 4 matches the i -th digit $T_v^{(i)}$ of T 's v -th element T_v . According to the definition, $t = \bigoplus_{0 \leq j \leq n-1} e_j^{(r)} T_{c \oplus j}$, and since the inner product on the right-hand side turns 0 for all terms but $e_r^{(r)} T_{c \oplus r}$, resulting in $t = T_v$.

Protocol 4 Table lookup of size $N = 2^k$, from $\langle\cdot\rangle$ to $\llbracket \cdot \rrbracket$ sharing

Functionality: $\llbracket T_v \rrbracket \leftarrow \mathcal{F}_{\text{LUT}}(\langle\langle v \rangle\rangle, T)$

Input: Share $\langle\langle v \rangle\rangle$ of $v \in GF(2^k)$, table $T : GF(2^k) \rightarrow GF(2^\ell)$

Output: Share $\llbracket T_v \rrbracket$ of the value $T_v \in GF(2^\ell)$

Subfunctionality: $\mathcal{F}_{\text{RandOHV}}$

- 1: $(\{\llbracket r_i \rrbracket\}_{0 \leq i < k}, \{\llbracket e_j^{(r)} \rrbracket\}_{0 \leq j < N}) \leftarrow \mathcal{F}_{\text{RandOHV}}(k)$
 - 2: $\langle\langle 0 \rangle\rangle \leftarrow \mathcal{F}_{\text{Zero}}(GF(2^k))$
 - 3: $c := \text{Reconst}(\langle\langle v \rangle\rangle + \text{ToAdditive}(\llbracket r \rrbracket) + \langle\langle 0 \rangle\rangle)$ $\triangleright 1$
 - round, $2k$ bits
 - 4: $\llbracket t \rrbracket := \bigoplus_{j=0}^{N-1} \llbracket e_j^{(r)} \rrbracket \cdot T_{c \oplus j}$
 - 5: $\llbracket v \rrbracket := \llbracket r \rrbracket \oplus c$
 - 6: **return** $(\llbracket t \rrbracket, \llbracket v \rrbracket)$
-

We prove the following in Appendix B.2.

Lemma 3.3 *The protocol LUT in Protocol 4 securely computes $\mathcal{F}_{\text{LUT}}^{\langle\cdot\rangle \rightarrow \llbracket \cdot \rrbracket}$ in the $\{\mathcal{F}_{\text{RandOHV}}, \mathcal{F}_{\text{Zero}}\}$ -hybrid model in the presence of a malicious adversary.*

3.4 General One-Hot Vector Protocol

We now show how to compute the random one-hot vector that was required in the table lookup. We start with a general protocol Ohv to securely compute the one-hot vector of a shared input and turn it to a random one-hot vector protocol by inputting a set of random shared bits (see Protocol 5). In

Functionality $\mathcal{F}_{\text{RandOHV}}$
Input: $N = 2^k$
Output: k random bits $\{\llbracket r_i \rrbracket\}_{i=0}^{k-1}$ and length- N one-hot vector $\llbracket e^{(r)} \rrbracket$ for $r = r_{k-1} \dots r_0$
<ol style="list-style-type: none"> 1. $\mathcal{F}_{\text{RandOHV}}$ receives from the adversary \mathcal{A} the shares $\{\llbracket r_i \rrbracket^C\}_{i=0}^{k-1}$ and $\llbracket e^{(r)} \rrbracket^C$ considered as the shares of randomness and the corresponding share of the one-hot vector held by corrupted parties. 2. $\mathcal{F}_{\text{RandOHV}}$ samples $\{r_i\}_{i=0}^{k-1}$ then computes $r = \sum_{i=0}^{k-1} 2^i \cdot r_i$ and $e^{(r)}$. 3. $\mathcal{F}_{\text{RandOHV}}$ generates $\llbracket r_i \rrbracket^{\mathcal{H}}$ from $(r_i, \llbracket r_i \rrbracket^C)$ for $i \in \{0, 1, \dots, k-1\}$. 4. $\mathcal{F}_{\text{RandOHV}}$ generates $\llbracket e^{(r)} \rrbracket^{\mathcal{H}}$ from $e^{(r)}$ and $\llbracket e^{(r)} \rrbracket^C$. 5. $\mathcal{F}_{\text{RandOHV}}$ distributes the shares $\llbracket r_i \rrbracket^{\mathcal{H}}$ and $\llbracket e^{(r)} \rrbracket^{\mathcal{H}}$ to the honest parties.

Figure 6: The functionality $\mathcal{F}_{\text{RandOHV}}$ to create a random one-hot vector of length N .

Appendix A, we also give a specialized protocol for vectors of length 16 with lower round complexity.

We sketch the procedure of Ohv, which is based on the dishonest majority protocol from [40]. It takes $(\llbracket v_{k-1} \rrbracket, \dots, \llbracket v_0 \rrbracket)$ as input. First, it selects bit $\llbracket v_0 \rrbracket$ and creates a one-hot vector with length 2, $(1 - \llbracket v_0 \rrbracket, \llbracket v_0 \rrbracket)$. Then, it selects bit $\llbracket v_1 \rrbracket$ and computes a one-hot vector $((1 - \llbracket v_1 \rrbracket) \cdot (1 - \llbracket v_0 \rrbracket, \llbracket v_0 \rrbracket), \llbracket v_1 \rrbracket \cdot (1 - \llbracket v_0 \rrbracket, \llbracket v_0 \rrbracket))$ with length 4. This is repeated until bit $\llbracket v_{k-1} \rrbracket$ computes a one-hot vector $((1 - \llbracket v_{k-1} \rrbracket) \cdot \vec{f}, \llbracket v_{k-1} \rrbracket \cdot \vec{f})$ with length 2^k , where \vec{f} is a one-hot vector with length 2^{k-1} from the previous iteration.

The communication complexity with three parties is as follows. For general N , the communication cost is $N - \log N - 1$ bits within $\log N - 1$ rounds. For $N = 256$, the communication cost will be 247 bits.

3.5 Approaches Using Large Lookup Tables

In this section, we explore alternative ways of securely computing the AES S-box, using a single lookup table of size 256. We present two protocols for secure AES evaluation with different tradeoffs in communication complexity and round complexity. We also present a third protocol, in Section 3.5.4, which is less efficient for AES, but allows securely evaluating an arbitrary lookup table of size N on $\llbracket \cdot \rrbracket$ -shared values, with a communication cost in $O(\sqrt{N})$ instead of $O(N)$. The cost of this protocol (Protocol 10) is shown in Table 4, together with other protocols for comparison.

3.5.1 New $\mathcal{F}_{\text{LUT}}^{ss1 \mapsto ss2}$ Instantiations.

As building blocks, we use three variants of the $\mathcal{F}_{\text{LUT}}^{ss1 \mapsto ss2}$ functionality with different combinations of secret sharing schemes, namely, $\mathcal{F}_{\text{LUT}}^{\llbracket \cdot \rrbracket}$, $\mathcal{F}_{\text{LUT}}^{\langle\cdot\rangle \rightarrow \llbracket \cdot \rrbracket}$ and $\mathcal{F}_{\text{LUT}}^{\llbracket \cdot \rrbracket \rightarrow \langle\cdot\rangle}$.

Protocol 5 Random One-hot Vector

Functionality: $(\llbracket r \rrbracket, \llbracket e^{(r)} \rrbracket) \leftarrow \mathcal{F}_{\text{RandOHV}}(N = 2^k)$
Input: \perp
Output: Shared random bits $\{\llbracket r_i \rrbracket\}_{i=0}^{k-1}$, and length- N one-hot vector $\llbracket e^{(r)} \rrbracket$ for $r = \sum_{i=0}^{k-1} 2^i r_i$

```

1:  $(\llbracket r_{k-1} \rrbracket, \dots, \llbracket r_0 \rrbracket) := \mathcal{F}_{\text{rand}}(k)$ 
2:  $\llbracket e^{(r)} \rrbracket := \text{OHV}(\llbracket r_{k-1} \rrbracket, \dots, \llbracket r_0 \rrbracket; N)$ 
3: Execute  $\mathcal{F}_{\text{verify}}$  for the following multiplication triplets
   from Ohv:  $(v_{i-1}, (f_0, \dots, f_{2^{i-1}-2}), (e_0, \dots, e_{2^{i-1}-2}))$  for all  $i \in \{2, \dots, k\}$ 
4: return  $(\{\llbracket r_i \rrbracket\}_{i=0}^k, \llbracket e^{(r)} \rrbracket)$ 

5: procedure Ohv( $(\llbracket v_{k-1} \rrbracket, \dots, \llbracket v_0 \rrbracket; N = 2^k)$ )
6:   if  $N = 2$  then
7:     return  $(1 - \llbracket v_0 \rrbracket, \llbracket v_0 \rrbracket)$ 
8:   else  $\triangleright \log N - 2$  recursive calls with  $N > 2$ 
9:      $\llbracket f_0 \rrbracket, \dots, \llbracket f_{N/2-1} \rrbracket := \text{Ohv}(\llbracket v_{k-2} \rrbracket, \dots, \llbracket v_0 \rrbracket; N/2)$ 
10:     $\llbracket e_0 \rrbracket, \dots, \llbracket e_{N/2-2} \rrbracket := \text{Mult}(\llbracket v_{k-1} \rrbracket, (\llbracket f_0 \rrbracket, \dots, \llbracket f_{N/2-2} \rrbracket))$ 
     $\triangleright 1$  round,  $N/2 - 1$  bits; store triple  $v_{k-1}, (f_0, \dots, f_{N/2-2}), (e_0, \dots, e_{N/2-2})$ 
11:     $\llbracket e_{N/2-1} \rrbracket := \llbracket v_{k-1} \rrbracket - \bigoplus_{i=0}^{N/2-2} \llbracket e_i \rrbracket$ 
12:     $\llbracket e^{(v)} \rrbracket := (\llbracket f_0 \rrbracket - \llbracket e_0 \rrbracket, \dots, \llbracket f_{N/2-1} \rrbracket - \llbracket e_{N/2-1} \rrbracket, \llbracket e_0 \rrbracket, \dots, \llbracket e_{N/2-1} \rrbracket)$ 
13:    return  $\llbracket e^{(v)} \rrbracket$   $\triangleright e^{(v)} = ((1 - v_{k-1}) \cdot \vec{f}, v_{k-1} \cdot \vec{f})$ 
14:   end if
15: end procedure

```

Table 4: Comparison of protocols for table lookup of size $N = 2^k$, with communication complexity for $n = 3, t = 1$. # Mult is the length of the input to $\mathcal{F}_{\text{verify}}$ needed for malicious security

Protocol	Offline	Online	Rounds	# Mult
$\langle\langle \cdot \rangle\rangle \mapsto \llbracket \cdot \rrbracket$ (Prot. 4)	$N - k - 1$	$2k$	1	$k - 1$
$\llbracket \cdot \rrbracket \mapsto \llbracket \cdot \rrbracket$ (Prot. 4 variant)	$N - k - 1$	k	1	$k - 1$
$\langle\langle \cdot \rangle\rangle \mapsto \langle\langle \cdot \rangle\rangle$ (Prot. 6)	$2(\sqrt{N} - \frac{k}{2} - 1)$	$2k$	2	$k - 2$
$\llbracket \cdot \rrbracket \mapsto \langle\langle \cdot \rangle\rangle$ (Prot. 6 variant)	$2(\sqrt{N} - \frac{k}{2} - 1)$	k	1	$k - 2$
$\llbracket \cdot \rrbracket \mapsto \llbracket \cdot \rrbracket$ (Prot. 10)	$2(\sqrt{N} - \frac{k}{2} - 1)$	$2k$	2	N

In $\mathcal{F}_{\text{LUT}}^{\llbracket \cdot \rrbracket}$, both the input $\llbracket v \rrbracket$ and output $\llbracket T_v \rrbracket$ are given as replicated sharings. This is a stronger requirement than previously, where v was only given additively shared, allowing an adversary to add an error to the input. The simplest way to realize $\mathcal{F}_{\text{LUT}}^{\llbracket \cdot \rrbracket}$ is with a slight tweak to the LUT protocol (Protocol 4): since the input is given in replicated shares, we now run the Reconst procedure (step 3) on the replicated sharing $\llbracket v + r \rrbracket$ instead of $\langle\langle v + r \rangle\rangle$. For a small number of parties, this reduces communication since opening replicated shares is cheaper. For instance, with $n = 3, t = 1$, the cost is reduced from $2k$ bits per party down to just k . Note that the preprocessing cost — generating a random one-hot vector of length $N = 2^k$ via Protocol 5 — is identical to that of \mathcal{F}_{LUT} .

$\mathcal{F}_{\text{LUT}}^{\langle\langle \cdot \rangle\rangle}$ can be implemented using Protocol 6. This protocol is very similar to that for $\mathcal{F}_{\text{LUT}}^{\langle\langle \cdot \rangle\rangle \mapsto \llbracket \cdot \rrbracket}$ (Protocol 4), except the one-hot vector only needs to be generated in additive shares, rather

than replicated shares. This allows for a much more efficient preprocessing protocol: the parties can run the replicated one-hot vector functionality $\mathcal{F}_{\text{RandOHV}}$ twice on input length $2^{k/2}$, obtaining two one-hot vectors $\llbracket e^{(r)} \rrbracket, \llbracket e^{(r')} \rrbracket$. Then, they can locally compute additive shares of the tensor product vector $e^{(r)} \times e^{(r')}$, giving a one-hot vector of length 2^k . Note that, since $\mathcal{F}_{\text{RandOHV}}$ gives replicated shares of the non-zero index, the same shares can still be used to obtain replicated shares of the index of the length 2^k vector.

Protocol 6 Table Lookup of size $N = 2^k$ in additive sharing

Functionality: $(\langle\langle T_v \rangle\rangle, \llbracket v \rrbracket) \leftarrow \mathcal{F}_{\text{LUT}}(\langle\langle v \rangle\rangle, T)$
Input: Share $\langle\langle v \rangle\rangle$ of $v \in GF(2^k)$, table $T : GF(2^k) \rightarrow GF(2^\ell)$
Output: Share $\langle\langle T_v \rangle\rangle$ of the value $T_v \in GF(2^\ell)$, and share $\llbracket v \rrbracket$
Subfunctionality: $\mathcal{F}_{\text{RandOHV}}$

```

1: Call  $\mathcal{F}_{\text{RandOHV}}(k/2)$  twice to get
    $(\{\llbracket r_i \rrbracket\}_{i=0}^{k/2-1}, \{\llbracket e_j^{(r)} \rrbracket\}_{j=0}^{\sqrt{N}-1}), (\{\llbracket r'_i \rrbracket\}_{i=0}^{k/2-1}, \{\llbracket e'_j \rrbracket\}_{j=0}^{\sqrt{N}-1})$ 
2:  $\llbracket r \rrbracket := (\llbracket r_0 \rrbracket, \dots, \llbracket r_{k/2-1} \rrbracket, \llbracket r'_0 \rrbracket, \dots, \llbracket r'_{k/2-1} \rrbracket)$ 
3:  $\langle\langle 0 \rangle\rangle \leftarrow \mathcal{F}_{\text{zero}}(GF(2^k))$ 
4:  $c \leftarrow \text{Reconst}(\langle\langle v \rangle\rangle + \text{ToAdditive}(\langle\langle r \rangle\rangle) + \langle\langle 0 \rangle\rangle)$   $\triangleright 1$ 
   round,  $2k$  bits
5:  $\llbracket v \rrbracket := \llbracket r \rrbracket + c$ 
6:  $\langle\langle \vec{f} \rangle\rangle \leftarrow \left( \text{MultLocal}(\llbracket e_i^{(r)} \rrbracket, \llbracket e'_j \rrbracket) \right)_{i,j=0}^{\sqrt{N}-1}$ 
7:  $\langle\langle t \rangle\rangle \leftarrow \bigoplus_{j=0}^{N-1} \langle\langle \vec{f}_j \rangle\rangle \cdot T_{c \oplus j}$ 
8: return  $(\langle\langle t \rangle\rangle, \llbracket v \rrbracket)$ 

```

The protocol securely realises the functionality $\mathcal{F}_{\text{LUT}}^{\langle\langle \cdot \rangle\rangle}$ from Figure 4. Despite being maliciously secure, one must still take care when composing this protocol with others, since the ideal functionality inherently allows an adversary to cheat by simply changing its additive share of the input or output. In Section 3.5.3, we show how to overcome this issue for the case of AES. We omit the proof of the following, which is very similar to Lemma 3.3.

Lemma 3.4 Protocol 6 securely realizes the functionality $\mathcal{F}_{\text{LUT}}^{\langle\langle \cdot \rangle\rangle}$ in the $(\mathcal{F}_{\text{RandOHV}}, \mathcal{F}_{\text{zero}})$ -hybrid model.

Finally, $\mathcal{F}_{\text{LUT}}^{\llbracket \cdot \rrbracket \mapsto \langle\langle \cdot \rangle\rangle}$ can be implemented using Protocol 6, except the opening of c in step 4 is done on replicated sharings instead of additive. This achieves the strongest performance characteristics of all variants: only k bits of communication per party (for $n = 3, t = 1$) and a cheap preprocessing phase that only requires two replicated, random one-hot vectors of length $2^{k/2}$.

3.5.2 AES Protocol Based on Replicated Sharing

Given the $\mathcal{F}_{\text{LUT}}^{\llbracket \cdot \rrbracket}$ functionality, our protocol for AES evaluation is straightforward. We assume the parties start with replicated shares of the input and expanded key. Then, each S-box is

evaluated with a single call to $\mathcal{F}_{\text{LUT}}^{[\cdot]}$ of length $N = 256$, and the linear layers are evaluated locally on the shares. Assuming a maliciously secure implementation of $\mathcal{F}_{\text{LUT}}^{[\cdot]}$, this protocol is maliciously secure, since $\mathcal{F}_{\text{LUT}}^{[\cdot]}$ does not allow any errors to be introduced by the adversary.

With $n = 3, t = 1$, the cost of this protocol in the online phase is just 8 bits of communication per party per S-box, or a total of $10 \cdot 8 \cdot 16 = 1280$ bits for one block of AES. The total round complexity is 10 rounds. The preprocessing phase, however, is much more expensive, due to the need to generate a large, replicated one-hot vector for each S-box. This costs 247 bits per party, per S-box, for a total of 39520 bits. Achieving malicious security can be done by batch verifying each of the multiplications in the RndOhv protocol using $\mathcal{F}_{\text{verify}}$.

3.5.3 AES Protocol Based on Additive Sharing

By relying on $\mathcal{F}_{\text{LUT}}^{\langle\langle\cdot\rangle\rangle}$ instead of $\mathcal{F}_{\text{LUT}}^{[\cdot]}$, we can reduce the preprocessing cost of the previous protocol by more than 10x. This is because $\mathcal{F}_{\text{LUT}}^{\langle\langle\cdot\rangle\rangle}$ can be realized by generating only two replicated one-hot vectors of length 16, instead of one of length 256. Running the whole protocol on additive instead of replicated shared inputs, we get a 3-party, passively secure protocol with an online communication complexity of 16 bits per S-box (2560 bits overall) and only 22 bits per S-box in the preprocessing phase (3200 overall). The main challenge is now how to add malicious security since we are now dealing with additive shares which are easily tampered with.

We consider two approaches to malicious security. First, we describe a specialized method tailored to AES, which needs no additional communication except for one call to $\mathcal{F}_{\text{verify}}$. The core idea is to exploit the fact that $\mathcal{F}_{\text{LUT}}^{\langle\langle\cdot\rangle\rangle}$ outputs replicated shares of its input x , and use this to obtain replicated shares of the $\mathcal{F}_{\text{LUT}}^{\langle\langle\cdot\rangle\rangle}$ output from the previous round, by evaluating the AES linear layer backwards. We combine this with a cheap way of verifying input/output S-box pairs by verifying two multiplication triples, relying on the algebraic structure of the S-box.

Our second approach is more general and can be used to realize $\mathcal{F}_{\text{LUT}}^{[\cdot]}$ with malicious security for arbitrary lookup tables of domain size 2^k . The preprocessing cost is the same as the AES-specific protocol, but the online phase has slightly more communication and requires using $\mathcal{F}_{\text{verify}}$ to verify a length- 2^k dot product triple, instead of just two multiplications.

AES-Optimized Protocol. We present the full protocol for AES evaluation in Protocol 7. The protocol begins with the inputs and round keys distributed as replicated shares. The first round of S-boxes is computed with $\mathcal{F}_{\text{LUT}}^{[\cdot] \rightarrow \langle\langle\cdot\rangle\rangle}$. For each subsequent round, we proceed to evaluate the linear layer, denoted L_i , and round key addition, followed by the S-box with $\mathcal{F}_{\text{LUT}}^{\langle\langle\cdot\rangle\rangle}$ to get an S-box output $y_i = \text{SubBytes}(x_i)$. We then invert the linear layer on the replicated shares of x_i to recover

Protocol 7 $\langle\langle\cdot\rangle\rangle$ -LUT based AES

Input: Message $\llbracket x \rrbracket$, round keys $\{\llbracket k^{(i)} \rrbracket, \langle\langle k^{(i)} \rangle\rangle\}_{i=0}^{10}$

Output: $\llbracket z \rrbracket$, where $z = \text{AES}_k(x)$

```

1:  $\llbracket x^0 \rrbracket \leftarrow \llbracket x \rrbracket + \llbracket k^{(0)} \rrbracket$  ▷AddRoundKey
    $\triangleright x_b^i, y_b^i$  is byte  $b$  of  $x^i, y^i$ 
2:  $\{\langle\langle y_b^0 \rangle\rangle\}_{b=0}^{15} \leftarrow \{\mathcal{F}_{\text{LUT}}^{[\cdot] \rightarrow \langle\langle\cdot\rangle\rangle}(\llbracket x_b^0 \rrbracket)\}_{b=0}^{15}$  ▷SubBytes
3: for  $i = 1, \dots, 9$  do
4:    $\langle\langle x^i \rangle\rangle \leftarrow L_i(\langle\langle y^{i-1} \rangle\rangle) + \langle\langle k^{(i)} \rangle\rangle$  ▷Shift/Mix/AddRK
5:    $\{\langle\langle y_b^i \rangle\rangle\}_{b=0}^{15} \leftarrow \{\mathcal{F}_{\text{LUT}}^{\langle\langle\cdot\rangle\rangle}(\langle\langle x_b^i \rangle\rangle)\}_{b=0}^{15}$  ▷SubBytes
6:    $\llbracket y^{i-1} \rrbracket \leftarrow L_i^{-1}(\llbracket x^i \rrbracket - \llbracket k^{(i)} \rrbracket)$ 
7: end for
8:  $\llbracket y^9 \rrbracket \leftarrow \text{Reshare}(\langle\langle y^9 \rangle\rangle)$ 
9:  $\llbracket z \rrbracket \leftarrow L_{10}(\llbracket y^9 \rrbracket) + \llbracket k^{(10)} \rrbracket$  ▷ShiftRows/AddRoundKey
10:  $\text{triples} \leftarrow \bigcup_{i=0}^9 \bigcup_{b=0}^{15} \text{VERIFYSBX}(\llbracket x_b^i \rrbracket, \text{Affine}^{-1}(\llbracket y_b^i \rrbracket))$ 
11: Run  $\mathcal{F}_{\text{verify}}$  to check triples
12: return  $\llbracket z \rrbracket$ 
```

```

13: procedure VERIFYSBX( $\llbracket x \rrbracket, \llbracket y \rrbracket$ )
14:    $\llbracket x^2 \rrbracket = \text{Square}(\llbracket x \rrbracket)$ 
15:    $\llbracket y^2 \rrbracket = \text{Square}(\llbracket y \rrbracket)$ 
16:   return  $\{(\llbracket x^2 \rrbracket, \llbracket y \rrbracket, \llbracket x \rrbracket), (\llbracket x \rrbracket, \llbracket y^2 \rrbracket, \llbracket y \rrbracket)\}$ 
17: end procedure
```

replicated shares of y_{i-1} . Finally, each S-box input/output pair (x_i, y_i) is verified with $\mathcal{F}_{\text{verify}}$, by first inverting the affine component of the S-box, denoted Affine , to get \hat{y}_i , and checking the two equations: $x_i^2 \hat{y}_i = x_i$, and $x_i \hat{y}_i^2 = \hat{y}_i$ which hold if and only if $\hat{y}_i = x_i^{254}$ in $GF(2^8)$. This idea was recently proposed for zero-knowledge proofs of AES [8].

One subtlety of the security proof (in Appendix B.3) is that when using $\mathcal{F}_{\text{verify}}$, we *cannot* allow the adversary to learn the errors in the multiplication triples, e.g. the values $d_i = x_i^2 \hat{y}_i - x_i$. This is because an error in an S-box input corresponds to an error in x_i , which would lead to a non-zero value of d_i that leaks information on \hat{y}_i . While at first glance, it may seem that even the *presence* of input-dependent errors would leak information to the adversary, in our case it is not a security issue to reveal whether *some* error occurred: if any error d_i is non-zero then at least one of $x_i^2 \hat{y}_i = x_i$ or $x_i \hat{y}_i^2 = \hat{y}_i$ must be false. The key point is that we cannot reveal the value or location of this error, which is why we need the stronger $\mathcal{F}_{\text{verify}}$ functionality from Section 2.7.

We prove the following in Appendix B.3.

Lemma 3.5 *Protocol 7 securely realizes the functionality \mathcal{F}_{AES} in the $(\mathcal{F}_{\text{LUT}}^{\langle\langle\cdot\rangle\rangle}, \mathcal{F}_{\text{verify}})$ -hybrid model with malicious security.*

3.5.4 Improved Protocol for $\mathcal{F}_{\text{LUT}}^{[\cdot]}$

In Protocol 10, shown in Appendix E, we present an alternative protocol for $\mathcal{F}_{\text{LUT}}^{[\cdot]}$, which generalizes the ideas of the

previous protocol to arbitrary lookup tables. Compared with the naive protocol for $\mathcal{F}_{\text{LUT}}^{\llbracket \cdot \rrbracket}$ discussed in the previous section, we reduce the communication cost of the preprocessing to $O(\sqrt{N})$. The online phase, however, has roughly double the cost in terms of communication and rounds. This is not as efficient as the AES-specific protocol in the previous section, but may still be useful in other applications.

The protocol follows a similar approach to Protocol 6, building a length- N one-hot vector taking the tensor product of two length- \sqrt{N} vectors, except it works on replicated shared inputs. Recall that Protocol 6 computes the output $v = \bigoplus_j \tilde{f}_j \cdot T_{c \oplus j}$, where c is a masked version of the input and \tilde{f} is the one-hot vector. We observe that, when \tilde{f} is decomposed into a tensor product of two smaller one-hot vectors, v can be seen as an inner product of two *secret*, $\llbracket \cdot \rrbracket$ -shared vectors of length N . This means we can tweak this to obtain *replicated* shares of the output, by using the $\mathcal{F}_{\text{weakDotProduct}}$ functionality, followed by $\mathcal{F}_{\text{verify}}$ to obtain malicious security.

4 Performance

We implemented the proposed protocols and two most related protocols from the state of the art in the same software framework for a fair comparison. Our code in the Rust programming language is available². Non-linear operations in small fields, e.g., $GF(2^8)$, $GF(2^4)$, are implemented via table lookups. Networking and I/O is done in a separate thread where each channel between two parties is encrypted and mutually authenticated using TLS1.3 with client/server certificates. Local randomness, for instance, to implement $\mathcal{F}_{\text{rand}}$, comes from a PRNG based on ChaCha20, the hash function we use for compare-view is SHA-256.

For semi-honest security, we implement Chida et al.’s $GF(2^8)$ -Circuit [21, Algorithm 5] as the baseline. We implement LUT-16 (Protocol 3 using Protocol 8 in the offline phase to generate random one-hot vectors), $GF(2^4)$ -Circuit (Protocol 3 but with the $GF(2^4)$ inverse $\llbracket v^{-1} \rrbracket$ as $\llbracket v^2 \rrbracket \cdot \llbracket v^4 \rrbracket \cdot \llbracket v^8 \rrbracket$), (2,3) LUT-256 (Protocol 4 using Protocol 5 to generate length-256 random one-hot vectors) and (3,3) LUT-256 (Protocol 7 with Protocol 6 and Protocol 8 in the offline phase to generate two length-16 random one-hot vectors).

For malicious security, we check the multiplications for correctness using Protocol 2 in all variants described above except (2,3) LUT-256³. Protocol 2 is implemented in $GF(2^{64})$ to achieve an acceptable level of soundness of at least 40 bits and utilize hardware support for carry-less multiplication (CLMUL). We also gain some efficiency by doing modular reduction only once at the end when computing inner products. We include a baseline from previous work, $GF(2^8)$ -Circuit ([21] + [31]), that uses bucket cut-and-choose adapted to $GF(2^8)$ with bucket size $B = 3$ and $C = 3$ triples to open (for

$\geq 2^{19}$ multiplications, i.e., ≥ 820 AES blocks). All described optimizations to reduce the number of multiplications triples to check have been implemented. Note that we estimate that Protocol 2 leads to about a factor 3 improvement in local computation compared to the techniques from [15, 44]. Thus, the performance of $GF(2^8)$ -Circuit + Protocol 2 should have equal or better performance⁴.

4.1 Experimental Setup

We experimentally evaluate the performance of the proposed protocols in different settings using two sets of three machines. Within each set, the machines have identical specifications: the first set (16-core Intel Core i9-9900 3.10GHz, 128GB RAM) is used in the 10 Gbit/s setting (≈ 9.47 Gbit/s, < 1 ms latency), and the second set (16-core Intel XEON E5-2650v2 2.60GHz, 128 GB RAM) is used for the remaining network settings, i.e., the 1 Gbit/s, 200 Mbit/s, 100 Mbit/s and WAN (50 Mbit/s) setting. The network throughput/latency was altered using `tc`. In all settings, 16 computation threads were used. We measure execution time (wall clock time) and the total amount of bytes sent per party during the computation of s many parallel AES block ciphers without the keyschedule to amortize performance. The reported benchmark data is the execution time/data communication of the slowest party, averaged over at least 10 iterations of the protocol. The throughput (denoted in AES blocks per second) is computed as $\lfloor s/t \rfloor$ where t is the execution time in seconds.

4.2 Benchmark Results

We defer a detailed description of the benchmark results to Appendix F and only highlight the main results for the high-throughput setting in Table 5. Table 8 in Appendix F summarizes the computation latency for one AES block. While our protocols improve the online phase and total throughput by a factor of 1.3 to 1.7 for semi-honest security compared to the state-of-the-art, we cannot name a clear winner. The journey to reduce communication creates trade-offs, such as introducing preprocessing (LUT-16), more rounds ($GF(2^4)$ -Circuit) or increased local computation (LUT-256). The saved communication translates into higher throughput, particularly benefiting lower bandwidth networks, but we cannot realize the full potential of a very efficient online phase for the LUT-256 variants due to the much more involved local computation.

For active security, replacing the bucket cut-and-choose with sublinear checks increases the total throughput by factor 4 to 8, even in high bandwidth networks. Considering only the throughput of the online phase in the high-bandwidth settings, we find that linear communication techniques for checking multiplications, e.g., the triple sacrifice of [31], outperforms

²<https://github.com/KULEuven-COSIC/maestro>

³Its costly preprocessing phase makes this variant unappealing.

⁴In the LAN setting, [44] report a performance of 10000 AES blocks in 1.08s, i.e., a total throughput of ≈ 9260 , so in this rough estimate our protocol $GF(2^8)$ -Circuit + Protocol 2 is factor ≈ 4.9 faster.

Table 5: The throughput in AES blocks per second for different network settings. We denote the best value for online and total throughput in bold. All maliciously secure protocols except $GF(2^8)$ -Circuit ([21] + [31]) have ≥ 37 bit statistical security. $GF(2^8)$ -Circuit ([21] + [31]) has ≥ 40 bit.

Protocol	Malicious	Throughput (blocks/s)									
		Online	Total	Online	Total	Online	Total	Online	Total	Online	Total
$GF(2^8)$ -Circuit [21]	✗	568 504	568 504	124 290	124 290	24 917	24 917	13 392	13 392	6 743	6 743
LUT-16	✗	775 697	318 498	180 606	99 656	35 063	20 174	16 617	10 840	11 052	6 739
$GF(2^4)$ -Circuit	✗	729 822	729 822	179 369	179 369	32 866	32 866	18 497	18 497	11 590	11 590
(2,3) LUT-256	✗	381 521	37 611	100 214	9 979	17 410	2 025	14 583	1 379	7 513	970
(3,3) LUT-256	✗	641 302	108 114	117 991	36 462	25 756	6 844	16 574	5 216	9 358	4 676
$GF(2^8)$ -Circuit ([21] + [31])	✓	152 127	6 044	32 185	2 703	7 889	1 481	4 707	1 039	2 832	743
$GF(2^8)$ -Circuit ([21] + Protocol 2)	✓	33 709	33 709	13 235	13 235	5 696	5 696	4 694	4 694	4 716	4 716
LUT-16 + Protocol 2	✓	50 189	44 883	19 113	17 344	6 472	5 707	5 273	4 481	5 844	4 365
$GF(2^4)$ -Circuit + Protocol 2	✓	48 924	48 924	18 829	18 829	6 429	6 429	5 396	5 396	6 073	6 073
(3,3) LUT-256 + Protocol 7	✓	32 508	25 409	13 547	10 560	3 732	2 652	3 370	2 345	5 056	3 289
Network		10 Gbit/s ≤ 1 ms RTT		1 Gbit/s ≤ 1 ms RTT		200 MBit/s 15ms RTT		100 MBit/s 30ms RTT		50 MBit/s 100ms RTT	
Batch size		250 000		250 000		100 000		100 000		100 000	

our protocols using the sublinear checks due to the increased computational cost. This advantage decreases as bandwidth decreases. In our studied settings, the bandwidth around 100 MBit/s is the tipping point where even the online phase of [31] (which can only be efficient due to the costly preprocessing) becomes slower than Protocol 2 due to the communication overhead. Further communication-saving techniques from the semi-honest setting only show a moderate effect of 6% to 28% improvement in throughput since Protocol 2's local computation becomes the main bottleneck.

4.3 Discussion, Strengths and Limitations

We conclude by discussing strengths and limitations of the presented protocols. They explore different trade-offs between local computation and communication, thus the concrete performance will depend on the computational and network resources in a specific scenario. This makes a generic statement with concrete numbers difficult, however we will sketch strengths and limitations in various settings in a qualitative way.

LUT-16. In settings with high bandwidth and low latency, this protocol performs best in terms of online phase throughput. However, for malicious security, previous work, $GF(2^8)$ -Circuit ([21] + [31]), is to be preferred to reach high online phase throughput. The necessary preprocessing and larger batch sizes limit LUT-16's utility for high total throughput and low computation latency, respectively.

$GF(2^4)$ -Circuit. This protocol is most effective in all network settings when a high total throughput is desired. Its comparatively simple structure without preprocessing makes it attractive for both semi-honest and malicious security. The protocol shares the downside of higher computation latency for each block with LUT-16 since it works more efficiently

on large batches.

(2,3) LUT-256. This protocol is a straight-forward variant of [40] adapted to replicated secret sharing. It does not have significant advantages over the improved variant (3,3) LUT-256 and shares its limitations. Although the online bandwidth cost is lower than (3,3) LUT-256, this is not reflected in its performance in our network settings due to the high cost of local computation.

(3,3) LUT-256. This protocol should be used to compute a small number of AES blocks in networks with latency when low computation latency is desired. Due to the higher local computation complexity, the throughput is limited when larger batches are needed.

Overall, both the LUT-16 and the $GF(2^4)$ -Circuit variant emerge as a valuable trade-off for scenarios with high throughput goals while the (3,3) LUT-256 protocol minimizes computation latency for a small batch size. Moreover, we closed the gap between semi-honest and malicious security in the WAN setting and achieve an overhead for malicious security of only 11% compared to the previous semi-honest state-of-the-art.

Acknowledgments

This work was partly supported by the Flemish Government through FWO SBO project MOZAIK S003321N, by CyberSecurity Research Flanders with reference number VR20192203, by the Danish Independent Research Council under Grant-ID DFF-0165-00107B (C3PO), by the Digital Research Center Denmark (DIREC), by the DIGIT Aarhus University Centre for Digitalisation, Big Data and Data Analytics, by JSPS KAKENHI Grant Number JP21H05052, and by JST CREST Grant Number JPMJCR22M1.

Ethics Considerations

This work does not involve any experiments involving humans, nor the collection or analysis of any real-world data. Therefore, the research and development tasks involved in carrying out the work have no direct ethical concerns. We can also consider the ethical implications of the results and general topic of our research. The primary use-case for our work is threshold cryptography, which aims to provide an extra layer of defence for cryptographic key storage, reducing the likelihood of key theft or loss of encrypted data. Our results share the ethical considerations of AES and encryption in general. Encryption, and in particular AES as used in Internet communication protocols increase privacy and security for sensitive data. This is important for individuals at risk of surveillance, e.g., activists under an authoritarian regime. Similarly, strong encryption can help to protect business information or intellectual property. On the flip side, encryption can help malicious actors to hide their nefarious activities from authorities. However, we strongly believe that the benefits of encryption strongly outweigh the abuse potential. It is possible that the techniques we develop will have secondary applications in other areas of secure multi-party computation.

Open Science

All artifacts (code of the implemented protocols and raw benchmark data of the results in Sect. 4) are publicly available at <https://doi.org/10.5281/zenodo.14719154> under MIT License and at <https://github.com/KULeuven-COSIC/maestro>.

References

- [1] Aysajan Abidin, Enzo Marquet, Jerico Moeyersons, Xhulio Limani, Erik Pohle, Michiel Van Kenhove, Johann M. Márquez-Barja, Nina Slamník-Krijestorac, and Bruno Volckaert. MOZAIK: an end-to-end secure data sharing platform. In *DEC 2023*, 2023.
- [2] Damiano Abram, Ivan Damgård, Peter Scholl, and Sven Trieflinger. Oblivious TLS via multi-party computation. In *CT-RSA 2021*, May 2021.
- [3] Martin R. Albrecht, Lorenzo Grassi, Christian Rechberger, Arnab Roy, and Tyge Tiessen. MiMC: Efficient encryption and cryptographic hashing with minimal multiplicative complexity. In *ASIACRYPT 2016, Part I*, December 2016.
- [4] Martin R. Albrecht, Christian Rechberger, Thomas Schneider, Tyge Tiessen, and Michael Zohner. Ciphers for MPC and FHE. In *EUROCRYPT 2015, Part I*, April 2015.
- [5] Toshinori Araki, Assi Barak, Jun Furukawa, Tamar Lichter, Yehuda Lindell, Ariel Nof, Kazuma Ohara, Adi Watzman, and Or Weinstein. Optimized honest-majority MPC for malicious adversaries - breaking the 1 billion-gate per second barrier. In *2017 IEEE Symposium on Security and Privacy*, May 2017.
- [6] Toshinori Araki, Jun Furukawa, Yehuda Lindell, Ariel Nof, and Kazuma Ohara. High-throughput semi-honest secure three-party computation with an honest majority. In *ACM CCS 2016*, October 2016.
- [7] Nuttapon Attrapadung, Hiraku Morita, Kazuma Ohara, Jacob C. N. Schuldt, and Kazunari Tozawa. Memory and round-efficient MPC primitives in the pre-processing model from unit vectorization. In *ASIACCS 22*, May / June 2022.
- [8] Carsten Baum, Ward Beullens, Shibam Mukherjee, Emanuela Orsini, Sebastian Ramacher, Christian Rechberger, Lawrence Roy, and Peter Scholl. One tree to rule them all: Optimizing GGM trees and OWFs for post-quantum signatures. In *ASIACRYPT 2024, Part I*, December 2024.
- [9] Carsten Baum, Tore Kasper Frederiksen, Julia Hesse, Anja Lehmann, and Avishay Yanai. PESTO: proactively secure distributed single sign-on, or how to trust a hacked server. In *EuroS&P*, 2020.
- [10] Amit Singh Bhati, Erik Pohle, Aysajan Abidin, Elena Andreeva, and Bart Preneel. Let’s go eevvee! A friendly and suitable family of AEAD modes for IoT-to-cloud secure computation. In *ACM CCS 2023*, November 2023.
- [11] Dan Boneh, Elette Boyle, Henry Corrigan-Gibbs, Niv Gilboa, and Yuval Ishai. Zero-knowledge proofs on secret-shared data via fully linear PCPs. In *CRYPTO 2019, Part III*, August 2019.
- [12] Joan Boyar, Philip Matthews, and René Peralta. Logic minimization techniques with applications to cryptology. *Journal of Cryptology*, (2), April 2013.
- [13] Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing: Improvements and extensions. In *ACM CCS 2016*, October 2016.
- [14] Elette Boyle, Niv Gilboa, and Yuval Ishai. Secure computation with preprocessing via function secret sharing. In *TCC 2019, Part I*, December 2019.
- [15] Elette Boyle, Niv Gilboa, Yuval Ishai, and Ariel Nof. Practical fully secure three-party computation via sub-linear distributed zero-knowledge proofs. In *ACM CCS 2019*, November 2019.

- [16] Elette Boyle, Niv Gilboa, Yuval Ishai, and Ariel Nof. Efficient fully secure computation via distributed zero-knowledge proofs. In *ASIACRYPT 2020, Part III*, December 2020.
- [17] Luís T. A. N. Brandão, Nicolas Christin, George Danezis, and Anonymous. Toward mending two nation-scale brokered identification systems. *PoPETs*, (2), April 2015.
- [18] A. Brüggemann, R. Hundt, T. Schneider, A. Suresh, and H. Yalame. Flute: Fast and secure lookup table evaluations. In *2023 IEEE Symposium on Security and Privacy (SP)*, may 2023.
- [19] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, October 2001.
- [20] D. Canright. A very compact S-box for AES. In *CHES 2005*, August / September 2005.
- [21] Koji Chida, Koki Hamada, Dai Ikarashi, Ryo Kikuchi, and Benny Pinkas. High-throughput secure AES computation. In *WAHC@CCS 2018*, 2018.
- [22] Ronald Cramer, Ivan Damgård, and Yuval Ishai. Share conversion, pseudorandom secret-sharing and applications to secure computation. In *TCC 2005*, February 2005.
- [23] Ivan Damgård and Marcel Keller. Secure multiparty AES. In *FC 2010*, January 2010.
- [24] Ivan Damgård, Marcel Keller, Enrique Larraia, Christian Miles, and Nigel P. Smart. Implementing AES via an actively/covertly secure dishonest-majority MPC protocol. In *SCN 12*, September 2012.
- [25] Ivan Damgård, Jesper Buus Nielsen, Michael Nielsen, and Samuel Ranellucci. The TinyTable protocol for 2-party secure computation, or: Gate-scrambling revisited. In *CRYPTO 2017, Part I*, August 2017.
- [26] Ivan Damgård and Rasmus Winther Zakarias. Fast oblivious AES A dedicated application of the MiniMac protocol. In *AFRICACRYPT 16*, April 2016.
- [27] F. Betül Durak and Jorge Guajardo. Improving the efficiency of AES protocols in multi-party computation. In *FC 2021, Part I*, March 2021.
- [28] Morris Dworkin, Elaine Barker, James Nechvatal, James Foti, Lawrence Bassham, E. Roback, and James Dray. Advanced encryption standard (aes), 2001-11-26 2001.
- [29] J.L. Fan and C. Paar. On efficient inversion in tower fields of characteristic two. In *IEEE International Symposium on Information Theory*, 1997.
- [30] Michael J. Freedman, Yuval Ishai, Benny Pinkas, and Omer Reingold. Keyword search and oblivious pseudo-random functions. In *TCC 2005*, February 2005.
- [31] Jun Furukawa, Yehuda Lindell, Ariel Nof, and Or Weinstein. High-throughput secure three-party computation for malicious adversaries and an honest majority. In *EUROCRYPT 2017, Part II*, April / May 2017.
- [32] Vipul Goyal and Yifan Song. Malicious security comes free in honest-majority MPC. Cryptology ePrint Archive, Report 2020/134, 2020.
- [33] Vipul Goyal, Yifan Song, and Chenzhi Zhu. Guaranteed output delivery comes free in honest majority MPC. In *CRYPTO 2020, Part II*, August 2020.
- [34] Lorenzo Grassi, Christian Rechberger, Dragos Rotaru, Peter Scholl, and Nigel P. Smart. MPC-friendly symmetric key primitives. In *ACM CCS 2016*, October 2016.
- [35] Shay Gueron, Yehuda Lindell, Ariel Nof, and Benny Pinkas. Fast garbling of circuits under standard assumptions. In *ACM CCS 2015*, October 2015.
- [36] Carmit Hazay and Yehuda Lindell. Efficient protocols for set intersection and pattern matching with security against malicious and covert adversaries. In *TCC 2008*, March 2008.
- [37] Yan Huang, David Evans, Jonathan Katz, and Lior Malka. Faster secure two-party computation using garbled circuits. In *USENIX Security 2011*, August 2011.
- [38] Mitsuru Ito, Akira Saito, and Takao Nishizeki. Secret sharing scheme realizing general access structure. *Electronics and Communications in Japan (Part III: Fundamental Electronic Science)*, 72(9), 1989.
- [39] Toshiya Itoh and Shigeo Tsujii. A fast algorithm for computing multiplicative inverses in $\text{gf}(2^m)$ using normal bases. *Inf. Comput.*, (3), sep 1988.
- [40] Marcel Keller, Emmanuela Orsini, Dragos Rotaru, Peter Scholl, Eduardo Soria-Vazquez, and Srinivas Vivek. Faster secure multi-party computation of AES and DES using lookup tables. In *ACNS 2017*, July 2017.
- [41] Marcel Keller, Dragos Rotaru, Nigel P. Smart, and Tim Wood. Reducing communication channels in MPC. In *SCN 18*, September 2018.
- [42] John Launchbury, Iavor S. Diatchki, Thomas DuBuisson, and Andy Adams-Moran. Efficient lookup-table protocol in secure multiparty computation. In *ACM SIGPLAN*, 2012.

- [43] Sven Laur, Riivo Talviste, and Jan Willemson. From oblivious AES to efficient and secure database join in the multiparty setting. In *ACNS 2013*, June 2013.
- [44] Yun Li, Yufei Duan, Zhicong Huang, Cheng Hong, Chao Zhang, and Yifan Song. Efficient 3pc for binary circuits with application to maliciously-secure DNN inference. In *USENIX Security 2023*, 2023.
- [45] Payman Mohassel, Mike Rosulek, and Ye Zhang. Fast and secure three-party computation: The garbled circuit approach. In *ACM CCS 2015*, October 2015.
- [46] Sean Murphy and Matthew J. B. Robshaw. Essential algebraic structure within the AES. In *CRYPTO 2002*, August 2002.
- [47] Peter Sebastian Nordholt and Meilof Veeningen. Minimising communication in honest-majority MPC by batchwise multiplication verification. In *ACNS 2018*, July 2018.
- [48] Satsuya Ohata and Koji Nuida. Communication-efficient (client-aided) secure two-party protocols and its application. In *FC 2020*, February 2020.
- [49] Benny Pinkas, Thomas Schneider, Nigel P. Smart, and Stephen C. Williams. Secure two-party computation is practical. In *ASIACRYPT 2009*, December 2009.
- [50] Mike Rosulek and Lawrence Roy. Three halves make a whole? Beating the half-gates lower bound for garbled circuits. In *CRYPTO 2021, Part I*, August 2021.
- [51] Dragos Rotaru, Nigel P. Smart, and Martijn Stam. Modes of operation suitable for computing on encrypted data. *IACR Trans. Symm. Cryptol.*, (3), 2017.
- [52] Akashi Satoh, Sumio Morioka, Kohji Takano, and Seiji Munetoh. A compact Rijndael hardware architecture with S-box optimization. In *ASIACRYPT 2001*, December 2001.
- [53] Johannes Wolkerstorfer, Elisabeth Oswald, and Mario Lamberger. An ASIC implementation of the AES S-boxes. In *CT-RSA 2002*, February 2002.
- [54] Samee Zahur, Mike Rosulek, and David Evans. Two halves make a whole - reducing data transfer in garbled circuits using half gates. In *EUROCRYPT 2015, Part II*, April 2015.

A Random One-Hot Vector Protocol for Length 16

We propose a protocol to securely compute the ideal functionality $\mathcal{F}_{\text{RandOHV}}$ with length-16 output in Protocol 8. Compared

Protocol 8 Random One-hot Vector (RndOhv) with Length 16

Functionality: $(\{\llbracket r_i \rrbracket\}_i, \{\llbracket e_j^{(r)} \rrbracket\}_j) \leftarrow \mathcal{F}_{\text{RandOHV}}(\perp)$

Input: \perp

Output: Shared random boolean values $(\llbracket r_3 \rrbracket, \llbracket r_2 \rrbracket, \llbracket r_1 \rrbracket, \llbracket r_0 \rrbracket)$ for $r_0, r_1, r_2, r_3 \in \{0, 1\}$ and the corresponding shared one-hot vector $\llbracket e^{(r)} \rrbracket$ for $r = \sum_{i=0}^3 2^i r_i$, where $|e^{(r)}| = 16$

Subfunctionality: $\mathcal{F}_{\text{rand}}$

- 1: $\llbracket r_k \rrbracket \leftarrow \mathcal{F}_{\text{rand}}(\perp)$ for $k \in [0, 3]$
 - 2: $\llbracket r_i r_j \rrbracket \leftarrow \mathcal{F}_{\text{weakMult}}(\llbracket r_i \rrbracket, \llbracket r_j \rrbracket)$ for all $3 \geq i > j \geq 0$
 - 3: $\llbracket r_i r_j r_k \rrbracket \leftarrow \mathcal{F}_{\text{weakMult}}(\llbracket r_i \rrbracket, \llbracket r_j r_k \rrbracket)$ for all $3 \geq i > j > k \geq 0$
 - 4: $\llbracket r_3 r_2 r_1 r_0 \rrbracket \leftarrow \mathcal{F}_{\text{weakMult}}(\llbracket r_3 r_2 \rrbracket, \llbracket r_1 r_0 \rrbracket)$
 $\triangleright 2$ offline rounds, 11 bits
 - 5: Servers locally compute $\llbracket e_j^{(r)} \rrbracket$ from the shares of the products as in Eq.(4), for $j \in [0, 15]$
 - 6: Execute $\mathcal{F}_{\text{verify}}$ for the following multiplication triplets:
 1. $(\llbracket r_i \rrbracket, \llbracket r_j \rrbracket, \llbracket r_i r_j \rrbracket)$ for all $3 \geq i > j \geq 0$
 2. $(\llbracket r_i \rrbracket, \llbracket r_j r_k \rrbracket, \llbracket r_i r_j r_k \rrbracket)$ for all $3 \geq i > j > k \geq 0$
 3. $(\llbracket r_3 r_2 \rrbracket, \llbracket r_1 r_0 \rrbracket, \llbracket r_3 r_2 r_1 r_0 \rrbracket)$
 - 7: **return** $\{\llbracket r_i \rrbracket\}_{0 \leq i \leq 3}, \{\llbracket e_j^{(r)} \rrbracket\}_{0 \leq j \leq 15}$
-

to Protocol 5, it has the same communication complexity but fewer rounds. The fundamental idea is based on the two-party Unitv-prep protocol for secure random unit vectorization protocol proposed in [7]. However, our proposed approach differs in that it allows the multi-party setting and it outputs sharings that are suitable for our construction.

The idea behind the construction is based on the fact that for a single random bit b , the pair $(b \oplus 1, b)$ forms a one-hot vector of length 2. Additionally, two one-hot vectors of length t can be tensor-multiplied to generate a one-hot vector of length $2t$. In the proposed approach, the constructed one-hot vector $e^{(r)}$ satisfies the following equation

$$e_j^{(r)} = \bigwedge_{0 \leq i \leq 3} (j[i] \oplus r[i] \oplus 1), \quad (4)$$

where $j[i]$ (resp., $r[i]$) represents the i -th bit of $j \in \mathbb{Z}_{16}$ (resp., $r \in \mathbb{Z}_{16}$). Note that, by the distributive property, the terms on the right-hand side for any $j \in \mathbb{Z}_{16}$ can be expressed as the sum of partial products of $\{r_0, r_1, r_2, r_3\}$. The proposed protocol achieves the one-hot encoding by generating random shares of $r[i] \in \mathbb{F}_2$ and securely computing all their partial products using Eq. (4).

Lemma A.1 *The protocol RndOhv in Protocol 8 securely computes $\mathcal{F}_{\text{RandOHV}}$ for $k = 4$ with abort in the $\{\mathcal{F}_{\text{rand}}, \mathcal{F}_{\text{weakMult}}, \mathcal{F}_{\text{verify}}\}$ -hybrid model in the presence of a malicious adversary under the honest majority setting.*

Proof: Simulation of RndOhv. \mathcal{S} emulates $\mathcal{F}_{\text{rand}}$ and receives from \mathcal{A} the share $\llbracket r_3 \rrbracket^C, \dots, \llbracket r_0 \rrbracket^C$ held by corrupted parties. For 11 invocations of Mult, \mathcal{S} emulates $\mathcal{F}_{\text{weakMult}}$ and

sends \mathcal{A} corrupt parties' input shares $(\llbracket r_i \rrbracket^C, \llbracket r_j \rrbracket^C)$ for all $3 \geq i > j \geq 0$, $(\llbracket r_i \rrbracket^C, \llbracket r_j r_k \rrbracket^C)$ for all $3 \geq i > j > k \geq 0$, and $(\llbracket r_3 r_2 \rrbracket^C, \llbracket r_1 r_0 \rrbracket^C)$. \mathcal{S} receives from \mathcal{A} the pairs of the error and shares $(d_{ij}, \llbracket r_i r_j \rrbracket^C)$ for all $3 \geq i > j \geq 0$, $(d_{ijk}, \llbracket r_i r_j r_k \rrbracket^C)$ for all $3 \geq i > j > k \geq 0$, and $(d, \llbracket r_3 r_2 r_1 r_0 \rrbracket^C)$. \mathcal{S} computes $\llbracket e_j^{(r)} \rrbracket^C$ for $j \in [0, 15]$ using the above shares held by corrupted parties. \mathcal{S} sends $\llbracket r \rrbracket^C := \llbracket r_{k-1} \rrbracket^C \dots \llbracket r_0 \rrbracket^C$ and $\llbracket e^{(r)} \rrbracket^C := \llbracket e_0^{(r)} \rrbracket^C \dots \llbracket e_{15}^{(r)} \rrbracket^C$ to $\mathcal{F}_{\text{RandOHV}}$. \mathcal{S} emulates $\mathcal{F}_{\text{verify}}$ and sends \mathcal{A} the multiplication triples, $(\llbracket r_i \rrbracket^C, \llbracket r_j \rrbracket^C, \llbracket r_i r_j \rrbracket^C)$ for all $3 \geq i > j \geq 0$, $(\llbracket r_i \rrbracket^C, \llbracket r_j r_k \rrbracket^C, \llbracket r_i r_j r_k \rrbracket^C)$ for all $3 \geq i > j > k \geq 0$, and $(\llbracket r_3 r_2 \rrbracket^C, \llbracket r_1 r_0 \rrbracket^C, \llbracket r_3 r_2 r_1 r_0 \rrbracket^C)$. If there exists a non-zero error among d_{ij}, d_{ijk} or d , \mathcal{S} sets $b = \text{abort}$, and otherwise \mathcal{S} sets $b = \text{accept}$. If $b = \text{accept}$ and \mathcal{A} replies **continue**, \mathcal{S} proceeds to the next step. Otherwise, \mathcal{S} sends **abort** to $\mathcal{F}_{\text{RandOHV}}$ and aborts.

We now show that the ideal execution and the real execution are indistinguishable. The view of \mathcal{A} consists of the corrupt parties' input shares to $\mathcal{F}_{\text{weakMult}}$, which are computed the same since they are obtained through linear operations from $\llbracket r_3 \rrbracket^C, \dots, \llbracket r_0 \rrbracket^C$. We also show that the output shares of all parties are distributed the same in both executions. The corrupted parties output shares in the ideal world are computed the same way as those in the real world. For the honest parties' shares, they are determined by using $\llbracket e^{(r)} \rrbracket^C$ and $e^{(r)}$ in the ideal world conditioned on $\llbracket e^{(r)} \rrbracket^{\mathcal{H}}$ is a shared one-hot vector of r . In the real world, $\llbracket e^{(r)} \rrbracket^{\mathcal{H}}$ is computed from $\llbracket r_3 \rrbracket^{\mathcal{H}}, \llbracket r_2 \rrbracket^{\mathcal{H}}, \llbracket r_1 \rrbracket^{\mathcal{H}}, \llbracket r_0 \rrbracket^{\mathcal{H}}$ as defined in Eq. (4) which satisfies $e^{(r)}$ is a one-hot vector of r . Here, $\llbracket e^{(r)} \rrbracket^{\mathcal{H}}$ is distributed uniformly since it is computed via linear combination of corrupted parties' shares of output from $\mathcal{F}_{\text{weakMult}}$. \square

Reducing the number of multiplication checks. Instead of verifying all 11 AND gates separately, we observe that it suffices to check 2 multiplications over a sufficiently large extension field $GF(2^k) = GF(2)[X]/f(X)$. The first multiplication verifies all the pairwise products $r_i r_j$:

$$\begin{aligned} (r_0 + r_1 X + r_2 X^2) \cdot (r_3 + r_0 X^3 + r_1 X^6) = \\ r_0 r_3 + r_1 r_3 X + r_2 r_3 X^2 \\ + X^3(r_0 + r_0 r_1 X + r_0 r_2 X^2) \\ + X^6(r_0 r_1 + r_1 X + r_1 r_2 X^2). \end{aligned} \quad (5)$$

The second multiplication verifies the remaining products:

$$\begin{aligned} (r_0 r_1 + r_0 r_3 X + r_1 r_3 X^2) \cdot (r_2 + r_3 X^3 + r_2 r_3 X^6) = \\ r_0 r_1 r_2 + r_0 r_2 r_3 X + r_1 r_2 r_3 X^2 \\ + X^3(r_0 r_1 r_3 + r_0 r_3 X + r_1 r_3 X^2) \\ + X^6(r_0 r_1 r_2 r_3 + r_0 r_2 r_3 X + r_1 r_2 r_3 X^2). \end{aligned} \quad (6)$$

Lemma A.2 *If the checking equations Eq. (5) and (6) both hold, then the multiplication triples $(\llbracket r_i \rrbracket, \llbracket r_j \rrbracket, \llbracket r_i r_j \rrbracket)$ for all*

$3 \geq i > j \geq 0$, $(\llbracket r_i \rrbracket, \llbracket r_j r_k \rrbracket, \llbracket r_i r_j r_k \rrbracket)$ for all $3 \geq i > j > k \geq 0$ and $(\llbracket r_3 r_2 \rrbracket, \llbracket r_1 r_0 \rrbracket, \llbracket r_3 r_2 r_1 r_0 \rrbracket)$ are all correct.

Proof: Let $d_{ij} \in \mathbb{F}_2$ for all $3 \geq i > j \geq 0$ be the additive errors introduced by \mathcal{A} in $\mathcal{F}_{\text{weakMult}}$ for all pairwise products. Then we denote $d_{ijk} \in \mathbb{F}_2$ the additive error from $\llbracket r_i r_j r_k \oplus d_{ijk} \rrbracket \leftarrow \mathcal{F}_{\text{weakMult}}(\llbracket r_i \rrbracket, \llbracket r_j r_k \oplus d_{jk} \rrbracket)$, for all $3 \geq i > j > k \geq 0$, and further we denote $d \in \mathbb{F}_2$ the additive error from $\llbracket r_3 r_2 r_1 r_0 \oplus d \rrbracket \leftarrow \mathcal{F}_{\text{weakMult}}(\llbracket r_3 r_2 \oplus d_{32}, r_1 \oplus d_{10} \rrbracket)$.

For the check with reduced number of multiplication checks, the parties interact with $\mathcal{F}_{\text{verify}}$ which now accepts triples in $GF(2^p)$, $p \geq 9$. The parties compute shares of $t_1, t_2, t_3, h_1, h_2, h_3 \in GF(2^p)$ as

$$\begin{aligned} \llbracket t_1 \rrbracket &:= \llbracket r_0 \rrbracket + \llbracket r_1 \rrbracket X + \llbracket r_2 \rrbracket X^2, \\ \llbracket t_2 \rrbracket &:= \llbracket r_3 \rrbracket + \llbracket r_0 \rrbracket X + \llbracket r_1 \rrbracket X^6, \\ \llbracket t_3 \rrbracket &:= \llbracket r_0 r_3 \rrbracket + \llbracket r_1 r_3 \rrbracket X + \llbracket r_2 r_3 \rrbracket X^2 \\ &\quad + X^3(\llbracket r_0 \rrbracket + \llbracket r_1 r_0 \oplus d_{10} \rrbracket X + \llbracket r_2 r_0 \oplus d_{20} \rrbracket X^2) \\ &\quad + X^6(\llbracket r_1 r_0 \oplus d_{10} \rrbracket + \llbracket r_1 \rrbracket X + \llbracket r_2 r_1 \oplus d_{21} \rrbracket X^2), \end{aligned}$$

and

$$\begin{aligned} \llbracket h_1 \rrbracket &:= \llbracket r_1 r_0 \oplus d_{10} \rrbracket + \llbracket r_3 r_0 \oplus d_{30} \rrbracket X + \llbracket r_3 r_1 \oplus d_{31} \rrbracket X^2, \\ \llbracket h_2 \rrbracket &:= \llbracket r_2 \rrbracket + \llbracket r_3 \rrbracket X^3 + \llbracket r_3 r_2 \oplus d_{32} \rrbracket X^6, \\ \llbracket h_3 \rrbracket &:= \llbracket r_2 r_1 r_0 \oplus d_{210} \rrbracket + \llbracket r_3 r_2 r_0 \oplus d_{320} \rrbracket X + \llbracket r_3 r_2 r_1 \oplus d_{321} \rrbracket X^2 \\ &\quad + X^3(\llbracket r_3 r_1 r_0 \oplus d_{310} \rrbracket + \llbracket r_3 r_0 \oplus d_{30} \rrbracket X + \llbracket r_3 r_1 \oplus d_{31} \rrbracket X^2) \\ &\quad + X^6(\llbracket r_3 r_2 r_1 r_0 \oplus d \rrbracket + \llbracket r_3 r_2 r_0 \oplus d_{320} \rrbracket X + \llbracket r_3 r_2 r_1 \oplus d_{321} \rrbracket X^2). \end{aligned}$$

The parties send $(\llbracket t_1 \rrbracket, \llbracket t_2 \rrbracket, \llbracket t_3 \rrbracket)$ and $(\llbracket h_1 \rrbracket, \llbracket h_2 \rrbracket, \llbracket h_3 \rrbracket)$ to $\mathcal{F}_{\text{verify}}$ which checks $t_1 \cdot t_2 = t_3$ and $h_1 \cdot h_2 = h_3$, thus obtains

$$\begin{aligned} 0 &= d_{30} + d_{31} X + d_{32} X^2 + d_{10} X^4 + d_{20} X^5 + d_{10} X^6 + d_{21} X^8, \\ 0 &= d_{210} + d_{320} X + d_{321} X^2 + d_{310} X^3 + d X^6 + d_{320} X^7 + d_{321} X^8, \end{aligned}$$

respectively. Thus, $\mathcal{F}_{\text{verify}}$ accepts if and only if all errors d_{ij}, d_{ijk} and d are zero. \square

B Deferred Proofs

B.1 Proof of Lemma 3.1

Lemma 3.1 (restated) *The protocol Inv in Protocol 3 securely computes \mathcal{F}_{Inv} with abort in the $\{\mathcal{F}_{\text{LUT}}^{\langle \cdot \rangle \rightarrow \llbracket \cdot \rrbracket}, \mathcal{F}_{\text{weakMult}}, \mathcal{F}_{\text{verify}}\}$ -hybrid model in the presence of a malicious adversary under the honest majority setting.*

Proof: \mathcal{S} receives $\llbracket x \rrbracket^C$ from \mathcal{F}_{Inv} and computes $\llbracket v_x \rrbracket^C$ through Step 1–7 using $\llbracket x \rrbracket^C$. \mathcal{S} emulates $\mathcal{F}_{\text{LUT}}^{\langle \cdot \rangle \rightarrow \llbracket \cdot \rrbracket}$ and receives $\langle \langle v \rangle \rangle^C, \llbracket v^{-1} \rrbracket^C$ and $\llbracket v \rrbracket^C$ from \mathcal{A} , and defines the error $d_v = \llbracket v_x \rrbracket^C \oplus \llbracket v \rrbracket^C$. \mathcal{S} emulates $\mathcal{F}_{\text{weakMult}}$, computes $\llbracket a_h \rrbracket^C$ and $\llbracket a_h \rrbracket^C \oplus \llbracket a_\ell \rrbracket^C$ using $\llbracket x \rrbracket^C$, and sends

$(\llbracket a_h \rrbracket^C, \llbracket v^{-1} \rrbracket^C), (\llbracket a_h \rrbracket^C \oplus \llbracket a_\ell \rrbracket^C, \llbracket v^{-1} \rrbracket^C)$ to \mathcal{A} . \mathcal{S} receives $(d_1, \llbracket a'_h \rrbracket^C), (d_2, \llbracket a'_\ell \rrbracket^C)$ from \mathcal{A} . \mathcal{S} emulates $\mathcal{F}_{\text{verify}}$ and sends \mathcal{A} the multiplication triples $(\llbracket a_h \rrbracket^C, \llbracket a_\ell \rrbracket^C, \llbracket a_h \times a_\ell \rrbracket^C), (\llbracket a_h \rrbracket^C, \llbracket v^{-1} \rrbracket^C, \llbracket a'_h \rrbracket^C)$ and $(\llbracket a_h \oplus a_\ell \rrbracket^C, \llbracket v^{-1} \rrbracket^C, \llbracket a'_\ell \rrbracket^C)$. Note that the first triple $(\llbracket a_h \rrbracket, \llbracket a_\ell \rrbracket, \llbracket a_h \times a_\ell \rrbracket)$ can indirectly prove that $\mathcal{F}_{\text{LUT}}^{\langle \cdot \rangle \rightarrow \llbracket \cdot \rrbracket}$ outputs the correct $\llbracket v \rrbracket$, that is, c is correctly computed. Here, $\llbracket v^{-1} \rrbracket$ is computed locally and we don't need to verify.

If there exists a non-zero error among d_v, d_1 or d_2 , \mathcal{S} sets $b = \mathbf{abort}$, and otherwise \mathcal{S} sets $b = \mathbf{accept}$. If $b = \mathbf{accept}$ and if \mathcal{A} replies **continue**, \mathcal{S} proceeds to the next step. Otherwise, \mathcal{S} sends **abort** to \mathcal{F}_{Inv} and aborts.

\mathcal{S} computes $\llbracket x^{-1} \rrbracket^C$ using $\llbracket a'_h \rrbracket^C$ and $\llbracket a'_\ell \rrbracket^C$. \mathcal{S} sends $\llbracket x^{-1} \rrbracket^C$ to \mathcal{F}_{Inv} .

We state why the ideal execution is indistinguishable from the real execution. The view of the adversary consists of the corrupt parties' shares of the first inputs to multiplicative inversions $\llbracket a_h \rrbracket$ and $\llbracket a_h \oplus a_\ell \rrbracket$, but these are uniformly random in both executions because they are obtained by applying a non-zero affine map to $\llbracket x \rrbracket^C$. We also need to show that the output shares of all parties are distributed the same in both executions. The corrupted parties' output shares are the same in both executions. For the honest parties' output shares in the ideal execution, they are sampled at random conditioned on that they can be reconstructed to x^{-1} . In the real execution, the honest parties' output share $\llbracket y \rrbracket^{\mathcal{H}}$ is obtained as $\llbracket x^{-1} \rrbracket^{\mathcal{H}}$ as the correctness was shown in Sect. 2.1, and it is uniformly distributed since it is computed by applying a non-zero affine map to the output of multiplications that were sampled uniformly. \square

B.2 Proof of Lemma 3.3

Lemma 3.3 (restated) *The protocol LUT in Protocol 4 securely computes $\mathcal{F}_{\text{LUT}}^{\llbracket \cdot \rrbracket \rightarrow \langle \cdot \rangle}$ in the $\{\mathcal{F}_{\text{RandOHV}}, \mathcal{F}_{\text{zero}}\}$ -hybrid model in the presence of a malicious adversary.*

Proof: Let \mathcal{A} denote the adversary. We will construct a simulator, \mathcal{S} , to simulate the honest parties' behaviour in the real execution.

Simulation of LUT. \mathcal{S} emulates $\mathcal{F}_{\text{RandOHV}}$ and receives $\llbracket r \rrbracket^C$ and $\llbracket e_j^{(r)} \rrbracket^C$ from \mathcal{A} , and defines $\langle \langle r \rangle \rangle^C = \text{ToAdditive}(\llbracket r \rrbracket^C)$. \mathcal{S} emulates $\mathcal{F}_{\text{zero}}$ and receives $\llbracket 0 \rrbracket^C$ from \mathcal{A} . \mathcal{S} receives $\langle \langle c \rangle \rangle^C$ from the adversary and computes $\langle \langle v \rangle \rangle^C = \langle \langle c \rangle \rangle^C \oplus \langle \langle r \rangle \rangle^C \oplus \langle \langle 0 \rangle \rangle^C$. \mathcal{S} samples at random a set of shares $\langle \langle c \rangle \rangle^{\mathcal{H}}$ held by honest parties and sends them to the adversary. \mathcal{S} computes the set of shares $\llbracket t \rrbracket^C := \bigoplus_{j=0}^{n-1} \llbracket e_j^{(r)} \rrbracket^C \cdot T_{c \oplus j}$ and $\llbracket v \rrbracket^C := \llbracket r \rrbracket^C \oplus c$. \mathcal{S} sends to $\mathcal{F}_{\text{LUT}}^{\langle \cdot \rangle \rightarrow \llbracket \cdot \rrbracket}$ the corrupted parties' input shares $\langle \langle v \rangle \rangle^C$, and the outputs shares $\llbracket t \rrbracket^C, \llbracket v \rrbracket^C$.

We now argue why the ideal execution is indistinguishable from the real execution. The view of the adversary consists of the honest parties' shares of the reconstructed c , but

these are uniformly random in both executions, thanks to the masking with $\langle \langle 0 \rangle \rangle$. We also need to show that the output shares of all parties are distributed the same in both worlds. The corrupted parties' shares are computed exactly the same way in both executions. For the honest parties' shares, in the ideal world they are sampled at random conditioned on $t = T_v$. In the real protocol, since $e_j^{(r)}$ has a 1 in position r , we have $t = T_{v \oplus r \oplus r} = T_v$. Furthermore, since the shares $\llbracket e^{(r)} \rrbracket^{\mathcal{H}}, \llbracket r \rrbracket^{\mathcal{H}}$ sampled by $\mathcal{F}_{\text{RandOHV}}$ are sampled uniformly, the output shares $\llbracket t \rrbracket^{\mathcal{H}}, \llbracket v \rrbracket^{\mathcal{H}}$ are also uniformly distributed, since they are each obtained by applying a non-zero affine map to $\llbracket e^{(r)} \rrbracket^{\mathcal{H}}$ and $\llbracket r \rrbracket^{\mathcal{H}}$. \square

B.3 Proof of Lemma 3.5

Lemma 3.5 (restated) *Protocol 7 securely realizes the functionality \mathcal{F}_{AES} in the $(\mathcal{F}_{\text{LUT}}^{\langle \cdot \rangle}, \mathcal{F}_{\text{verify}})$ -hybrid model with malicious security.*

Proof: We construct a simulator, \mathcal{S} , as follows. First, \mathcal{S} receives from \mathcal{F}_{AES} the corrupted parties' shares $\llbracket x \rrbracket^C$ and $\llbracket k^{(i)} \rrbracket^C$, and defines $\langle \langle k^{(i)} \rangle \rangle^C = \text{ToAdditive}(\llbracket k^{(i)} \rrbracket^C)$. For the first set of calls to $\mathcal{F}_{\text{LUT}}^{\llbracket \cdot \rrbracket \rightarrow \langle \cdot \rangle}$, \mathcal{S} sends to \mathcal{A} the appropriate shares of x_b^0 , and receives the output shares $\langle \langle y_b^0 \rangle \rangle^C$. For each subsequent round, for $i = 1, \dots, 9$, \mathcal{S} does as follows:

- Apply the linear layer L_i to the corrupted parties' shares to obtain $\langle \langle x^i \rangle \rangle^C$.
- Receive $\langle \langle x^i \rangle \rangle^C$ from \mathcal{A} , as input to $\mathcal{F}_{\text{LUT}}^{\langle \cdot \rangle}$.
- Define the error $\delta^i = \bigoplus_{j \in C} (\langle \langle x^i \rangle \rangle^j \oplus \langle \langle x^i \rangle \rangle^j)$.
- Receive the adversary's output sharings $\langle \langle y^i \rangle \rangle^C, \llbracket x_b^i \rrbracket^C$ for $\mathcal{F}_{\text{LUT}}^{\langle \cdot \rangle}$.
- Compute the shares $\llbracket y^{i-1} \rrbracket^C$ according to the protocol.

Finally, \mathcal{S} emulates Reshare, and computes the corresponding error δ^{10} in the new sharing of y^9 . To emulate $\mathcal{F}_{\text{verify}}$, \mathcal{S} first sends to \mathcal{A} the corresponding shares of the triples. Then, if any δ^i is non-zero, \mathcal{S} sends **abort** to \mathcal{A} and aborts; otherwise, \mathcal{S} sends **accept**, and if \mathcal{A} responds with **continue**, \mathcal{S} sends to \mathcal{F}_{AES} the corrupted parties' shares of the outputs, z .

We claim that the ideal execution is distributed identically to that of the real execution. Note that the real protocol aborts if any of the S-box input/output pairs are incorrect; otherwise, the output z must be the result of a correct AES evaluation. In the ideal execution, the protocol aborts if any error δ^i is non-zero. Since δ^i is the sum of all corrupt parties' shares of the x^i value which was meant to be input into \mathcal{F}_{LUT} , and x^i was used as input, any non-zero δ^i means that an incorrect S-box input was used in round i . Since this input is used to derive the sharings $\llbracket x^i \rrbracket$, this will cause the S-box verification for round $i - 1$ to fail, and the protocol will abort. \square

C Details on AES

C.1 Encryption Algorithm for Each Block

The encryption of a block in AES is a deterministic algorithm that takes a 128-bit array and an encryption key as input and produces a 128-bit array as output. The algorithm can be defined as a function $\text{Enc} : \{0, 1\}^{128} \times \{0, 1\}^{128} \rightarrow \{0, 1\}^{128}$.

The protocol proceeds as follows.

1. **Initialization:** The input values are divided into 8-bit segments, each of which is considered an element of $GF(2^8)$. These elements are represented in a 4×4 array with column-first order, denoted as $\{s_{r,c}\}_{0 \leq r,c \leq 3}$. The same process is applied to the encryption key.
2. **Key Expansion:** The round keys $\{k_{r,c}^{(i)}\}_{0 \leq r,c \leq 3, 0 \leq i \leq 10}$ to be used in each round $i \in \{1, \dots, 10\}$ are generated from the encryption keys $\{k_{r,c}\}_{0 \leq r,c \leq 3}$. This is done as follows:

$$k_{r,c}^{(i)} := \begin{cases} k_{r,c} & \text{if } i = 0, \\ k_{r,0}^{(i-1)} \oplus \text{Sbox}(k_{(r+1 \bmod 4),3}^{(i-1)}) \oplus rc_r^{(i)} & \text{if } i \neq 0, c = 0, \\ k_{r,c}^{(i-1)} \oplus k_{r,c-1}^{(i)} & \text{otherwise.} \end{cases}$$

Here, $rc_r^{(i)} \in GF(2^8)$ is defined as $rc_0^{(i)} := (\{02\}_{16})^{i-1}$ and $rc_r^{(i)} := \{00\}_{16}$ for $1 \leq r \leq 3$. Sbox represents a substitution according to a predefined table (Section C.2).

3. Each element of the array is computed as $s_{r,c} := s_{r,c} \oplus k_{r,c}^{(0)}$.
4. **Round Processing:** For $i = 1, \dots, 10$, the following steps are repeated:

- **SubBytes:** Each element of the array is substituted according to a predefined table (AES S-box):

$$s_{r,c} := \text{Sbox}(s_{r,c}).$$

- **ShiftRows:** Each row is shifted according to the following rule: $s_{r,c} := s_{r,(c+r \bmod 4)}$.
- **MixColumns:** For each column, the following are calculated:

$$\begin{pmatrix} s_{0,c} \\ s_{1,c} \\ s_{2,c} \\ s_{3,c} \end{pmatrix} := \begin{pmatrix} \{02\}_{16} & \{03\}_{16} & \{01\}_{16} & \{01\}_{16} \\ \{01\}_{16} & \{02\}_{16} & \{03\}_{16} & \{01\}_{16} \\ \{01\}_{16} & \{01\}_{16} & \{02\}_{16} & \{03\}_{16} \\ \{03\}_{16} & \{01\}_{16} & \{01\}_{16} & \{02\}_{16} \end{pmatrix} \begin{pmatrix} s_{0,c} \\ s_{1,c} \\ s_{2,c} \\ s_{3,c} \end{pmatrix}$$

Note that this step is omitted when $i = 10$.

- **AddRoundKey:** Each element of the array is XORed with the round key: $s_{r,c} := s_{r,c} \oplus k_{r,c}^{(i)}$.

5. **Finalization:** The 4×4 array $\{s_{r,c}\}$ is concatenated in column-first order to produce the output.

C.2 AES S-Box

The AES S-Box is a substitution table used in the key expansion and SubBytes step to ensure the non-linearity of encryption. Specifically, for an input $s \in GF(2^8)$, it produces an output $\{a_7 \dots a_0\}_2 \in GF(2^8)$ defined as follows:

$$\begin{aligned} a_0 &:= b_0 \oplus b_4 \oplus b_5 \oplus b_6 \oplus b_7 \oplus 1, \\ a_1 &:= b_0 \oplus b_1 \oplus b_5 \oplus b_6 \oplus b_7 \oplus 1, \\ a_2 &:= b_0 \oplus b_1 \oplus b_2 \oplus b_6 \oplus b_7, \\ a_3 &:= b_0 \oplus b_1 \oplus b_2 \oplus b_3 \oplus b_7, \\ a_4 &:= b_0 \oplus b_1 \oplus b_2 \oplus b_3 \oplus b_4, \\ a_5 &:= b_1 \oplus b_2 \oplus b_3 \oplus b_4 \oplus b_5 \oplus 1, \\ a_6 &:= b_2 \oplus b_3 \oplus b_4 \oplus b_5 \oplus b_6 \oplus 1, \\ a_7 &:= b_3 \oplus b_4 \oplus b_5 \oplus b_6 \oplus b_7. \end{aligned} \tag{7}$$

Here, $\{b_7 \dots b_0\}_2$ represents the bit sequence that denotes the *multiplicative inverse* $s^{-1} \in GF(2^8)$ of the input value $s \in GF(2^8)$. Note that when $s = \{00\}_{16}$, all b_i are set to 0 for all i .

C.3 Oblivious AES

Protocol 9 describes the whole AES algorithm that is computed in MPC. The main body of the paper focused on computing SubBytes.

Protocol 9 Oblivious AES

Functionality: $\{\llbracket z_i \rrbracket\} \leftarrow \mathcal{F}_{\text{AES}}(\{\llbracket x_i \rrbracket\}, \{\llbracket k_i \rrbracket\})$

Input: 128-bit Boolean shared values $\{\llbracket x_i \rrbracket\}_{i=0,\dots,127}$, $\{\llbracket k_i \rrbracket\}_{i=0,\dots,127}$

Output: 128-bit Boolean shared value $\{\llbracket z_i \rrbracket\}_{i=0,\dots,127}$

- 1: Servers locally perform initialization step to obtain $\{\llbracket x_{r,c} \rrbracket\}_{0 \leq r,c \leq 3}$ and $\{\llbracket k_{r,c} \rrbracket\}_{0 \leq r,c \leq 3}$
 - 2: $\{\llbracket k_{r,c}^{(i)} \rrbracket\}_{0 \leq i \leq 10, 0 \leq r,c \leq 3} \leftarrow \text{KeyExpansion}(\{\llbracket k_{r,c} \rrbracket\}_{0 \leq r,c \leq 3})$
▷ one-time operation for each key
 - 3: $\{\llbracket x_{r,c} \rrbracket\} \leftarrow \text{AddRoundKey}(\{\llbracket x_{r,c} \rrbracket\}, \{\llbracket k_{r,c}^{(0)} \rrbracket\})$
 - 4: **for** $i = 1, \dots, 9$ **do**
 - 5: $\{\llbracket x_{r,c} \rrbracket\} \leftarrow \text{SubBytes}(\{\llbracket x_{r,c} \rrbracket\})$
 - 6: $\{\llbracket x_{r,c} \rrbracket\} \leftarrow \text{ShiftRows}(\{\llbracket x_{r,c} \rrbracket\})$
 - 7: $\{\llbracket x_{r,c} \rrbracket\} \leftarrow \text{MixColumns}(\{\llbracket x_{r,c} \rrbracket\})$
 - 8: $\{\llbracket x_{r,c} \rrbracket\} \leftarrow \text{AddRoundKey}(\{\llbracket x_{r,c} \rrbracket\}, \{\llbracket k_{r,c}^{(i)} \rrbracket\})$
 - 9: **end for**
 - 10: $\{\llbracket x_{r,c} \rrbracket\} \leftarrow \text{SubBytes}(\{\llbracket x_{r,c} \rrbracket\})$
 - 11: $\{\llbracket x_{r,c} \rrbracket\} \leftarrow \text{ShiftRows}(\{\llbracket x_{r,c} \rrbracket\})$
 - 12: $\{\llbracket x_{r,c} \rrbracket\} \leftarrow \text{AddRoundKey}(\{\llbracket x_{r,c} \rrbracket\}, \{\llbracket k_{r,c}^{(10)} \rrbracket\})$
 - 13: Servers locally perform finalization step to obtain $\{\llbracket z_i \rrbracket\}_{i=0,\dots,127}$
 - 14: **return** $\{\llbracket z_i \rrbracket\}_{i=0,\dots,127}$
-

D Batch Verification Protocol

We now prove security of Protocol 2.

Theorem 2.3 (restated) *Protocol 2 (Verify) securely realizes the functionality $\mathcal{F}_{\text{verify}}$ (see Fig. 3), in the $(\mathcal{F}_{\text{weakMult}}, \mathcal{F}_{\text{coin}}, \mathcal{F}_{\text{rand}})$ -hybrid model. The failure probability in the simulation is at most $(m + 2 \log N)/|\mathbb{F}|$.*

Proof: We begin by proving the base case when $N = 1$, namely, Protocol 1.

Proposition D.1 *Protocol 1 securely realizes $\mathcal{F}_{\text{verify}}$ for a single multiplication triple.*

Proof: The simulator, \mathcal{S} , receives the corrupted parties' shares $\llbracket x \rrbracket^C, \llbracket y \rrbracket^C, \llbracket z \rrbracket^C$ from $\mathcal{F}_{\text{verify}}$, and receives the outcome $b \in \{\text{accept}, \text{abort}\}$. It emulates $\mathcal{F}_{\text{rand}}$ and $\mathcal{F}_{\text{weakMult}}$, receiving shares $\llbracket x' \rrbracket^C, \llbracket r \rrbracket^C$ and $\llbracket z' \rrbracket^C$, plus an additive error d , and then sends a random $t \leftarrow \mathbb{F}$ for $\mathcal{F}_{\text{coin}}$. It simulates the first Reconst by sending a random value ρ . For the second $\mathcal{F}_{\text{weakMult}}$, it receives the shares $\llbracket \sigma \rrbracket^C$, together with another additive error f . For the second Reconst, if $b = \text{accept}$, $d = 0$ and $f = 0$ then \mathcal{S} opens σ by sending honest shares corresponding to the secret 0, and sends **continue** to $\mathcal{F}_{\text{verify}}$. Otherwise, it samples honest shares corresponding to a random σ , sends these to the adversary and sends **abort** to $\mathcal{F}_{\text{verify}}$.

First, notice that in the real world, ρ is statistically close to uniform, because of the tx' term that masks x . Secondly, if $e = z - xy$ is the error in the triple, then we have

$$z + tz' - \rho y = z - xy + t(z' - x'y) = e + td.$$

So, if the triple is correct and the adversary chooses $d = 0, f = 0$ then we always have $\sigma = 0$ in both real and ideal worlds, and furthermore the output in both cases will be **accept**. On the other hand, if there are any errors then in the real world we have $\sigma = (e + td)r + f$. If e is non-zero, then $e + td$ is non-zero except with probability $1/|\mathbb{F}|$, since d is fixed before the sampling of t . It follows that σ is statistically close to uniform, since r is uniformly random and unknown to the adversary. Finally, this implies that the protocol output will be **abort** except with probability $1/|\mathbb{F}|$, which is statistically close to the ideal world. \square

Next, we analyze the VerifyDotProduct procedure with an inductive argument. Namely, we show that if the recursive call to VerifyDotProduct of length $N/2$ securely realizes the functionality, then so does the main procedure.

Proposition D.2 *Suppose that VerifyDotProduct on input of an inner product triple of length $N/2$ securely implements $\mathcal{F}_{\text{verify}}$, with $m = 1$ and length $N/2$. Then, VerifyDotProduct securely implements $\mathcal{F}_{\text{verify}}$ with $m = 1$ and length N . The failure probability in the simulation is $2/|\mathbb{F}|$.*

Proof: We construct a simulator, \mathcal{S} , as follows. \mathcal{S} receives from $\mathcal{F}_{\text{verify}}$ the corrupted shares of the triple $\llbracket \vec{x} \rrbracket^C, \llbracket \vec{y} \rrbracket^C, \llbracket z \rrbracket^C$, and the result $b \in \{\text{accept}, \text{abort}\}$. \mathcal{S} computes the shares of f_i and g_i , and uses these to simulate $\mathcal{F}_{\text{weakDotProduct}}$. It receives the adversary's shares $h(1), h(2)$, and locally computes the shares of $h(0)$, to define shares of the polynomial $\llbracket h(X) \rrbracket^C$. It also receives from \mathcal{A} errors, which define via interpolation an error polynomial $e(X)$ such that $h(X) = \tilde{f}(X) \cdot \tilde{g}(X) + e(X)$.

Next, \mathcal{S} sends a random $r \leftarrow \mathbb{F}$ to \mathcal{A} . It then emulates the recursive call to VerifyDotProduct; if $b = \text{abort}$ or $e(X) \neq 0$, it sends **abort** to the adversary, followed by **abort** to the length- N $\mathcal{F}_{\text{verify}}$. Otherwise, it sends **accept** to the adversary; if it responds with continue, then send **accept** to $\mathcal{F}_{\text{verify}}$, otherwise send **abort**.

We now argue indistinguishability. In the real world, if the inner VerifyDotProduct check succeeds then $h(r) = \tilde{f}(r) \cdot \tilde{g}(r)$. Since $h(X)$ is degree at most 2, this implies that $h(X)$ and $\tilde{f}(X) \cdot \tilde{g}(X)$ are equal as polynomials, except with probability $2/|\mathbb{F}|$. Since $h(0) + h(1) = z$, by construction, it follows that, except with negligible probability, if the protocol accepts then $z = \tilde{f}(0) \cdot \tilde{g}(0) + \tilde{f}(1) \cdot \tilde{g}(1) = \vec{x} \cdot \vec{y}$, as required.

Meanwhile, in the ideal world, the simulator always aborts if the triple is incorrect, or if $e(X) \neq 0$. The only possible differences between the two worlds are the cases: (i) the triple is correct, but $e(X) \neq 0$ and VerifyDotProduct accepts, or (ii) the triple is incorrect, but the real protocol accepts. Each of these cases would require VerifyDotProduct to accept in the real world, even though $h(X) \neq \tilde{f}(X) \cdot \tilde{g}(X)$. As argued above, this happens with probability at most $2/|\mathbb{F}|$. \square

Proposition D.3 *If at least one triple input to Protocol 2 is incorrect, then so is the input to VerifyDotProduct, except with probability at most $m/|\mathbb{F}|$.*

Proof: Suppose that $z_i = \vec{x}_i \cdot \vec{y}_i + \delta_i$, and at least one $\delta_i \neq 0$. Then, if

$$(\vec{x}_0, r\vec{x}_1, \dots, r^{m-1}\vec{x}_{m-1}) \cdot (\vec{y}_0, \dots, \vec{y}_{m-1}) = \sum_{i=0}^{m-1} r^i z_i,$$

then it holds that

$$(1, r, \dots, r^{m-1}) \cdot (\delta_0, \dots, \delta_{m-1}) = 0.$$

Viewing the δ_i 's as coefficients of a non-zero, degree $m-1$ polynomial, this holds with probability at most $m/|\mathbb{F}|$, for a random r . \square

The claim and final bound in the theorem follows by a hybrid argument over the $\log N$ recursive calls to VerifyDotProduct and the final base case. \square

E Improved Protocol for $\mathcal{F}_{\text{LUT}}^{\llbracket \cdot \rrbracket}$

We present the improved protocol for $\mathcal{F}_{\text{LUT}}^{\llbracket \cdot \rrbracket}$ in Protocol 10. Since the protocol operates entirely on $\llbracket \cdot \rrbracket$ -shared data, its

security is straightforward and we omit the proof of the following.

Lemma E.1 *Protocol 10 securely realizes the functionality $\mathcal{F}_{\text{LUT}}^{\llbracket \cdot \rrbracket}$ with malicious security.*

Protocol 10 Table Lookup of size $N = 2^k$ in replicated sharing

Functionality: $\llbracket T_v \rrbracket \leftarrow \mathcal{F}_{\text{LUT}}(\llbracket v \rrbracket, T)$

Input: Share $\llbracket v \rrbracket$ of $v \in GF(2^k)$, table $T : GF(2^k) \rightarrow GF(2^\ell)$

Output: Share $\llbracket T_v \rrbracket$ of the value of T at v

Subfunctionality: $\mathcal{F}_{\text{RandOHV}}$

- 1: Call $\mathcal{F}_{\text{RandOHV}}(k/2)$ twice to get $(\{\llbracket r_i \rrbracket\}_{i=0}^{k/2-1}, \{\llbracket e_j^{(r)} \rrbracket\}_{j=0}^{\sqrt{N}-1})$ and $(\{\llbracket r'_i \rrbracket\}_{i=0}^{k/2-1}, \{\llbracket e_j^{(r')} \rrbracket\}_{j=0}^{\sqrt{N}-1})$
 - 2: $\llbracket r \rrbracket := (\llbracket r_0 \rrbracket, \dots, \llbracket r_{k/2-1} \rrbracket, \llbracket r'_0 \rrbracket, \dots, \llbracket r'_{k/2-1} \rrbracket)$
 - 3: $c \leftarrow \text{Reconst}(\llbracket v \rrbracket + \llbracket r \rrbracket) \triangleright 1 \text{ round, } k \text{ bits}$
 $\triangleright \vec{f}^{(0)}, \vec{f}^{(1)} \in \{0, 1\}^N$
 - 4: $\llbracket \vec{f}^{(0)} \rrbracket := (\llbracket e_0^{(r)} \rrbracket, \dots, \llbracket e_{\sqrt{N}-1}^{(r)} \rrbracket, \dots, \llbracket e_{\sqrt{N}-1}^{(r)} \rrbracket)$
 - 5: $\llbracket \vec{f}^{(1)} \rrbracket := (\llbracket e_0^{(r')} \rrbracket, \dots, \llbracket e_{\sqrt{N}-1}^{(r')} \rrbracket, \dots, \llbracket e_{\sqrt{N}-1}^{(r')} \rrbracket)$
 - 6: $\llbracket \vec{g}^{(0)} \rrbracket := (T_c \cdot \llbracket \vec{f}_0^{(0)} \rrbracket, \dots, T_{c \oplus (N-1)} \cdot \llbracket \vec{f}_{N-1}^{(0)} \rrbracket)$
 - 7: $\llbracket v \rrbracket \leftarrow \mathcal{F}_{\text{weakDotProduct}}(\llbracket \vec{g}^{(0)} \rrbracket, \llbracket \vec{f}^{(1)} \rrbracket) \triangleright 1 \text{ round, } k \text{ bits}$
 - 8: Run $\mathcal{F}_{\text{verify}}$ on input $(\llbracket \vec{g}^{(0)} \rrbracket, \llbracket \vec{f}^{(1)} \rrbracket, \llbracket v \rrbracket)$
 - 9: **return** $\llbracket v \rrbracket$
-

F Detailed Benchmark Results

F.1 Implementation Details

For LUT-16 and LUT-256, the offline phase computes on bit shares. We implemented the generation of random one-hot vectors using bit-slicing where we operate on a pack of 16 bits. This improves both local computation and efficiency for I/O. The inner product required in oblivious table lookups is realized using 16 and 256 hard-coded tables, respectively, that contain the lookup table permuted by the reconstructed public $c \oplus j$ (see Step 4 in Protocol 4). Specifically, the table entry contains 4 and 8 bitvectors, respectively, where each encodes the i -th output bit of the permuted table. Then, given the random one-hot vector \vec{e} as bitvector, the inner product for each output bit can be computed as $(\vec{e} \ \& \ t[i]).\text{parity}() \bmod 2$ where $\&$ denotes bit-wise AND. This approach neither requires branching nor multiplication instructions and improves local computation of Protocol 4 by about 10 times compared to a naive approach.

F.2 Benchmark Results

Tables 6 and 7 give detailed numbers, including preprocessing and online phase communication for passive and active security. The time/communication for preprocessing and online

phase includes all necessary checks for malicious security (e.g., Protocol 2 and compare-view).

For semi-honest security, two of our protocols, LUT-16 and $GF(2^4)$ -Circuit outperform the state-of-the-art $GF(2^8)$ -Circuit protocol. LUT-16 offers the fastest online phase which improves online throughput by factor 1.36 compared to $GF(2^8)$ -Circuit, while $GF(2^4)$ -Circuit has the highest, overall throughput resulting in a factor 1.28 improvement compared to $GF(2^8)$ -Circuit. The LUT-256 protocol variants allows for a potentially rapid online phase due to the few communication rounds and low amount of data. Our current implementation cannot fully realize this potential. The bottleneck is the local computation of the inner product between the random one-hot bitvector and the permuted 256-element lookup table. Further optimization is required for this step. This poor performance coupled with the high cost of the preprocessing makes the (2,3) variant not attractive for the malicious security setting, so we didn't implement it. The (3,3) LUT-256 protocol overcomes the expensive preprocessing phase at the cost of doubling the communication in the online phase compared to (2,3) LUT-256 and can thus improve throughput by factor ≈ 2.8 . However, it still falls short to $GF(2^4)$ -Circuit since the more expensive local computation dominates in this setting.

For malicious security, we first note that our implementation of the multiplication correctness check of Protocol 2 is comparatively much slower than, e.g., the triple post-sacrifice step in [31], despite optimizations using carry-less multiplication for $GF(2^{64})$ multiplication and inner products (about 3 to 8 times). This results in a significantly slower online phase for our protocols. However, regarding overall throughput, all protocols using Protocol 2 outperform the check of [31] with improvements ranging from factor 2 to 4, respectively. Naturally, these protocol variants also use much less communication, decreasing the number of sent bytes by up to factor 12.

While our protocol implementation was geared towards high-throughput, it is possible to get an approximation of the computation latency by evaluating only one AES block (see Table 8). In network settings with latency, lookup-table based approaches with a lower number of rounds therefore have a lower latency. LUT-16 and LUT-256 reduce latency by factor 1.2 to 2.4 compared to $GF(2^8)$ -Circuit, respectively. Our $GF(2^4)$ -Circuit variant requires 4 rounds per S-box and thus increases computation latency.

Chida et al. [21] also report on another setup where machines are connected in a ring topology with dual connections between each machine. This allows for optimizations where for each step, half of the data is sent to one party, and half is sent to the other party, essentially rotating the parties' roles to improve throughput. Moreover, they also implement counter-mode caching as a mode-level optimization. Both optimizations can be implemented in our protocols and are expected to enhance the performance.

Table 6: Benchmark results for passive security on batches of 250 000 AES blocks in the LAN setting with ≈ 9.42 Gbits/sec bandwidth. Time and communicated data is reported per batch, the throughput is reported as AES blocks per second. We denote the best value for online phase and throughput in bold.

Protocol	Preprocessing		Online		Throughput (blocks/s)		
	Time (s)	Data (MB)	Time (s)	Data (MB)	Preprocessing	Online	Total
$GF(2^8)$ -Circuit [21]	-	-	0.44	160	-	568 504	568 504
LUT-16	0.46	55	0.32	80	540 374	775 697	318 498
$GF(2^4)$ -Circuit	-	-	0.34	100		729 822	729 822
(2,3) LUT-256	5.99	1235	0.66	40	41 724	381 521	37 611
(3,3) LUT-256	1.92	110	0.39	80	130 036	641 302	108 114

Table 7: Benchmark results for active security on batches of 100 000 AES blocks in the LAN setting with ≈ 9.42 Gbits/sec bandwidth. Time and communicated data is reported per batch, the throughput is reported as AES blocks per second. We denote the best value per metric in bold.

Protocol	Preprocessing		Online		Throughput (blocks/s)		
	Time (s)	Data (MB)	Time (s)	Data (MB)	Preprocessing	Online	Total
$GF(2^8)$ -Circuit ([21] + [31])	9.56	≈ 470	0.72	≈ 192	10 459	138 832	9 727
$GF(2^8)$ -Circuit ([21] + Protocol 2)	-	-	2.17	≈ 64	-	46 081	46 081
LUT-16 + Protocol 2	0.23	22	2.24	≈ 32	442 095	44 624	40 533
$GF(2^4)$ -Circuit + Protocol 2	-	-	2.34	≈ 40	-	42 799	42 799
(3,3) LUT-256 + Protocol 7	0.84	44	3.65	≈ 32	119 745	27 373	22 280

Table 8: Computation latency for one AES block, reported as the execution time of the online phase in various network settings.

Protocol	Malicious	Latency (in ms)			
		1 Gbit/s ≤ 1 ms RTT	200 MBit/s 15ms RTT	100 MBit/s 30ms RTT	50 MBit/s 100ms RTT
$GF(2^8)$ -Circuit [21]	\times	3	454	818	2489
LUT-16	\times	19	489	801	1955
$GF(2^4)$ -Circuit	\times	4	592	1246	3577
(2,3) LUT-256	\times	1	174	308	1002
(3,3) LUT-256	\times	1	167	357	1024
$GF(2^8)$ -Circuit ([21] + [31])	\checkmark	17	449	872	2761
$GF(2^8)$ -Circuit ([21] + Protocol 2)	\checkmark	11	718	1449	4317
LUT-16 + Protocol 2	\checkmark	103	729	1303	3414
$GF(2^4)$ -Circuit + Protocol 2	\checkmark	12	862	1712	5161
(3,3) LUT-256 + Protocol 7	\checkmark	39	415	783	2475