IACR Communications in Cryptology
ISSN 3006-5496, Vol. 3, No. 3, 41 pages.

https://doi.org/10.62056/a3z11w9p

Check for updates

TLShare: Private Authenticated MPC and FHE Inputs Over TLS

Manuel B. Santos[1], Dimitris Mouris[a,2], Xiang Xie[3], Miguel de Vega[1] and Andrei Lapets[1]

[1] Nillion Labs, Switzerland
[2] Snap Inc., USA
[3] Primus Labs, China

Abstract. Transport Layer Security (TLS) allows clients to establish secure and private channels with servers. DECO (CCS'20) and follow-up works proposed protocols that enable proving the provenance of a TLS response, i.e., that a payload came from a particular server, without needing server-side modifications. Unfortunately, these works are limited to proving Boolean statements over the payload (e.g., age ≥ 18) and cannot combine payloads from multiple clients.

We introduce TLShare, a framework that extracts authenticated data from a TLS connection and imports it into secure multiparty computation (MPC) or fully homomorphic encryption (FHE), without requiring server-side changes or exposing client credentials. Unlike prior work, TLShare allows the payload itself, not just a predicate about it, to serve as private input to secure downstream computation. TLShare supports combining verifiable inputs across multiple clients and servers, enabling new applications such as privacy-preserving financial risk assessment and collaborative analytics. We design three protocols for TLShare: one for MPC using verifiable secret sharing, and two for FHE using interactive and non-interactive zero-knowledge proofs, each ensuring input authenticity, integrity, and end-to-end privacy. We evaluate all three protocols of TLShare over both LAN and WAN settings, comparing their trade-offs and demonstrating their practicality.

Keywords: Fully Homomorphic Encryption · Multiparty Computation · TLS · Zero-Knowledge Proofs

1 Introduction

Transport Layer Security (TLS) is the standard for securing communication over the Internet, providing confidentiality, integrity, and authentication guarantees between clients and servers [DR08, Res18]. However, reusing data transmitted over a TLS connection in a separate application context (e.g., proving a property about the payload to a third party) poses significant challenges [ZCC+16]. The TLS payload is cryptographically bound to the session between the client and the server, and its authenticity is not easily verifiable outside that context. One obvious solution is to modify the TLS server (e.g., a bank) to support cryptographic proofs or attestation mechanisms; however, this is highly impractical, as organizations are unlikely to implement and maintain custom extensions for arbitrary downstream applications [BDM+21]. Conversely, simply handing over client credentials

E-mail: manuel.batalha.santos@gmail.com (Manuel B. Santos), jimouris@udel.edu (Dimitris Mouris), xiexiangiscas@gmail.com (Xiang Xie), miguel@nillion.com (Miguel de Vega), andrei.lapets@nillion.com (Andrei Lapets)
[a]Work completed while at Nillion.

This work is licensed under a "CC BY 4.0" license.
Received: 2026-03-26 Accepted: 2026-09-25

to a third-party application breaks the fundamental trust model and compromises both security and privacy [MMS16]. These limitations highlight the need for mechanisms that allow clients to extract verifiable private data from TLS sessions without requiring server cooperation or revealing sensitive credentials.

DECO [ZMM+20] introduced the concept of a decentralized oracle, which partially addresses the problem of reusing TLS payload. It enables a client to prove statements about a TLS payload in zero knowledge without requiring any server modifications. DECO extracts data from the TLS session and converts it into a zero-knowledge proof (ZKP) friendly form, allowing clients to prove predicates on the payload (e.g., threshold conditions) [ELW+25]. However, DECO only outputs a Boolean value (i.e., whether the predicate holds), which limits expressiveness. Additionally, the proofs are often published on-chain [B+21], revealing information about the data they certify (e.g., age ≥ 18). Finally, decentralized oracles [ZMM+20, ELW+25, RWG+18, CCY23, XYWY24, CDH+25, LEFS25] do not allow combining data from multiple clients as the prover must possess all underlying secrets in the clear.

This motivates extending beyond existing oracle paradigms and exploring a new class of architectures that preserve both privacy and authenticity in a more expressive and composable way. We categorize existing oracles into *public-to-public* and *private-to-public* models. The former exposes both input and output data, limiting their applicability to non-sensitive domains [B+21, LK23]. Private-to-public oracles [ZMM+20, RWG+18, CCY23, LEFS25], allow clients to keep their input data private while releasing a public ZKP of a predicate over that data. However, these oracles suffer from two severe limitations: a) they cannot combine data across users without violating privacy, and b) they rely on simple ZKPs that reveal a single bit of information instead of offering support for any private downstream computation.

Private-to-private oracle. To address these limitations, we introduce a new oracle model, called *private-to-private*, which extracts data from a TLS connection and imports it into private computation servers for secure multiparty computation (MPC) or fully homomorphic encryption (FHE) while keeping the data private and maintaining data provenance. This opens the door to new use cases such as credit scoring for collateralized lending [JLP03] or aggregate statistical analysis across institutions [Cou17, LJA+18, Man21, MST24, MPD+25] without exposing any sensitive information. We introduce a novel framework called TLShare,[1] which is the first private-to-private oracle for MPC and FHE downstream computation. Contrary to approaches that require institutions (e.g., banks or hospitals) to participate in MPC or FHE, where clients have little to no say over their data [Goo20, MMT+24, CGB17, EFM+19], TLShare allows clients to control what data is used in which downstream computation. Additionally, secure computation between institutions requires significant trust and coordination, often hindered by regulatory or competitive concerns (e.g., competitive dynamics, GDPR constraints, operational overhead). In contrast, with TLShare, clients connect to TLS endpoints that serve as authentic and verifiable data sources without having to participate in any secure computation. Finally, in contrast to private-to-public oracles, where a ZKP proves a statement about a single client's data, TLShare allows data from multiple clients to be combined under secure computation while proving provenance.

1.1 Our Contribution

We propose TLShare, the first private-to-private oracle that allows private computation over authenticated and private TLS payloads. The client extracts authenticated data (e.g., financial information) from trusted sources (e.g., banks, credit bureaus) and inputs it into MPC or FHE servers for downstream computation. For instance, after importing

[1] Pronounced "TLS Share".

financial data from multiple institutions, the downstream MPC/FHE can perform a financial risk assessment for the user without revealing the raw data to lenders. TLShare minimizes adoption barriers for private computation, as institutions simply handle a standard TLS connection (e.g., a client accessing their bank account) and remain agnostic to the downstream MPC/FHE. To that end, in Section 6 we elaborate on two applications for TLShare.

The TLS connection is the foundation for clients to input their response payload with provenance and MPC/FHE servers to verify its authenticity. The main technical challenge is to authenticate the TLS response without revealing it to the attestor, and then carry that authentication into the private inputs consumed by MPC or FHE. We address the first step with Π_{AuthData} and the second with verifiable sharing or encryption proofs, as detailed in Sections 4, 5.1, and 5.2.

TLShare is agnostic to the TLS version as it only requires an underlying protocol Π_{AuthData} that realizes $\mathcal{F}_{\mathsf{AuthData}}$. Our implementation instantiates this protocol using GtP for TLS 1.2 [XYWY24] and we discuss support for TLS 1.3 in Section 7.4. Below, we outline the three novel protocols of TLShare.

1.1.1 MPC mode

This protocol (Section 4) focuses on private downstream computation within an MPC network. At a high level, the user and an attestor first perform a two-party computation and connect to a TLS server. This 2PC is secure against malicious adversaries and ends up with commitments on each bit of the TLS server's response. Next, we use these commitments as the base commitments for verifiable secret sharing, associating them with the randomness and the commitments produced during the TLS phase. Finally, each MPC node verifies the correctness of its share against the original TLS commitment using consistency checks.

1.1.2 FHE mode with Non-Interactive ZK Proofs

Our second protocol (Section 5.1) begins similarly to our first one, in that we start with commitments for each bit of the TLS response. Next, the client encrypts the TLS response using a public key version of an FHE scheme based on the learning with errors problem. We focus on the CGGI scheme [CGGI20], but our techniques can be generalized to other schemes. To prove that the ciphertext is a valid encryption of the same message committed in the TLS phase, the client and the MPC nodes run a non-interactive ZK (NIZK) proof. The client starts by constructing a binary vector encoding the message bits, encryption randomness and noise, and commits to this vector using a pairing-based vector commitment scheme (see Section 2.4). Finally, we introduce a NIZK argument of vector membership to show that the value committed in each commitment corresponds to the message bits contained within the vector commitment. This ensures that the ciphertext is well-formed and tied to the TLS response.

1.1.3 FHE mode with Interactive ZK Proofs

Finally, our third protocol (Section 5.2) relies on information-theoretic message authentication codes (IT-MACs) using an interactive zero-knowledge (IZK) [YSWW21] protocol to prove the validity of Regev [Reg05] FHE encryptions. More specifically, the client proves in zero knowledge to the attestor that the ciphertext is computed correctly using the response plaintext, a binary randomness vector, and bounded noise terms. The attestor signs the ciphertext and an FHE server verifies that it corresponds to the original TLS response, without seeing the plaintext.

1.2 Related Work

Proving various properties of data obtained from TLS connections has been an increasingly popular research area. We classify these works as private-to-public oracles since private data from a TLS connection (e.g., age) results in a public proof (e.g., age ≥ 18), usually used in smart contracts. Another crucial distinction between private-to-public oracles and our work is that the ZKP in the former is quite limited (True/False statements) compared to the capabilities of TLShare (see Section 6). This line of work is also referred to as designated-commitment TLS protocols and focuses on ensuring data authenticity and integrity while providing privacy guarantees.

Town Crier (TC) [ZCC+16] was an early pioneer, relying on a trusted execution environment (TEE), specifically Intel Software Guard Extensions (SGX), to act as a trusted bridge between smart contracts and external websites. However, its reliance on SGX comes with a multitude of attacks [YF14, VMW+18, CCX+19]. Unlike TC, our approach does not depend on hardware-based trust assumptions. The notable work of DECO [ZMM+20] enables clients to demonstrate data provenance and prove statements about it selectively or in zero-knowledge, all without requiring any protocol modifications on the TLS server side. DECO's design is tailored for TLS 1.2, which remains widely supported across many real-world systems.[2] Despite its utility, DECO is susceptible to attacks during the multiplicative-to-additive protocol [TS21, MP24, CDH+25]. Dido [CCY23] builds upon DECO by introducing a round-optimal three-party key exchange along with enhanced X25519-based protocols. TLSNotary [PSE22] further extends this model by separating the verifier into two roles: a verifier and a notary who co-signs the client's transcript, enabling auditability and trust separation. Recently, Garble-then-Prove (GtP) [XYWY24] demonstrated substantial performance improvements over DECO by optimizing several two-party protocols and transforming authenticated garbling bits into zkSNARK-friendly commitments. While all these protocols focus largely on TLS 1.2, Janus [LEFS25] introduces performance enhancements and extends support to both 1.2 and 1.3. However, Janus requires clients to pre-select cipher suites, reducing negotiation flexibility and increasing implementation complexity. Origo [ELW+25] focuses on TLS 1.3 and follows the GtP approach by converting the commitment into a ZK-friendly one. It introduces a two-party computation (2PC)-free oracle solution, effectively reducing 2PC overheads, but at the cost of requiring additional security assumptions, such as resilience to a proxy that allows the client to perform man-in-the-middle attacks, a notably strong assumption.

Several works explore private features over TLS connections, often requiring modifications on the server side: TLS-N [RWG+18] introduces a mechanism for non-repudiation of origin, Oblivious TLS [ADST21] enables a client running a standard TLS endpoint to communicate securely with a set of MPC servers. An orthogonal approach, Compute but Verify [DGPS24], presents a method to execute an MPC protocol on authenticated inputs. As these works rely on server-side extensions, the cooperation of servers in adopting their protocols is essential, and something we see as a significant lift. In contrast, TLShare offers a novel approach to authenticate MPC inputs without requiring any changes to the server infrastructure.

PECO [San22] and DiStefano [CDH+25] focus on providing an additional property: client's privacy. PECO offers two modes: one that provides ZKP of valid signatures for authentication among a set of k verifier-approved servers (k-anonymity), and another that verifies the validity of the server's signature in zero-knowledge. However, PECO uses a broken primitive [CLR24] and lacks an implementation. DiStefano, on the contrary, generates secure private commitments over TLS 1.3 data (relying on modified AES) and introduces a new primitive called "ZKP of knowledge of valid signatures (ZKPVS)" that proves in zero-knowledge that a signature is valid in relation to a set of public keys. This

[2] DECO has explored TLS 1.3 but has introduced vulnerabilities [CDH+25], largely attributable to the technical complexities introduced by TLS 1.3's revised handshake protocol and key derivation mechanisms.

property allows a client using DiStefano to prove that they are connecting with the intended server without revealing the server's identity (which will reveal the client's browsing history over time). ZKPVS is orthogonal to TLShare and compatible with our approach.

All the aforementioned works implement private-to-public oracles, focusing on proving a TLS property in zero knowledge. In contrast, TLShare introduces the first *private-to-private oracle*, supporting three distinct modes for secure downstream computation. We stress that the above works involve three parties (server, verifier, and client) with the security assumption that the verifier and the client do not collude. In TLShare, we use the same architecture and threat model but refer to the *verifier* as *attestor* to avoid confusion between the TLS verifier and the MPC/FHE verifiers.

MPCAuth [TCDP23] is a related work that enables a client to authenticate to multiple TLS servers independently with the work of only one authentication. Although MPCAuth uses a similar technique of encrypting the TLS channel within MPC, it solves a different problem than private-to-public oracles, that of authentication. Likewise, zkcreds [RWGM23] focus on anonymous credentials such as proving properties from a passport to anonymously access age-restricted videos. Lastly, zkLogin [BCJ+24] leverages identity tokens to authenticate blockchain transactions and ensure that the off-chain and on-chain identities of users remain hidden.

Finally, there exists a special class of ZKP systems over secret shared data where the servers can verify some property of their shares without learning the underlying data [CGB17, BBC+19, BBC+21, DPRS23, MPD+25]. These ZKPs have been increasingly used in the realm of privacy-preserving statistics, such as histograms and heavy-hitters, but they do not support any form of input authenticity from a third party.

2 Preliminaries

2.1 Multiparty Computation (MPC)

MPC enables multiple parties to compute a function $f(x_1, \ldots, x_n)$ over their inputs $x_1, \ldots, x_n$ without revealing their private data [Yao82, Yao86]. MPC protocols provide different security guarantees: semi-honest security assumes participants follow the protocol honestly but may try to infer additional information, while malicious security accounts for participants who deviate from the protocol. We view these protocols in two categories: garbled circuits [Yao86, GMW87, KS08, WRK17b, YWZ20] and secret-sharing protocols [DPSZ12, KOS16, Kel20]. Garbled circuits encode the function f as a Boolean circuit (e.g., AES, SHA-256) comprising `AND`, `XOR`, etc. gates and evaluate an encrypted variant of it. Secret-sharing techniques focus on arithmetic circuits (consisting of additions and multiplications) and split the inputs into seemingly random values, which are then distributed between the parties.

One popular scheme is Shamir's secret sharing [Sha79] that constructs a random polynomial $F(x) = \mu + a_1 x + a_2 x^2 + \cdots + a_{\mathsf{T}-1} x^{\mathsf{T}-1}$ over a finite field $\mathbb{F}_q$, where μ is the secret message and T is the reconstruction threshold. The share for party i is a pair $(i, F(i))$, and any T of shares can recover μ via Lagrange interpolation. We use $[\![x]\!]$ to refer to shares of x and $[\![x]\!]_i$ to the share of x owned by party i. In TLShare, we use Pedersen verifiable secret sharing (Pedersen VSS) [Ped92], where each share is also accompanied by a proof verifying its validity. In Pedersen VSS, the dealer commits to the secret using Pedersen commitments [Ped92] (i.e., $g^{\mu} \cdot h^{r}$ for generators g and h and random r) and distributes Shamir shares $([\![\mu]\!]_i, [\![r]\!]_i)$ along with Pedersen commitments to the coefficients of the corresponding Shamir polynomials. We provide more details about Pedersen VSS and its integration with TLShare in Section 4.

2.2 Fully Homomorphic Encryption (FHE)

Fully homomorphic encryption (FHE) enables computation directly on encrypted data, supporting arbitrary encrypted additions and multiplications without the need to decrypt intermediate results [Gen09, VJH21, GMT23]. Modern FHE constructions rely on the hardness of the Learning With Errors (LWE) problem [Reg05] or its ring-based variant, RLWE [LPR10], which encrypt messages using noisy inner products. The LWE problem, introduced by Regev [Reg05], is based on the difficulty of recovering a secret $\boldsymbol{s}$ given samples of the form $\boldsymbol{b} = \boldsymbol{A} \cdot \boldsymbol{s} + \boldsymbol{e}$, where $\boldsymbol{A} \in \mathbb{Z}_q^{m \times n}$ is a public matrix, $\boldsymbol{s} \in \mathbb{Z}_q^n$ is a secret vector, and $\boldsymbol{e} \in \mathbb{Z}^m$ is a small noise vector. From the LWE assumption, one can construct a public-key encryption scheme consisting of three algorithms $(\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ such that $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}$ and $\mathsf{Dec}(\mathsf{sk}, \mathsf{Enc}(\mathsf{pk}, \mu)) = \mu$. Furthermore, an LWE-based FHE scheme supports both addition and multiplication over encrypted data. Specifically, $\mathsf{Enc}(\mathsf{pk}, \mu_1) \oplus \mathsf{Enc}(\mathsf{pk}, \mu_2) = \mathsf{Enc}(\mathsf{pk}, \mu_1 + \mu_2), \quad \mathsf{Enc}(\mathsf{pk}, \mu_1) \otimes \mathsf{Enc}(\mathsf{pk}, \mu_2) = \mathsf{Enc}(\mathsf{pk}, \mu_1 \cdot \mu_2)$, where the operations $\oplus$ and $\otimes$ on ciphertexts correspond to plaintext addition and multiplication, respectively.

Regev Scheme. Choose the dimension of the secret-key n, the ciphertext modulus q, and a noise distribution χ (e.g., discrete Gaussian) based on a security parameter λ. The Regev encryption scheme [Reg05] is defined as follows. The key generation algorithm KeyGen samples a secret key uniformly $\boldsymbol{s} \leftarrow \mathbb{Z}_q^n$, a matrix $\boldsymbol{A} \leftarrow \mathbb{Z}_q^{n \times m}$, and noise $\boldsymbol{e} \leftarrow \chi^m$, then sets $\boldsymbol{y}^\top = \boldsymbol{s}^\top \boldsymbol{A} + \boldsymbol{e}^\top$. The public key is $(\boldsymbol{A}, \boldsymbol{y})$ and the secret key is $\boldsymbol{s}$. Here, $m = \Omega(n \log q)$. To encrypt a bit μ, sample $\boldsymbol{r} \leftarrow \{0,1\}^m$ and compute $(\boldsymbol{a}, b) = (\boldsymbol{A}\boldsymbol{r},\ \boldsymbol{y}^\top \boldsymbol{r} + \lfloor \mu \cdot q/2 \rfloor) \mod q$. To decrypt, compute $b - \boldsymbol{s}^\top a \mod q$ and output 0 if the result is closer to 0 than to $q/2$, and 1 otherwise.

Subsequent work improved on Regev's initial scheme by decreasing the size of the public key [Joy24, LPR10, LP11]. The LPR [LPR10] scheme decreases the matrix $\boldsymbol{A}$ to an $n \times n$ matrix. This is achieved by adding noise to the ciphertext components, i.e., $(\boldsymbol{a}, b) = (\boldsymbol{A}\boldsymbol{r} + \boldsymbol{e},\ \boldsymbol{y}^\top \boldsymbol{r} + e_0 + \lfloor \mu \cdot q/2 \rfloor) \mod q$, where $\boldsymbol{r}, \boldsymbol{e} \leftarrow \chi^n$ and $e_0 \leftarrow \chi$.

CGGI Scheme. Encrypted operations result in noise growth, which eventually prevents decryption. To allow unlimited encrypted operations, a technique called bootstrapping resets noise down to minimal levels [DM15, CGGI20]. The DM [DM15] and CGGI [CGGI20] (TFHE) schemes support arbitrary univariate functions during bootstrapping, enabling them to replace sequences of additions and multiplications [LMP22, GUM+24]. CGGI has the most efficient bootstrapping in terms of latency and is thus the cryptosystem we use as one of our back-ends.

Let us now recall the CGGI scheme in more detail. For a modulus $q \in \mathbb{N}$, define the discretized torus as $\mathbb{T}_q = q^{-1}\mathbb{Z}/\mathbb{Z}$ [Joy22]. The plaintext space is $\mathbb{Z}_p$, and the secret key space is $\{0,1\}^n \subset \mathbb{Z}^n$, where $p \mid q$ and $\Delta = q/p$. The CGGI secret key is defined as $(s_1, \ldots, s_n) \in \{0,1\}^n$. To encrypt a message $\mu \in \mathbb{Z}_p$ under the secret key, one samples a random vector $(a_1, \ldots, a_n) \in \mathbb{T}_q^n$ uniformly and samples an error term $e \in \mathbb{Z}$ from a suitable distribution. The ciphertext is defined as $(a_1, \ldots, a_n, b) \in \mathbb{T}_q^{n+1}$, where: $b = \sum_{j=1}^n s_j \cdot a_j + \mu^*$ and $\mu^* = (\Delta \cdot \mu + e)/q \in \mathbb{T}_q$. This encryption is additively homomorphic and supports bootstrapping to maintain decryptability even after many homomorphic operations.

The public-key version presented in [Joy21] follows the method proposed by Rothblum [Rot11], where the public key encodes encryptions of the zero message under different randomness. In summary, it constructs a public key matrix $\boldsymbol{A} \in \mathbb{T}_q^{m \times (n+1)}$ by sampling m independent encryptions of 0 under $\boldsymbol{s}$. The public key is $\boldsymbol{A}$, and the secret key is $\boldsymbol{s}$. To encrypt a message $\mu \in \mathbb{T}_p$, sample a random vector $\boldsymbol{r} \leftarrow \{0,1\}^m$ and compute $(\boldsymbol{a}, b) = \boldsymbol{r}^\top \boldsymbol{A} + (0, \ldots, 0, \mu) \in \mathbb{T}_q^{n+1}$. Joye [Joy24] uses the LPR [LPR10] approach to reduce the size of the public key by adding small noise to the ciphertext.

2.3 Zero-Knowledge Proofs (ZKP)

A ZKP allows a prover $\mathcal{P}$ to convince a verifier $\mathcal{V}$ that some given statement is true, without revealing any information beyond the fact that the statement is true [Mic00, BCCT13, BCG+13, GGPR13, BBHR19, MT21, BCI+22]. Informally, ZKPs must be: *complete*, meaning that valid proofs are always accepted; *sound*, meaning that it is infeasible to produce a convincing proof for a false statement; and *zero-knowledge*, meaning that the proof reveals nothing about the underlying witness [Mic00, Lib24]. In the context of this work, we focus on two specific ZKPs: a non-interactive zero-knowledge (NIZK) for proving correct encryptions, and an interactive zero-knowledge (IZK) system called QuickSilver [YSWW21].

NIZK. We consider NIZK arguments as a tool for generating verifiable CGGI ciphertexts (Section 5.1.1). Let $\{R_\lambda\}_\lambda$ be a family of NP relations parameterized by the security parameter λ [Lib24]. A NIZK argument for this family consists of three algorithms:

- $\mathsf{NIZK.GenCRS}(1^\lambda, \mathsf{lpp}) \to (pp, \tau)$: generates a common reference string pp and a simulation trapdoor τ, possibly using language-dependent parameters lpp. The relation to be proven becomes $\mathcal{R}_{\mathsf{pp}}$.
- $\mathsf{NIZK.Prove}_{\mathsf{pp}}(x, w) \to \pi$: given a statement x and witness w such that $(x, w) \in \mathcal{R}_{\mathsf{pp}}$, returns a proof π.
- $\mathsf{NIZK.Verify}_{\mathsf{pp}}(x, \pi) \to \{0, 1\}$: checks whether π is a valid proof for x under pp.

In addition to the aforementioned ZKP algorithms, we include a commitment algorithm $\mathsf{NIZK.Com}_{\mathsf{pp}}(w) \to V$ allowing $\mathcal{P}$ to commit to the witness such that it can be linked to the proof.

IZK with QuickSilver. QuickSilver [YSWW21] is an IZK protocol [WYKW21, BMRS21] and leverages Information-Theoretic Message Authentication Codes (IT-MACs) to authenticate Boolean and arithmetic circuits. IT-MACs are also widely used in MPC [BDOZ11, WRK17a, DILO22]. QuickSilver involves a prover and a verifier where $\mathcal{V}$ holds a private key $\Delta \in \mathbb{F}_{2^\lambda}$ and a local key $K[x] \in \mathbb{F}_{2^\lambda}$, while $\mathcal{P}$, holds a bit $x \in \{0, 1\}$ and a MAC tag of the form $M[x] = K[x] + x \cdot \Delta \in \mathbb{F}_{2^\lambda}$. In QuickSilver, the prover commits to a witness and to all wire values of the circuit using IT-MACs and then evaluates the circuit interactively with $\mathcal{V}$ to obtain IT-MACs on output values. We use $\langle\langle x \rangle\rangle$ to refer to such verifiable shares with IT-MACs, where $\mathcal{P}$ holds x and $M[x]$, and $\mathcal{V}$ holds $K[x]$ and Δ. QuickSilver proves circuit satisfiability and verifies correctness with batched checks [NNOB12, WYKW21] and achieves constant-size communication.

2.4 Vector Commitments over Pairings

Vector commitments allow a sender to commit to an ordered sequence $\boldsymbol{w} = (w_1, \ldots, w_n) \in \mathbb{Z}_p^n$ and later open the value at a specific position i along with a proof that w_i is indeed the committed message at that index [CF13]. One of the key advantages of vector commitments is that the size of the commitment remains constant, regardless of the size of the vector. In this work, we focus on the vector commitment scheme from [Lib24], which is based on bilinear pairings and includes the additional functionality of proving inner product relations. Specifically, it enables the prover to demonstrate that a committed private vector $\boldsymbol{w}$ satisfies a linear relation $\langle \boldsymbol{w}, \boldsymbol{t} \rangle = x$, where $\boldsymbol{t}$ and x are public.

A bilinear map, or asymmetric pairing, is a function $e : \mathbb{G} \times \hat{\mathbb{G}} \to \mathbb{G}_T$ defined over two source groups $\mathbb{G}$ and $\hat{\mathbb{G}}$. We assume a structured common reference string (CRS) that defines the group elements $(g, \{g_i = g^{\alpha^i}\}_{i \in [1,2n] \setminus \{n+1\}})$ and $(\hat{g}, \{\hat{g}_i = \hat{g}^{\alpha^i}\}_{i=1}^n)$, for a hidden scalar $\alpha \in \mathbb{Z}_p$ [LY10, Lib24]. This structure relies on the assumed hardness of the $(2n, n)$-Discrete Logarithm problem, which assumes computing α from the CRS is

hard [Lib24, Definition 1]. We analyze algebraic PPT adversaries in the Algebraic Group Model (AGM) [FKL18]. For either source group $\mathbb{H} \in \{\mathbb{G}, \hat{\mathbb{G}}\}$, whenever an adversary outputs $Y \in \mathbb{H}$, it also generates $\mathsf{repr}(Y) = (a_1, \ldots, a_m) \in \mathbb{Z}_p^m$ such that $Y = \prod_{i=1}^m Y_i^{a_i}$, where $Y_1, \ldots, Y_m \in \mathbb{H}$ are group elements from the CRS.

To commit to a vector $\boldsymbol{w} = (w_1, \ldots, w_n) \in \mathbb{Z}_p^n$, the sender samples $\gamma \leftarrow \mathbb{Z}_p$ and computes the vector commitment:

$$V = g^{\gamma} \cdot \prod_{j=1}^n g_j^{w_j} = g^{\gamma + \sum_{j=1}^n w_j \cdot \alpha^j}. \tag{1}$$

To open a specific position i in the committed vector, the sender provides $\pi_i = g_{n+1-i}^{\gamma} \cdot \prod_{j=1, j\neq i}^n g_{n+1-i+j}^{w_j} = (V/g^{w_i \cdot \alpha^i})^{\alpha^{n+1-i}}$, which can be verified by checking $e(V, \hat{g}_{n+1-i}) = e(\pi_i, \hat{g}) \cdot e(g_1, \hat{g}_n)^{w_i}$, confirming that w_i is the message committed at index i ($i \in [1, n]$).

The use of pairings also enables proving that the committed vector satisfies a linear relation of the form $\langle \boldsymbol{w}, \boldsymbol{t} \rangle = x$, for public $\boldsymbol{t} \in \mathbb{Z}_p^n$ and $x \in \mathbb{Z}_p$. This is done by raising the verification equation to $t_i \in \mathbb{Z}_p$ and multiplying over all $i \in [1, n]$ as $e\big(V, \prod_{i=1}^n \hat{g}_{n+1-i}^{t_i}\big)$, which equals $e\big(\prod_{i=1}^n \pi_i^{t_i}, \hat{g}\big) \cdot e(g_1, \hat{g}_n)^{\sum_{i=1}^n w_i \cdot t_i}$.

3 Overview of TLShare

3.1 Functionality & Threat Model

We present an overview of TLShare, a framework that enables data extraction from TLS into secure computation. TLShare is a private-to-private oracle, allowing data transmitted over TLS to be consumed by MPC or FHE without any changes to the TLS server.

We present the ideal functionality of TLShare $\mathcal{F}_{\text{TLShare}}$ in Fig. 1, which defines the interactions between a client, an attestor, a TLS server, and a set of private computation nodes $\mathcal{N}_1, \ldots, \mathcal{N}_n$. We adopt the standard assumption used in prior work [ZMM+20, CDH+25, XYWY24], namely that the attestor does not collude with the client. The functionality $\mathcal{F}_{\text{TLShare}}$ has two phases: 1) the *TLS Phase* and 2) the *Share/Encryption Phase*. The former involves the TLS handshake, the client's query to the TLS server, and the server's response, while the latter transforms the server's response either to secret shares or to FHE ciphertexts along with a proof of provenance. The following properties follow from the functionality in Fig. 1:

- **Server:** assumed to be honest, so R is the ground truth.
- **Malicious client:** cannot change Q or R, that were attested.
- **Malicious attestor:** cannot make the client receive a different Q/R from that which was received/sent by the server.
- **Privacy:** A malicious attestor cannot learn any information about Q and R, except their length. This is aligned with [ZMM+20, XYWY24] as TLS is not length-hiding.

We introduce three novel protocols for TLShare and support two modes of operation (i.e., MPC and FHE), with one protocol for MPC (Section 4) and two different protocols for FHE (Section 5). We target two scenarios that combine data from multiple TLS servers, which is not feasible with any related work [ZMM+20, ELW+25, RWG+18, CCY23, XYWY24, CDH+25, LEFS25]. We outline both scenarios below and applications in Section 6.

Single-client multiple-sources. In this case, a client aggregates data from various institutions (e.g., banks, tax authorities) through TLS into secure computation (such as assessing their financial risk). We demonstrate an overview of this scenario in Fig. 2.

Multiple-clients multiple-sources. In this case, multiple users contribute data from different sources into a joint secure computation. Our motivating application for this

Ideal Functionality $\mathcal{F}_{\text{TLShare}}$

This functionality interacts with a server $\mathcal{S}$, a client, an attestor, $n > 1$ Nodes $\mathcal{N}_i$, and an adversary.

- **Public Parameters:** Parameter $\mathsf{mode} \in \{\mathsf{fhe}, \mathsf{mpc}\}$. Query template Query for $\mathcal{S}$ (e.g., a `GET` request) and a threshold T for secret sharing.
- **Server $\mathcal{S}$:** has no input ($\perp$) but hosts and operates a public website.
- **Client:** has a private input (e.g., a password) Inp.
- **Attestor:** has no input ($\perp$).
- **Node $\mathcal{N}_i$ (for $i \in [1, n]$):** In case of mpc mode, the nodes have no input ($\perp$). In case of fhe mode, the nodes have shares of the secret key of an FHE key pair $(\llbracket \mathsf{sk}_{\mathsf{fhe}} \rrbracket_i, \mathsf{pk}_{\mathsf{fhe}})$.

TLS Phase (Handshake, Query, and Response):

1. On input $(\mathcal{S}, \mathsf{Inp}, \mathsf{Query})$ from the client:
   a. Compute $\mathsf{Q} \leftarrow \mathsf{Query}(\mathsf{Inp})$.
   b. Send Q to $\mathcal{S}$.
2. Receive a response R from $\mathcal{S}$:
   a. Store the tuple (Q, R) and send R to the client.
   b. Send $(|\mathsf{Q}|, |\mathsf{R}|)$ to the adversary.

Share/Encryption Phase:

3. On input (mode) from the client, retrieve (Q, R) from memory. Define $\mathcal{R}_{\mathsf{Enc}}(\mathsf{pk}_{\mathsf{fhe}}, \mathsf{R}_j; \mathsf{ct}_j, \omega_j) = 1$ if $\mathsf{ct}_j = \mathsf{Enc}_{\mathsf{pk}_{\mathsf{fhe}}}(\mathsf{R}_j; \omega_j)$ and ω_j satisfies the public well-formedness conditions of the selected FHE scheme, or 0 otherwise.
4. If $\mathsf{mode} = \mathsf{fhe}$, then:
   a. If the client is honest, sample ω_j according to the encryption algorithm and set $\mathsf{ct}_j \leftarrow \mathsf{Enc}_{\mathsf{pk}_{\mathsf{fhe}}}(\mathsf{R}_j; \omega_j)$ for every j. If the client is corrupted, receive $\{(\mathsf{ct}_j, \omega_j)\}_j$ from the adversary and retain these exact ciphertexts only if $\mathcal{R}_{\mathsf{Enc}}(\mathsf{pk}_{\mathsf{fhe}}, \mathsf{R}_j; \mathsf{ct}_j, \omega_j) = 1$ for every j; otherwise, send $(\mathsf{sid}, \mathsf{abort})$ to $\mathcal{N}_1$ and stop.
   b. Send $(\mathsf{sid}, \{\mathsf{ct}_j\}_j)$ to $\mathcal{N}_1$ and to the client.
5. If $\mathsf{mode} = \mathsf{mpc}$, then:
   a. If the client is honest, sample degree-$(\mathsf{T} - 1)$ polynomials F_j with $F_j(0) = \mathsf{R}_j$ and set $\llbracket \mathsf{R}_j \rrbracket_i = F_j(i)$. If the client is corrupted, receive candidate polynomials F_j and shares $\{\llbracket \mathsf{R}_j \rrbracket_i\}_{j,i}$ from the adversary. For each node $\mathcal{N}_i$, retain its exact shares only if $\deg(F_j) < \mathsf{T}$, $F_j(0) = \mathsf{R}_j$, and $\llbracket \mathsf{R}_j \rrbracket_i = F_j(i)$ for every j; otherwise, send $(\mathsf{sid}, \mathsf{abort})$ to $\mathcal{N}_i$.
   b. Send $(\mathsf{sid}, \llbracket \mathsf{R}_j \rrbracket_i)$ for every j to each node $\mathcal{N}_i$ that was not sent an abort.

Figure 1: Ideal Functionality $\mathcal{F}_{\text{TLShare}}$ between a server $\mathcal{S}$, a client, an attestor, and nodes $\mathcal{N}_i$ for $i \in [1, n]$.

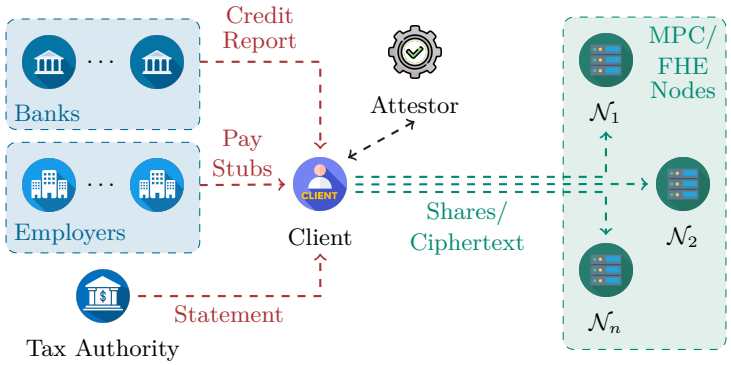

Figure 2: **Overview.** The client wants to: a) import their data from multiple TLS connections into MPC/FHE nodes, b) prove to the nodes that the imported data is authentic and valid, and c) keep the data private. In the depicted example, the MPC/FHE nodes assess the client's financial risk by combining bank statements, credit reports, pay stubs, tax reports, citizenship cards, etc, all from different sources while keeping all the data private. Some applications may further require multiple provers to import their data into the same MPC/FHE session, enabling computation over the joint data.

scenario is the Boston Women's Workforce Council (BWWC) study that measured wage gaps across demographic groups (e.g., race and gender, or inequality of opportunity for women and minority-owned small businesses) [Cou17, LJA+18, Man21]. More details in Section 6.2.

3.2 Modes of Operation

TLShare supports both MPC and FHE through three cryptographic protocols. Namely, TLShare supports MPC mode, FHE mode with non-interactive zero-knowledge proofs (NIZK), and FHE mode with interactive zero-knowledge proofs (IZK). For all these three protocols, we instantiate the TLS phase using GtP's Π_{AuthData} [XYWY24]. The client and attestor jointly process the TLS handshake under garbled two-party computation (2PC) to bind the response to the authenticated server session without revealing Q or R to the attestor beyond their lengths. The client holds the response bits and their IT-MAC tags, while the attestor holds the corresponding verification keys. Next, in the MPC and FHE NIZK modes, these authenticated bits are converted to Pedersen commitments that the attestor signs. In the FHE IZK mode, the IT-MACs serve directly as the encryption-proof witness (Fig. 4 and Sections 4 and 5). The different modes introduce different performance tradeoffs, such as the presence or absence of preprocessing, which we explore in Section 7.4. We use Encode to denote some length-prefixed encoding. Each protocol uses Encode to define its session context ctx in its preamble based on a protocol-specific domain-separation tag and binds it together with the protocol execution by authenticating it with the complete ordered output.

MPC Mode. In MPC mode, the client secret shares their authenticated TLS response among a network of MPC nodes. In the TLS phase, the client obtains an authenticated response R by initiating a TLS session with the server and the attestor. To securely distribute R to the MPC nodes, the client transforms it into an additively homomorphic commitment. We instantiate this using Pedersen commitments, which are compatible with Pedersen verifiable secret sharing (Pedersen VSS). The Pedersen commitments of R serve as the base commitments for Pedersen VSS, allowing the response to be secret-shared. The attestor also signs the Pedersen commitment, enabling each MPC node to verify the authenticity of the committed message. This ensures that the value under commitment is

exactly the authenticated TLS response generated in the protocol's initial phase.

FHE Mode with Non-Interactive ZK Proofs (NIZK). In this mode, the client encrypts the authenticated TLS response under a public FHE key and transmits the resulting ciphertext to a designated FHE node $\mathcal{N}_1$. The corresponding secret key is secret-shared among a set of nodes $\mathcal{N}_i$, enabling threshold decryption. Without loss of generality, the secret key may also be held by an external analyst interested in learning the final result of the computation. As in the MPC mode, the client first initiates a TLS session with the server and an attestor. This process yields signed commitments to the server's private response, which are shared between the client and the attestor. Using these commitments, the client encrypts the TLS response under the FHE public key and generates a non-interactive zero-knowledge (NIZK) argument that attests to the validity of the ciphertext. Specifically, the proof establishes that the ciphertext is well-formed and that it encrypts a value consistent with the previously committed TLS data. To achieve this, the client runs a novel NIZK argument of vector membership, which links the encrypted plaintext to the committed value. This proof guarantees that the ciphertext is a valid encryption of the committed response, thereby ensuring the consistency and correctness of the encryption.

FHE Mode with Interactive ZKPs (IZK). This variant supports FHE based on LWE schemes over a prime modulus, with a correctness proof implemented using an IZK protocol built on IT-MACs [YSWW21]. Compared with the previous two approaches, the IZK mode skips the IT-MAC conversion and uses them directly. These act as authenticated witnesses in the zero-knowledge circuit that proves the correctness and validity of the LWE-based encryption. The proof circuit encodes the encryption process and links it directly to the TLS response obtained earlier. Finally, the resulting ciphertext is signed by the attestor, enabling the FHE node to confirm that it encodes the original TLS response.

Table 1: Comparison of the three protocols during the Share/Encryption phase.

Protocol	No Preprocessing	No CRS	Modulus Type
MPC	✓	✓	Field
FHE with NIZK	✓	✗	Field/Ring
FHE with IZK	✗	✓	Field

Discussion. Table 1 shows a qualitative comparison between the proposed protocols. In MPC mode, the client does not require preprocessing or trusted setup, and there is only a one-shot communication with the MPC nodes during the share phase. The attestor is not involved after the TLS phase, leading to low latency and bandwidth overhead. Meanwhile, both MPC and FHE with NIZK require the client and attestor to run a conversion protocol during the TLS phase to generate commitments of R. In contrast, the FHE with IZK avoids this step by performing an interactive ZKP of valid encryption directly on the TLS-derived IT-MACs. During the share/encryption phase, the MPC and FHE with NIZK modes communicate only once with the nodes $\mathcal{N}_i$. At the same time, the FHE IZK-based protocol requires additional interaction with the attestor to generate the proof that ties the ciphertext to R.

Comparing the FHE mode protocols, the NIZK version protocol requires a single trusted setup and enables a fully non-interactive execution between the client and the FHE node $\mathcal{N}_1$, shifting proof verification entirely onto the FHE node and not involving the attestor after the TLS phase. In contrast, the FHE with IZK protocol eliminates trusted setup but requires preprocessing between the client and attestor during the encryption phase. Overall, the FHE with IZK achieves faster computation time during the encryption phase but is limited to prime-field LWE schemes, whereas the FHE with NIZK method supports ring-based LWE schemes.

3.3 Sequential Multi-Session Composition

The functionality in Fig. 1 describes one stand-alone execution. To capture the applications above, one can define its sequential multi-session extension for a polynomial number m of non-overlapping executions, possibly involving different clients and servers. Let aid identify the downstream computation and let sid_i be a fresh public identifier for execution $i \in [1, m]$. The extension stores state under $(\mathsf{aid}, \mathsf{sid}_i)$, rejects repeated identifiers, and applies the original transitions, corruption behavior, and abort conditions independently to each entry. If a required entry aborts, the application aborts. Otherwise, the downstream computation starts after all m entries complete. In the protocols, the authenticated context becomes $\mathsf{ctx}_i^* := \mathsf{Encode}(\mathsf{aid}, \mathsf{ctx}_i)$, preventing data from different executions or applications from being mixed.

Sessions used by the same downstream computation must have compatible parameters. In MPC mode, they use the same nodes, field, threshold, and sharing convention, where each node's tuple of verified Shamir shares becomes its already-shared input to the downstream MPC. In FHE mode, they use the same public key, plaintext encoding, and ciphertext parameters. The resulting ciphertexts become inputs to the homomorphic evaluation of the fixed function f. Resharing and key switching are outside this interface.

This extension follows the sequential modular-composition theorem for stand-alone protocols [Can00]. We use the stand-alone definition with arbitrary polynomial-length auxiliary input. If the downstream protocol is secure in the same model, under the same corruption restrictions, from the representation above, then m sequential TLShare executions followed by that protocol securely realize their sequentially composed functionality. Indeed, hybrids replace the TLShare executions one at a time, treating preceding transcripts and adversarial state as auxiliary input. Since m is polynomial, the accumulated advantage remains negligible.

4 MPC Mode Protocol

TLShare's MPC mode ensures the authenticated response obtained from a TLS session is not merely shared between the client and attestor but rather securely distributed as verifiable shares to an MPC network. This distinction is important: while the client and attestor share verified shares of the TLS response, private computation over that data must be carried out by the MPC nodes, who receive a different set of shares of the same underlying TLS response.

Garble-then-Prove (GtP) [XYWY24] produces authenticated Pedersen commitments to the TLS response, which are verified by the attestor and returned to the client. To allow the client to convince other parties of the authenticity of a TLS payload, the attestor produces signed commitments. However, this is not enough to enable a distributed MPC network to verify and consume the data securely for private computation. We unlock this functionality in TLShare with a construction that transfers the commitments and the secret shares directly to the MPC nodes in a way that maintains authenticity. We compose two cryptographic tools that share a common foundation: 1) Pedersen commitments, providing authenticated commitments to R; and 2) Pedersen VSS, supporting verifiability of shares under Pedersen commitments. The full protocol is described in Fig. 3.

The protocol is divided into two phases: *TLS phase* and *share/encryption phase*. In the TLS phase, the client and attestor jointly execute the TLS handshake with $\mathcal{S}$ and submit the query, resulting in the server's response committed with an additive homomorphic scheme. These steps follow the Π_{AuthData} protocol from GtP, instantiated with Pedersen commitments. Both the client and attestor hold commitments to each bit of R and, as shown in step 2 of Fig. 3, the attestor signs these commitments, thereby certifying their provenance and binding them to the payload. The only communication in the share phase

TLShare: MPC mode using Pedersen VSS

This protocol runs between a server $\mathcal{S}$, a client, an attestor, and an MPC network of n nodes $\mathcal{N}_1, \ldots, \mathcal{N}_n$. The server $\mathcal{S}$ behaves the same as in a standard TLS connection. $C = g^{\mu} \cdot h^{r}$ is a Pedersen commitment to μ for generators g and h and random r [Ped92].

- **Public Parameters:** a query template (Query) for the server $\mathcal{S}$ (e.g., a `GET` request) and a threshold T for Shamir secret sharing.
- **Session Context:** Before the TLS phase, the parties agree on a fresh sid. For response length ℓ, define $\mathsf{ctx} := \mathsf{Encode}(\texttt{TLShare-mpc}, \mathsf{sid}, \mathsf{cert}_{\mathcal{S}}, \mathsf{Query}, \ell, \mathsf{T}, \mathcal{N}_1, \ldots, \mathcal{N}_n)$.
- **Server $\mathcal{S}$:** has a secret key ($sk_{\mathcal{S}}$) for a signature scheme (e.g., ECDSA), a certificate $\mathsf{cert}_{\mathcal{S}}$, and hosts and operates a public website.
- **Client:** has a private input Inp (e.g., a password).
- **Attestor:** owns a key pair ($pk_{\mathcal{A}}, sk_{\mathcal{A}}$) to a signature scheme (e.g., for ECDSA).
- **MPC Node $\mathcal{N}_i$ (for $i \in [1, n]$):** has no input.

TLS Phase (Handshake, Query, and Response):

1. The client, $\mathcal{S}$, and the attestor run $\Pi_{\mathsf{AuthData}}(sk_{\mathcal{S}}, \mathsf{cert}_{\mathcal{S}}, \mathsf{Inp}, \mathsf{Query})$ instantiated with Pedersen commitments over group $\mathbb{G}$, i.e., converting the IT-MACs of $[\![\mathsf{Q}]\!]$ and $[\![\mathsf{R}]\!]$ to commitments over them. Let C_j be a commitment to the jth bit of R with randomness r_j. The client outputs $(\mathsf{sid}, [C_1, \ldots, C_\ell], \mathsf{Q}, \mathsf{R})$ and the attestor outputs $(\mathsf{sid}, [C_1, \ldots, C_\ell])$, where $\ell = |\mathsf{R}|$.
2. The attestor signs the context and ordered committed response, computing $\sigma = \mathsf{Sign}(sk_{\mathcal{A}}, (\mathsf{ctx}, [C_1, \ldots, C_\ell]))$ and sends it to the client.

Share/Encryption Phase:

3. The client broadcasts $(\sigma, [C_1, \ldots, C_\ell])$ to all nodes $\mathcal{N}_i$.
4. The nodes verify $\mathsf{Verify}(pk_{\mathcal{A}}, (\mathsf{ctx}, [C_1, \ldots, C_\ell]), \sigma)$ and abort in case it fails.
5. For $j \in [1, \ell]$, take each bit of the response R_j, its corresponding commitment C_j with randomness r_j, and run the following:
   a. The client selects two random polynomials $F(x)$ and $G(x)$ over $\mathbb{Z}_q$ of degree at most $\mathsf{T} - 1$ such that:
   $$F(x) = \mathsf{R}_j + \alpha_1 x + \alpha_2 x^2 + \cdots + \alpha_{\mathsf{T}-1} x^{\mathsf{T}-1},$$
   $$G(x) = r_j + \beta_1 x + \beta_2 x^2 + \cdots + \beta_{\mathsf{T}-1} x^{\mathsf{T}-1},$$
   where $\mathsf{R}_j = F(0)$ and $r_j = G(0)$ are the secrets. Recall that R_j is the message and r_j is the randomness used in the commitment C_j. For each node $\mathcal{N}_i$ ($i \in [1, n]$), the client computes the shares $[\![\mathsf{R}_j]\!]_i = F(i)$ and $[\![r_j]\!]_i = G(i)$.
   b. The client computes commitments to the coefficients as $g^{\alpha_l} \cdot h^{\beta_l}$ for $l \in \{1, \ldots, \mathsf{T} - 1\}$, and broadcasts them to all nodes.
   c. The client sends the pair $([\![\mathsf{R}_j]\!]_i, [\![r_j]\!]_i)$ privately to each node $\mathcal{N}_i$.
   d. Each $\mathcal{N}_i$ verifies the correctness of its share as $g^{[\![\mathsf{R}_j]\!]_i} \cdot h^{[\![r_j]\!]_i} = C_j \times \prod_{l=1}^{\mathsf{T}-1} \left(g^{\alpha_l} \cdot h^{\beta_l}\right)^{(i^l)}$. Each node outputs its vector of shares only if the signature and all share-verification checks succeed. Otherwise, abort.

Figure 3: TLShare protocol for $\mathsf{mode} = \mathsf{mpc}$.

is between the client and the MPC nodes, as the attestor is no longer involved. The client starts (Fig. 3 step 3) by distributing the signed commitments to all MPC nodes, allowing them to verify the provenance of the committed data (step 4). Then, the Pedersen commitments generated during the TLS phase are used as the commitments of the secrets in Pedersen VSS. We refer to these as *base commitments*. Specifically, treating each

TLS bit commitment $C_j = g^{\mathsf{R}_j} \cdot h^{r_j}$ as the base commitment, the client creates Shamir secret-sharing polynomials $F(\cdot)$ and $G(\cdot)$ for both R_j and its randomness r_j, in step 5.a) of Fig. 3. Then, at steps 5.b) and 5.c), the client generates and sends corresponding commitments of the polynomial coefficients $g^{\alpha_l} \cdot h^{\beta_l}$ and shares $(\llbracket \mathsf{R}_j \rrbracket_i, \llbracket r_j \rrbracket_i)$ to the nodes. Finally, using these commitments, each MPC node locally verifies its share by computing $g^{\llbracket \mathsf{R}_j \rrbracket_i} \cdot h^{\llbracket r_j \rrbracket_i} = C_j \times \prod_{l=1}^{\mathsf{T}-1} \left(g^{\alpha_l} \cdot h^{\beta_l}\right)^{(i^l)}$, where α_l and β_l are the coefficients of the Shamir polynomials. This expression ensures that each share is consistent with the attested commitment C_j. Notably, authenticated broadcast channels are used to ensure global consistency across nodes. Moreover, we assume that the underlying signature scheme satisfies Existential UnForgeability under Chosen-Message Attack (EUF-CMA).

Theorem 1. *If the discrete logarithm assumption holds and the underlying signature scheme is EUF-CMA secure, the MPC mode protocol from Fig. 3 securely realizes the MPC branch of the functionality $\mathcal{F}_{TLShare}$ (Fig. 1) in the stand-alone model in the $\mathcal{F}_{\mathsf{AuthData}}$-hybrid world.*

Given a functionality $\mathcal{F}_{\mathsf{AuthData}}$ for secure authenticated data [XYWY24], we show that the MPC mode protocol realizes $\mathcal{F}_{\mathrm{TLShare}}$ for malicious adversaries as stated in Theorem 1 in Section 4.1. The functionality $\mathcal{F}_{\mathsf{AuthData}}$ from [XYWY24] is defined in Fig. 4.

Functionality $\mathcal{F}_{\mathsf{AuthData}}$

This functionality interacts with a prover $\mathcal{P}$, a verifier $\mathcal{V}$, a server $\mathcal{S}$, and an adversary.

- Upon receiving $(\mathsf{sid}, \mathsf{Q}, \alpha, S)$ from $\mathcal{P}$ and $(\mathsf{sid}, \mathsf{Q})$ from $\mathcal{V}$, where sid is defined as a session identifier, Q is a query template and α is a private input for Q,
  1. Compute a query $Q := \mathsf{Q}(\alpha)$, and then sends a pair (sid, Q) to $\mathcal{S}$.
  2. Receives a response (sid, R) from $\mathcal{S}$ and then stores a tuple (sid, Q, R).
  3. Sends $(\mathsf{sid}, |Q|, |R|, S)$ to the adversary.
- Upon receiving $(\mathsf{commit}, \mathsf{sid}, \mathsf{cid})$ from $\mathcal{P}$, where cid is defined as a fresh commitment identifier, if a tuple (sid, Q, R) was previously stored, it is updated to $(\mathsf{sid}, \mathsf{cid}, Q, R)$, and send $(\mathsf{committed}, \mathsf{sid}, \mathsf{cid})$ to $\mathcal{V}$ and the adversary.
- Output $(\mathsf{sid}, \mathsf{cid}, Q, R)$ to $\mathcal{P}$ and $(\mathsf{sid}, \mathsf{cid}, S)$ to $\mathcal{V}$.

Figure 4: Functionality $\mathcal{F}_{\mathsf{AuthData}}$ for authenticating web data [XYWY24].

4.1 Security of MPC mode protocol

Recall Theorem 1. We now prove that the protocol in Fig. 3 securely realizes $\mathcal{F}_{\mathrm{TLShare}}$.

Proof. We need to address three separate corruption cases for the client, the attestor and $\mathsf{T}-1$ MPC nodes. In each case, we construct a PPT simulator given access to functionality $\mathcal{F}_{\mathrm{TLShare}}$ which runs a PPT adversary as a subroutine, and emulates functionality $\mathcal{F}_{\mathsf{AuthData}}$.

Case 1: Malicious client. Let $\mathcal{A}$ control the client. We consider the following sequence of hybrids.

- H_0: This is the real execution of the MPC mode protocol.
- H_1: Replace Π_{AuthData} by $\mathcal{F}_{\mathsf{AuthData}}$. The simulator obtains the authenticated response R, samples r_j, and sets $C_j = g^{\mathsf{R}_j} h^{r_j}$. Let $M = (\mathsf{ctx}, [C_1, \ldots, C_\ell])$ be the message signed by the honest attestor. The hybrids H_0 and H_1 are indistinguishable by the security of Π_{AuthData}.

- H_2: Proceed as in H_1, but abort if an honest node accepts (M', σ') for $M' \neq M$. A reduction $\mathcal{B}$, given access to the EUF-CMA signing oracle, obtains the signature on M and runs $\mathcal{A}$. Any accepted pair (M', σ') with $M' \neq M$ is then a forgery. Hence, H_1 and H_2 are indistinguishable under EUF-CMA security.

- H_3: Replace share delivery by the corrupt-client branch of $\mathcal{F}_{\text{TLShare}}$. The simulator performs the nodes' signature and share checks. For a completed sharing phase, let S contain the nodes that accept their complete share vectors.

  If $|S| < \mathsf{T}$, set $I = S$. Otherwise, choose $\mathsf{T}-1$ nodes from S. Using the authenticated openings known from H_1, the simulator interpolates polynomials F_j, G_j of degree less than T satisfying

  $$F_j(0) = \mathsf{R}_j, \qquad G_j(0) = r_j, \qquad F_j(i) = [\![\mathsf{R}_j]\!]_i, \quad G_j(i) = [\![r_j]\!]_i \quad (i \in I).$$

  These polynomials exist because there are at most T distinct interpolation points, including zero.

  If $I = S$, all accepted shares are already accounted for. Otherwise, the commitment polynomial induced by F_j, G_j agrees with the broadcast commitment polynomial at the T points in $\{0\} \cup I$. Vandermonde invertibility therefore implies equality of their coefficient commitments. Every additional accepting node consequently satisfies

  $$g^{[\![\mathsf{R}_j]\!]_i} h^{[\![r_j]\!]_i} = g^{F_j(i)} h^{G_j(i)}.$$

  A value share $[\![\mathsf{R}_j]\!]_i \neq F_j(i)$ would give two openings of the same Pedersen commitment to different values. Thus, computational binding ensures that all accepted shares lie on the constructed polynomials, except with negligible probability.

  The simulator submits F_j and the exact value shares for nodes in S to the functionality. For each remaining node, it submits the candidate values $F_j(i) + 1$ modulo q instead. Since the response is nonempty, these candidates cause that node to abort and are never released. The functionality therefore reproduces the accepted share vectors and local abort outcomes of H_2, except with negligible probability.

Thus, H_3 is the ideal execution, and the real and ideal executions are indistinguishable under the security of Π_{AuthData}, EUF-CMA security, and the discrete logarithm assumption.

Case 2: Malicious attestor. In this case the security reduces to the security of Π_{AuthData} since the attestor only interacts during this phase. The simulator finishes when receiving the signed commitments.

Case 3: Malicious $\mathsf{T}-1$ MPC nodes. Consider the following sequence of hybrids:

- H_0: This is the real execution of the MPC mode protocol in Fig. 3 with adversary controlling the $\mathsf{T}-1$ MPC nodes.

- H_1: The simulator computes fake commitments C'_i for a fake $\mathsf{R}' = 0$. Then, the simulator plays the role of the challenger in the EUF-CMA game and receives the signature σ corresponding to the commitments C'_i. Sends σ to the MPC nodes.

Indistinguishability between H_0 and H_1 comes from the hiding property of Pedersen commitments and the EUF-CMA property of the signature scheme.

- H_2: The simulator computes the Pedersen VSS for the fake $\mathsf{R}' = 0$ with base commitments given by C'_i.

Indistinguishability between H_1 and H_2 comes from the hiding property of the Pedersen VSS.

□

5 FHE Mode Protocols

We now describe how to generate an FHE ciphertext of a TLS payload in a verifiable and privacy-preserving manner. We present two solutions that can be instantiated with two popular types of FHE schemes: CGGI [CGGI20] and the LPR-style variant of Regev's scheme [Reg05]. As in the MPC mode, we leverage tools from GtP [XYWY24] to produce authenticated shares of the TLS payload.

5.1 FHE Mode with Non-Interactive ZKP (NIZK)

We present a protocol that generates an FHE ciphertext of a payload extracted from a TLS connection, instantiated with the CGGI scheme [CGGI20]. The protocol ensures two key security guarantees: 1) the validity of the CGGI resulting ciphertext, and 2) the consistency between the encrypted message and the corresponding payload from the TLS response R, as authenticated by the TLS phase.

To achieve these goals, our TLShare FHE with NIZK protocol integrates pairing-based vector commitments (see Section 2.4) as these can be used to prove that a tuple $(a_1, \ldots, a_n, b)$ is a valid CGGI encryption of a plaintext μ under an encryption key [LY10, FMRV22, BCD+25]. This construction exploits the algebraic structure of vector commitments to prove inner product relations between a secret vector (containing the message, encryption randomness, and noise) and a public vector (containing the public key and output ciphertext). Beyond proving the correctness of encryption, we also establish the consistency between the encrypted value and the TLS payload R. To this end, both the client and the attestor receive Pedersen commitments to each bit of R. However, this design introduces a technical challenge: the commitments used to encode the TLS payload are not directly compatible with the vector commitment constructions (Eq. (1)) required to prove encryption validity. In particular, the CGGI proof binds the ciphertext to a vector commitment but not to the attestor-signed TLS commitment.

To bridge this gap, we introduce a novel cryptographic component that we call *proof of vector membership*. This technique enables us to use Pedersen committed values with the NIZK proof of valid encryption, proving end-to-end provenance for FHE. Specifically, we allow the client to prove that the binary plaintext μ of a CGGI ciphertext is the same as the response R committed using Pedersen commitments, without revealing R or μ. Next, we describe the novel construction of our vector membership proof.

5.1.1 NIZK Proof of Vector Membership (VM)

We consider the same generators $\hat{g}$ and $\hat{g}_1$ derived by the CRS (see Section 2.4) to enable composability between the vector membership and CGGI encryption NIZKs. Consider a Pedersen commitment $\hat{C} = \hat{g}^r \cdot \hat{g}_1^{\mu}$ of a message μ in the second source pairing group $\hat{\mathbb{G}}$, where μ has binary expansion given by $\mu = (\mu_1, \ldots, \mu_l)$. Also, let $\boldsymbol{w} = (w_1, \ldots, w_n) \in \mathbb{Z}_p^n$ contain μ as a contiguous segment, where the CGGI argument additionally proves that $\boldsymbol{w}$ is binary. More formally, $\mu = \sum_{i=1}^{l} 2^{i-1} w_{k+i}$ for $0 \leq k$, $1 \leq l$, $k + l \leq n$, and $2^l < p$. We note that the prover does not choose k and l. Following the binary TFHE witness of Bootland et al. [BCD+25, Sections 4.11.2 and 7.6.3], we fix the canonical order $\boldsymbol{w} = (\boldsymbol{r} \parallel \boldsymbol{E}_1 \parallel \boldsymbol{e}_2 \parallel \mathsf{bits}_l(\mu) \parallel \boldsymbol{w}_{\mathsf{aux}})$, where d is the LWE dimension, $\boldsymbol{E}_1 \in \{0,1\}^{d \times (3d+2)}$, $\boldsymbol{e}_2 \in \{0,1\}^{3d+2}$, bits_l is least-significant-bit first, and $\boldsymbol{w}_{\mathsf{aux}}$ contains the auxiliary proof bits and padding. Hence, k is the number of bits preceding the message segment in this encoding, and the verifier rejects any other value. Finally, let $\hat{V}$ be a vector commitment to the binary vector $\boldsymbol{w}$ as $\hat{g}^{\gamma} \cdot \prod_{j=1}^{n} \hat{g}_j^{w_j}$ from Eq. (1).

In our NIZK proof, the prover aims to convince the verifier that the binary representation of μ (i.e., $\mu_1, \ldots, \mu_l$) corresponds to a contiguous segment within the binary vector $\boldsymbol{w}$ committed in the vector commitment $\hat{V}$, while keeping μ and $\boldsymbol{w}$ private. In other words,

the prover wants to convince the verifier that μ is a member of $\boldsymbol{w}$, with $\hat{C}$ and $\hat{V}$ being their public commitments, respectively. More formally, we define the vector membership (VM) relation as

$$\mathcal{R}_{\mathsf{pp}} = \Big\{ \big((\hat{C}, \hat{V}, k, l), (r, \gamma, \mu, \boldsymbol{w})\big) \;\Big|\; \hat{C} = \hat{g}^r \cdot \hat{g}_1^{\mu} \;\wedge\; \hat{V} = \hat{g}^{\gamma} \cdot \textstyle\prod_{j=1}^{n} \hat{g}_j^{w_j} \\ \wedge\; \mu = \textstyle\sum_{i=1}^{l} 2^{i-1} w_{k+i} \pmod p \Big\}.$$

where the public parameters pp denotes the CRS containing the commitment key $(\hat{g}, \{\hat{g}_i\}_{i=1}^n)$ and the pairing groups $(\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T)$. The NIZK argument for the VM relation $\mathcal{R}_{\mathsf{pp}}$ is:

- $\mathsf{VM.GenCRS}(1^\lambda, 1^n)$: on input a security parameter λ and the maximal bit length n, proceed as follows. Choose asymmetric bilinear groups $(\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T)$ of prime order $p > 2^{\ell(\lambda)}$, for some polynomial function $\ell : \mathbb{N} \to \mathbb{N}$, and sample $g \xleftarrow{\$} \mathbb{G}$, $\hat{g} \xleftarrow{\$} \hat{\mathbb{G}}$. Sample $\alpha \xleftarrow{\$} \mathbb{Z}_p$, set $g_i = g^{\alpha^i}$ and $\hat{g}_i = \hat{g}^{\alpha^i}$, and retain $\tau = \alpha$ as the simulation trapdoor. The public parameters are defined as:

$$\mathsf{pp} = \big((\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T),\; g,\; \hat{g}, \\ \{g_i\}_{i \in [1,2n] \setminus \{n+1\}},\; \{\hat{g}_i\}_{i \in [1,n]}\big).$$

  Return (pp, τ), where only pp is published and τ is given exclusively to the zero-knowledge simulator.

- $\mathsf{VM.Com}_{\mathsf{pp}}(\mu, (\hat{V}, \gamma))$: on input $\mu \in \mathbb{Z}_p$ and $\hat{V} \in \hat{\mathbb{G}}$, the prover chooses a random $r \xleftarrow{\$} \mathbb{Z}_p$ and computes the commitment $\hat{C} = \hat{g}^r \cdot \hat{g}_1^{\mu} \in \hat{\mathbb{G}}$. Return $\mathsf{com} = (\hat{C}, \hat{V}) \in \hat{\mathbb{G}}^2$ and the opening information $\mathsf{aux} = (r, \gamma) \in \mathbb{Z}_p^2$.

- $\mathsf{VM.Prove}_{\mathsf{pp}}((\mathsf{com}, k, l), (\mu, \boldsymbol{w}, \mathsf{aux}))$: given $\mathsf{com} = (\hat{C}, \hat{V}) \in \hat{\mathbb{G}}^2$, indexes k and l, and witness information (μ, w, aux), where $\mathsf{aux} = (r, \gamma)$ and μ has binary expansion $(\mu_1, \ldots, \mu_l)$ and $\boldsymbol{w}$ is a binary vector, compute the following proof π:

$$g_n^{-r} \cdot \textstyle\prod_{i \in L} (g_{n+1-i}^{\gamma} \cdot \prod_{j \in [1,n] \setminus \{i\}} g_{n+1+j-i}^{w_j})^{2^{i-1-k}}, \tag{2}$$

  where $L = \{k+1, \ldots, k+l\}$.

- $\mathsf{VM.Verify}_{\mathsf{pp}}((\mathsf{com}, k, l), \pi)$: given $\mathsf{com} = (\hat{C}, \hat{V}) \in \hat{\mathbb{G}}^2$ and proof π, return True if and only if

$$e(\textstyle\prod_{i \in L} g_{n+1-i}^{2^{i-1-k}}, \hat{V}) / e(g_n, \hat{C}) = e\,(\pi, \hat{g})\,. \tag{3}$$

Definition 1. Let $L = \{k+1, \ldots, k+l\}$ and $t_i = 2^{i-1-k}$. We say that the VM NIZK argument is *computationally sound with respect to fixed openings* if, for every algebraic PPT adversary $\mathcal{A}$, after sampling $(\mathsf{pp}, \cdot) \leftarrow \mathsf{VM.GenCRS}(1^\lambda, 1^n)$ and $(\hat{C}, \hat{V}, k, l, \pi, r, \mu, \gamma, \boldsymbol{w}) \leftarrow \mathcal{A}(\mathsf{pp})$,

$$\Pr\left[\begin{matrix} \mathsf{VM.Verify}_{\mathsf{pp}}((\hat{C}, \hat{V}, k, l), \pi) = 1 \;\wedge\; \hat{C} = \hat{g}^r \hat{g}_1^{\mu} \\ \wedge\; \hat{V} = \hat{g}^{\gamma} \textstyle\prod_{j=1}^{n} \hat{g}_j^{w_j} \;\wedge\; \mu \neq \sum_{i \in L} t_i w_i \end{matrix}\right] \leq \mathsf{negl}(\lambda).$$

Moreover, VM NIZK argument is *perfectly zero-knowledge in the AGM* if there exists a PPT simulator Sim such that, for $(\mathsf{pp}, \tau) \leftarrow \mathsf{VM.GenCRS}(1^\lambda, 1^n)$ and every algebraically generated $(x, w) \in \mathcal{R}_{\mathsf{pp}}$, where $x = (\hat{C}, \hat{V}, k, l)$, the distributions

$$\mathsf{VM.Prove}_{\mathsf{pp}}(x, w) \quad \text{and} \quad \mathsf{Sim}(\mathsf{pp}, \tau, x, \mathsf{repr}(x))$$

are identical, where $\mathsf{repr}(x)$ denotes the algebraic representations of the commitments $\hat{C}$ and $\hat{V}$.

Correctness of Proof of Membership. Recall that $\hat{C} = \hat{g}^r \hat{g}_1^{\mu}$ commits to the message μ, whereas $\hat{V} = \hat{g}^{\gamma} \prod_{j=1}^{n} \hat{g}_j^{w_j}$ commits to the vector $\boldsymbol{w}$. By bilinearity [Lib24], these commitments satisfy

$$\begin{aligned} e(g_{n+1-i}, \hat{V}) &= e(g_{n+1-i}^{\gamma} \cdot \textstyle\prod_{j\in[1,n]\setminus\{i\}} g_{n+1+j-i}^{w_j}, \hat{g}) \cdot e(g_1, \hat{g}_n)^{w_i}, \\ e(g_n, \hat{C}) &= e(g_n^r, \hat{g}) \cdot e(g_1, \hat{g}_n)^{\mu}. \end{aligned} \tag{4}$$

Proof. To check the correctness of the proof, we need to guarantee that Eq. (3) holds. Let $L = \{k+1, \ldots, k+l\}$. For a valid witness, the membership relation gives $\mu = \sum_{i\in L} 2^{i-1-k} w_i$. Observe that:

$$\begin{aligned} & e(\textstyle\prod_{i\in L} g_{n+1-i}^{2^{i-1-k}}, \hat{V}) \\ &= \textstyle\prod_{i\in L} e(g_{n+1-i}, \hat{g}^{\gamma} \cdot \prod_{j=1}^{n} \hat{g}_j^{w_j})^{2^{i-1-k}} \\ &= \textstyle\prod_{i\in L} e(g^{\gamma} \cdot \prod_{j=1}^{n} g_j^{w_j}, \hat{g}_{n+1-i})^{2^{i-1-k}} \\ &= \textstyle\prod_{i\in L} e(g_{n+1-i}^{\gamma} \cdot \prod_{j=1}^{n} g_{n+1+j-i}^{w_j}, \hat{g})^{2^{i-1-k}} \\ &= \textstyle\prod_{i\in L} e(g_{n+1-i}^{\gamma} \cdot g_{n+1}^{w_i} \cdot \prod_{j\in[1,n]\setminus\{i\}} g_{n+1+j-i}^{w_j}, \hat{g})^{2^{i-1-k}} \\ &= \textstyle\prod_{i\in L} e(g_{n+1}, \hat{g})^{2^{i-1-k}\cdot w_i} \cdot \prod_{i\in L} e(g_{n+1-i}^{\gamma} \cdot \prod_{j\in[1,n]\setminus\{i\}} g_{n+1+j-i}^{w_j}, \hat{g})^{2^{i-1-k}}. \end{aligned}$$

Since $e(g_{n+1}, \hat{g}) = e(g_1, \hat{g}_n)$, we have $\prod_{i\in L} e(g_{n+1}, \hat{g})^{2^{i-1-k} w_i} = e(g_1, \hat{g}_n)^{\sum_{i\in L} 2^{i-1-k} w_i} = e(g_1, \hat{g}_n)^{\mu}$. Moreover, Eq. (4) gives $e(g_n, \hat{C}) = e(g_n^r, \hat{g}) \cdot e(g_1, \hat{g}_n)^{\mu}$. Therefore, dividing the two expressions $e(\prod_{i\in L} g_{n+1-i}^{2^{i-1-k}}, \hat{V}) / e(g_n, \hat{C})$ yields

$$\begin{aligned} & \frac{e(g_1, \hat{g}_n)^{w} \cdot \prod_{i\in L} e(g_{n+1-i}^{\gamma} \cdot \prod_{j\in[1,n]\setminus\{i\}} g_{n+1+j-i}^{w_j}, \hat{g})^{2^{i-1-k}}}{e(g_1, \hat{g}_n)^{\mu} \cdot e(g_n^r, \hat{g})} \\ &= e(\textstyle\prod_{i\in L} [g_{n+1-i}^{\gamma} \cdot \prod_{j\in[1,n]\setminus\{i\}} g_{n+1+j-i}^{w_j}]^{2^{i-1-k}}, \hat{g}) \cdot e(g_n^{-r}, \hat{g}) \\ &= e(g_n^{-r} \cdot \textstyle\prod_{i\in L} [g_{n+1-i}^{\gamma} \cdot \prod_{j\in[1,n]\setminus\{i\}} g_{n+1+j-i}^{w_j}]^{2^{i-1-k}}, \hat{g}) = e(\pi, \hat{g}), \end{aligned}$$

where π is the proof defined in Eq. (2). Thus, Eq. (3) holds. □

Security of NIZK of Vector Membership. We prove the security of NIZK of Vector Membership in the Algebraic Group Model (AGM) [FKL18], which is an idealized model where all algorithms are assumed to be algebraic. More specifically, it is assumed that whenever a group element is generated, a representation as a function of previously observed group elements in the same group is also provided.

Theorem 2. *Under the $(2n, n)$-DLOG assumption, VM is computationally sound with respect to fixed openings against algebraic PPT adversaries and perfectly zero-knowledge in the AGM.*

Proof. We analyze the soundness and zero-knowledge properties separately.

Soundness. Let $L = \{k+1, \ldots, k+l\}$ and $t_i = 2^{i-1-k}$ for $i \in L$, and define $T(X) = \sum_{i\in L} t_i X^{n+1-i}$ and $W(X) = \sum_{j=1}^{n} w_j X^j$. Assume that an algebraic adversary outputs commitments with openings (r, μ) and $(\gamma, \boldsymbol{w})$, and an accepting proof π, but $\mu \neq \sum_{i\in L} t_i w_i$. The fixed openings imply $\hat{C} = \hat{g}^{r+\mu\alpha}$ and $\hat{V} = \hat{g}^{\gamma+W(\alpha)}$. By the AGM, the proof has a representation $\pi = g^{P(\alpha)}$ for

$$P(X) = p_0 + \sum_{d\in[1,2n]\setminus\{n+1\}} p_d X^d,$$

TLShare: FHE mode with Non-Interactive Zero-Knowledge (NIZK) Proofs

This protocol runs between $\mathcal{S}$, a client, an attestor, and FHE node $\mathcal{N}_1$. $\mathcal{S}$ behaves the same as in a standard TLS connection. The protocol requires a common reference string (CRS) computed in a setup phase by a trusted third party and accessible to both the client and $\mathcal{N}_1$. $\hat{V}$ is a vector commitment as in Eq. (1), and $\hat{C}$ is a Pedersen commitment [Ped92], over the second source group $\hat{\mathbb{G}}$ of pairing $e : \mathbb{G} \times \hat{\mathbb{G}} \to \mathbb{G}_T$.

- **Public Parameters:** a query template (Query) for the server $\mathcal{S}$ (e.g., a `GET` request).
- **Session Context:** Before the TLS phase, the parties agree on a fresh sid. For response length ℓ, define $\mathsf{ctx} \coloneqq \mathsf{Encode}(\texttt{TLShare-fhe-nizk}, \mathsf{sid}, \mathsf{cert}_{\mathcal{S}}, \mathsf{Query}, \ell, \mathsf{pk}_{\mathsf{fhe}}, \mathsf{pp}, \mathcal{N}_1)$.
- **Server $\mathcal{S}$:** has a secret key ($sk_{\mathcal{S}}$) for a signature scheme (e.g., ECDSA), a certificate $\mathsf{cert}_{\mathcal{S}}$, and hosts and operates a public website.
- **Client:** has a private input Inp (e.g., a password).
- **Attestor:** owns a key pair $(pk_{\mathcal{A}}, sk_{\mathcal{A}})$ to a signature scheme (e.g., for ECDSA).
- **FHE Node** $\mathcal{N}_1$ has no input.

TLS Phase (Handshake, Query, and Response):

1. The client, $\mathcal{S}$, and attestor run $\Pi_{\mathsf{AuthData}}(sk_{\mathcal{S}}, \mathsf{cert}_{\mathcal{S}}, \mathsf{Inp}, \mathsf{Query})$ instantiated with Pedersen commitments over group $\hat{\mathbb{G}}$, i.e., converting the IT-MACs of $[\![\mathsf{Q}]\!]$ and $[\![\mathsf{R}]\!]$ to commitments over them. Let $\hat{C}_j$ be a commitment to the jth bit of R with randomness r_j. The client outputs $(\mathsf{sid}, [\hat{C}_1, \ldots, \hat{C}_\ell], \mathsf{Q}, \mathsf{R})$ and the attestor outputs $(\mathsf{sid}, [\hat{C}_1, \ldots, \hat{C}_\ell])$, where $\ell \coloneqq |\mathsf{R}|$.
2. The attestor signs the context and ordered committed response, computing $\sigma \coloneqq \mathsf{Sign}(sk_{\mathcal{A}}, (\mathsf{ctx}, [\hat{C}_1, \ldots, \hat{C}_\ell]))$ and sends it to the client.

Share/Encryption Phase:

3. The client sends $(\sigma, [\hat{C}_1, \ldots, \hat{C}_\ell])$ to $\mathcal{N}_1$.
4. $\mathcal{N}_1$ verifies $\mathsf{Verify}(pk_{\mathcal{A}}, (\mathsf{ctx}, [\hat{C}_1, \ldots, \hat{C}_\ell]), \sigma)$ and aborts in case it fails.
5. For each bit of the response R_j, $j \in [1, \ell]$:
   a. The client locally runs the CGGI proving algorithm to obtain
   $$(\mathsf{ct}_j, \pi_{\mathsf{CGGI}}, \hat{V}_j, \boldsymbol{w}_j, \gamma_j) \leftarrow \mathsf{CGGI.Prove}_{\mathsf{pp}}(\mathsf{R}_j),$$
   where ct_j is the public CGGI ciphertext whose validity is proved by π_{CGGI}.
   b. Then, the client locally sets $\mathsf{com} = (\hat{C}_j, \hat{V}_j) \in \hat{\mathbb{G}}^2$ and the opening information $\mathsf{aux} = (r_j, \gamma_j) \in \mathbb{Z}_p^2$. Finally, the client locally runs $\pi_{\mathsf{VM}} \leftarrow \mathsf{VM.Prove}_{\mathsf{pp}}(\mathsf{com}, k, l, (\mathsf{R}_j, \boldsymbol{w}_j, \mathsf{aux}))$ for index k, bitlength l, binary vector $\boldsymbol{w}_j$ containing the response R_j, encryption randomness and noise, and π_{VM} given by Eq. (2).
   c. The client sends the tuple $(\mathsf{ct}_j, \hat{V}_j, \pi_{\mathsf{VM}}, \pi_{\mathsf{CGGI}}, k, l)$ to the FHE node $\mathcal{N}_1$ for verification.
   d. $\mathcal{N}_1$ runs $\mathsf{CGGI.Verify}_{\mathsf{pp}}(\mathsf{ct}_j, \hat{V}_j, \pi_{\mathsf{CGGI}})$ to check that ct_j is a valid CGGI ciphertext for the witness committed in $\hat{V}_j$; and finally runs the verification algorithm $\mathsf{VM.Verify}_{\mathsf{pp}}(\mathsf{com}, k, l, \pi_{\mathsf{VM}})$ by checking Eq. (3). $\mathcal{N}_1$ also checks that (k, l) is the pair prescribed by the public CGGI witness layout and aborts if any check fails.

Figure 5: TLShare protocol with $\mathsf{mode} = \mathsf{fhe}$ with NIZK for the CGGI scheme.

whose coefficient of X^{n+1} is zero because g_{n+1} is omitted from the CRS. Moreover, $\prod_{i \in L} g_{n+1-i}^{t_i} = g^{T(\alpha)}$. Substituting these expressions into Eq. (3) and using bilinearity gives

$$e(g, \hat{g})^{T(\alpha)(\gamma + W(\alpha)) - \alpha^n (r + \mu\alpha)} = e(g, \hat{g})^{P(\alpha)}.$$

Since $e(g, \hat{g})$ generates $\mathbb{G}_T$, it follows that $F(\alpha) = 0$ for $F(X) = T(X)(\gamma + W(X)) -$

$X^n(r+\mu X) - P(X)$. Writing $[X^d]Q$ for the coefficient of X^d in $Q(X)$, we have

$$[X^{n+1}]\big(T(X)W(X)\big) = \sum_{i\in L} t_i w_i \quad \text{and} \quad [X^{n+1}]\big(X^n(r+\mu X)\big) = \mu.$$

Since neither $\gamma T(X)$ nor $P(X)$ contains an X^{n+1} term, $[X^{n+1}]F = \sum_{i\in L} t_i w_i - \mu \neq 0$. Thus, F is a non-zero polynomial of degree at most $2n$.

We now construct a reduction $\mathcal{B}$ that, given a $(2n, n)$-DLOG instance, gives the adversary the corresponding CRS with g_{n+1} omitted. By the definition of soundness with respect to fixed openings, the openings are part of the adversary's output. Together with the public values k, l and the AGM representation of π, they determine every coefficient of F, although α remains unknown. If the adversary violates soundness, $\mathcal{B}$ finds the roots of F over $\mathbb{Z}_p$. Since $F(\alpha) = 0$, one of these roots is α. $\mathcal{B}$ identifies it by testing $g^\beta = g_1$ for each root β, and outputs the unique value that passes. Hence, a non-negligible soundness violation yields a non-negligible advantage against the $(2n, n)$-DLOG assumption.

Zero-knowledge. Let $(\mathsf{pp}, \alpha) \leftarrow \mathsf{VM.GenCRS}(1^\lambda, 1^n)$ and $x = (\hat{C}, \hat{V}, k, l)$ be an algebraically generated true statement. Under Definition 1, the simulator receives pp, α, x, and $\mathsf{repr}(x)$, but not the witness. We assume that setup samples nonidentity generators $g, \hat{g}$ and we consider $\hat{g}_0 = \hat{g}$.

The supplied representations give coefficients $c_j, v_j \in \mathbb{Z}_p$ such that $\hat{C} = \prod_{j=0}^n \hat{g}_j^{c_j}$ and $\hat{V} = \prod_{j=0}^n \hat{g}_j^{v_j}$. Define $f_C(X) = \sum_{j=0}^n c_j X^j$, $f_V(X) = \sum_{j=0}^n v_j X^j$, and $T(X) = \sum_{i\in L} t_i X^{n+1-i}$, where $L = \{k+1, \ldots, k+l\}$ and $t_i = 2^{i-1-k}$. Thus, $\hat{C} = \hat{g}^{f_C(\alpha)}$ and $\hat{V} = \hat{g}^{f_V(\alpha)}$. Using α, the simulator computes $\pi_{\mathsf{sim}} = g^{T(\alpha)f_V(\alpha) - \alpha^n f_C(\alpha)}$ in polynomial time, with all exponent arithmetic over $\mathbb{Z}_p$.

By bilinearity, $e(\pi_{\mathsf{sim}}, \hat{g}) = e(g^{T(\alpha)}, \hat{V})/e(g_n, \hat{C})$. Since $g^{T(\alpha)} = \prod_{i\in L} g_{n+1-i}^{t_i}$, the simulated proof satisfies Eq. (3). The real proof satisfies the same equation by correctness. Nondegeneracy and the fact that $\hat{g}$ is a generator make $Y \mapsto e(Y, \hat{g})$ injective, so $\pi_{\mathsf{sim}} = \pi_{\mathsf{real}}$.

This equality holds for every valid statement, witness, and supplied representation. Since simulation preserves the CRS distribution, the joint real and simulated views, including auxiliary information, are identical. Hence, VM is perfectly zero-knowledge under the stated AGM definition. □

5.1.2 TLShare for CGGI encryption

We now combine the NIZK proof of vector membership with the NIZK proof of valid CGGI encryption [FMRV22, Lib24, BCD+25], and the Pedersen commitments from the TLShare TLS phase. As a result, we perform a verifiable FHE encryption of a plaintext that can be provably linked to a TLS payload. We adopt the public-key NIZK argument for valid CGGI encryption as described by Bootland et al. [BCD+25, Section 7.6.3]. Let $(\mathsf{CGGI.Prove}_{\mathsf{pp}}, \mathsf{CGGI.Verify}_{\mathsf{pp}})$ denote the proving and verification algorithms of this NIZK proof. Our full TLShare protocol combining the VM and the CGGI NIZK proofs is summarized in Fig. 5.

In step 1 of Fig. 5, the server, client, and attestor run Π_{AuthData} from GtP [XYWY24], using Pedersen commitments. This results in both the client and the attestor receiving Pedersen commitments to the bits of the TLS payload R. To enable external verification of the CGGI encryption, the attestor signs the committed response in a signature σ and forwards it to the client in step 2. In step 3, the client sends the signature σ and the commitments $[\hat{C}_1, \ldots, \hat{C}_\ell]$ to the node. The node verifies them in step 4.

The encryption phase involves only one-way communication from the client to the FHE node. At this stage, the client locally generates the required proofs: the NIZK argument for valid CGGI encryption (step 5.*a*) and the NIZK argument of vector membership (step 5.*b*). Specifically, for each bit j of the payload R_j, the client performs the following steps.

It runs the CGGI proving algorithm to obtain $(\mathsf{ct}_j, \pi_{\mathsf{CGGI}}, \hat{V}_j, \boldsymbol{w}_j, \gamma_j) \leftarrow \mathsf{CGGI.Prove}_{\mathsf{pp}}(\mathsf{R}_j)$, where ct_j is the public ciphertext proved valid by π_{CGGI}, and $\boldsymbol{w}_j$ is the vector committed under $\hat{V}_j$ containing R_j and the encryption coins. The client takes the commitment $\hat{C}_j$ corresponding to R_j and the vector commitment $\hat{V}_j$ from the previous step. Then, the client uses them as the commitments $\mathsf{com} = (\hat{C}_j, \hat{V}_j)$ from $\mathsf{VM.Com}$ and the randomness of the commitments as $\mathsf{aux} = (r_j, \gamma_j)$. In other words, the NIZK argument for valid CGGI encryption [BCD+25] internally commits to the vector $\boldsymbol{w}$ through the value $\hat{V}$ during the proof generation. Therefore, we reuse the vector commitment $\hat{V}$ as part of the NIZK argument for vector membership. The client continues in step 5.*b* of Fig. 5 by generating the proof of vector membership via $\pi_{\mathsf{VM}} \leftarrow \mathsf{VM.Prove}_{\mathsf{pp}}(\mathsf{com}, k, l, (\mathsf{R}_j, \boldsymbol{w}_j, \mathsf{aux}))$. In step 5.*c*), the client sends $(\mathsf{ct}_j, \hat{V}_j, \pi_{\mathsf{VM}}, \pi_{\mathsf{CGGI}}, k, l)$ to the FHE node. Finally, in step 5.*d*), the node verifies that π_{CGGI} proves validity of this same ciphertext and vector commitment, and that π_{VM} links the committed plaintext to $\hat{C}_j$.

5.1.3 Security of TLShare FHE Mode with NIZK Proofs

We begin by observing that the vector-membership proof $\mathsf{Prove}_{\mathsf{pp}}$ presupposes that the vector commitment $\hat{V}$ indeed binds to a *binary* vector $\boldsymbol{w}$. This is ensured by the $\mathsf{CGGI.Prove}_{\mathsf{pp}}$ algorithm, which outputs both $\hat{V}$ and a proof component within π_{CGGI} that certifies each coordinate of $\boldsymbol{w}$ is in $\{0, 1\}$. We state the security of TLShare protocol with $\mathsf{mode} = \mathsf{fhe}$ with NIZK in Theorem 3 and prove its security below.

Theorem 3. *If the LWE and $(2n, n)-$DLOG assumptions hold, and the underlying signature scheme is EUF-CMA secure, then, in the AGM+ROM against algebraic PPT adversaries, the FHE mode protocol with NIZK from Fig. 5 securely realizes the functionality $\mathcal{F}_{TLShare}$ (Fig. 1) in the stand-alone model in the $\mathcal{F}_{AuthData}$-hybrid world.*

Proof. We need to address three separate corruption cases for the client, the attestor, and the FHE node $\mathcal{N}_1$. In each case, we construct a PPT simulator given access to functionality $\mathcal{F}_{\mathsf{TLShare}}$ which runs a PPT adversary as a subroutine, and emulates functionality $\mathcal{F}_{\mathsf{AuthData}}$.

Case 1: Malicious Client. Let $\mathcal{A}$ control the client. We consider the following sequence of hybrids.

- H_0: This is the real execution of the FHE mode with NIZK protocol.
- H_1: Replace Π_{AuthData} by $\mathcal{F}_{\mathsf{AuthData}}$. The simulator obtains the authenticated response R, samples r_j, and sets $\hat{C}_j = \hat{g}^{r_j} \hat{g}_1^{\mathsf{R}_j}$. Let $M = (\mathsf{ctx}, [\hat{C}_1, \ldots, \hat{C}_\ell])$ be the message signed by the honest attestor. The hybrids H_0 and H_1 are indistinguishable by the security of Π_{AuthData}.
- H_2: Proceed as in H_1, but abort if the honest FHE node accepts (M', σ') for $M' \neq M$. A reduction $\mathcal{B}$, given access to the EUF-CMA signing oracle, obtains the signature on M and runs $\mathcal{A}$. Any accepted pair (M', σ') with $M' \neq M$ is then a forgery. Hence, H_1 and H_2 are indistinguishable under EUF-CMA security.
- H_3: Replace the ciphertext delivery by the corrupt-client branch of $\mathcal{F}_{\mathsf{TLShare}}$. The simulator receives the tuples $(\mathsf{ct}_j, \hat{V}_j, \pi_{\mathsf{VM},j}, \pi_{\mathsf{CGGI},j}, k, l)$ sent to the honest FHE node and performs all its checks. If any check fails, it causes the functionality to abort. Otherwise, simulation-extractable knowledge soundness of the CGGI argument in the AGM+ROM [BCD+25] yields an opening $(\boldsymbol{w}_j, \gamma_j)$ of $\hat{V}_j$ and encryption coins ω_j for the same ciphertext ct_j. Since the simulator knows the authenticated opening (R_j, r_j) of $\hat{C}_j$, VM soundness with respect to these fixed openings implies that ct_j encrypts R_j, except with negligible probability. Thus, $\mathcal{R}_{\mathsf{Enc}}(\mathsf{pk}_{\mathsf{fhe}}, \mathsf{R}_j; \mathsf{ct}_j, \omega_j) = 1$. The simulator submits the exact pairs $(\mathsf{ct}_j, \omega_j)$ to the functionality, which forwards the same ciphertexts to $\mathcal{N}_1$.

Thus, H_3 is the ideal execution, and the real and ideal executions are indistinguishable under the security of Π_{AuthData}, EUF-CMA security, CGGI knowledge soundness, and VM soundness.

Case 2: Malicious Attestor. Similarly to the security proof of the MPC mode protocol under a malicious attestor (Section 4.1), the security reduces to the security of Π_{AuthData} as the attestor only interacts during this phase. The simulator finishes when receiving the signed commitments.

Case 3: Malicious FHE Node $\mathcal{N}_1$. Consider the following sequence of hybrids:

- H_0: This is the real execution of the FHE mode with NIZK protocol in Fig. 5 with the adversary controlling $\mathcal{N}_1$.
- H_1: The simulator computes fake Pedersen commitments $[\hat{C}'_1, \ldots, \hat{C}'_\ell]$ for a fake response $R' = 0^\ell$ (all zeros). Then, the simulator plays the role of the challenger in the EUF-CMA game and receives the signature σ' corresponding to the commitments $[\hat{C}'_1, \ldots, \hat{C}'_\ell]$. The simulator sends both $[\hat{C}'_1, \ldots, \hat{C}'_\ell]$ and σ' to the FHE node $\mathcal{N}_1$.

Indistinguishability between H_0 and H_1 comes from the hiding property of Pedersen commitments and the EUF-CMA property of the signature scheme. The malicious FHE node cannot distinguish between commitments to different values due to the information-theoretic hiding property of Pedersen commitments.

- H_2: For each bit $R'_j = 0$, the simulator computes a fake CGGI ciphertext (encoded in $\hat{V}'_j$) and generates valid NIZK proofs π'_{CGGI} and π'_{VM} for the dummy value. The simulator sends $(\hat{V}'_j, \pi'_{\mathsf{VM}}, \pi'_{\mathsf{CGGI}}, k, l)$ to the FHE node for verification.

Indistinguishability between H_1 and H_2 comes from the semantic security (IND-CPA) of the CGGI encryption scheme and the zero-knowledge property of the NIZK arguments. The malicious FHE node cannot distinguish between encryptions of different plaintexts under the LWE assumption, and the NIZK proofs reveal nothing about the witness due to zero-knowledge.

□

5.2 FHE Mode with Interactive ZKP (IZK)

In this section, we describe an interactive zero-knowledge (IZK) approach to securely encrypt a TLS payload message R into an FHE ciphertext. We focus on an LPR-style variant of Regev's LWE scheme [Pri25] with ciphertext and plaintext moduli q and p, and public key encryption of a message μ:

$$(\boldsymbol{a}, b) = (\boldsymbol{A}\boldsymbol{r} + \boldsymbol{e}, \boldsymbol{y}^\top \boldsymbol{r} + e_0 + \lfloor \mu \cdot \Delta \rfloor) \bmod q, \tag{5}$$

under public key $(\boldsymbol{A}, \boldsymbol{y})$, where $\Delta = q/p$ is the scaling factor between plaintext and ciphertext spaces, $\boldsymbol{r} \xleftarrow{\$} \{0,1\}^n$ is a binary randomness vector, $\boldsymbol{e} \xleftarrow{\$} \chi^n$, $e_0 \xleftarrow{\$} \chi$ are noise samples, and χ is a centered discrete Gaussian with standard deviation stdev over $\mathbb{Z}_q$.

At a high level, our FHE mode with IZK protocol derives authenticated shares (i.e., IT-MACs) of the server's private response from the TLS phase. Then, we compose it with an IT-MAC-based IZK protocol to prove the validity of an LWE encryption circuit. We use the efficient QuickSilver protocol [YSWW21] together with the GtP framework to achieve a constant number of rounds and a communication overhead near one field element per multiplication gate [YSWW21]. Our full IZK protocol is summarized in Fig. 6.

Recall, in the MPC mode and the FHE mode with NIZK, the Π_{AuthData} [XYWY24, Figs. 17–18] in the TLS phase outputs Pedersen commitments of the TLS payload R.

TLShare: FHE mode with Interactive Zero-Knowledge (IZK) Proofs

This protocol runs between a server $\mathcal{S}$, a client, an attestor, and FHE node $\mathcal{N}_1$. $\mathcal{S}$ behaves the same as in a standard TLS connection. We use $\langle\langle x \rangle\rangle$ to denote authenticated shares (i.e., IT-MACs) of x, as defined in Section 2.3.

- **Public Parameters:** a query template (Query) for the server $\mathcal{S}$ (e.g., a `GET` request). An LPR-style Regev public key $(\boldsymbol{A}, \boldsymbol{y})$ and a plaintext modulus p as defined in Section 2.2.
- **Session Context:** Before the TLS phase, the parties agree on a fresh sid. For response length ℓ, define $\mathsf{ctx} \coloneqq \mathsf{Encode}(\texttt{TLShare-fhe-izk}, \mathsf{sid}, \mathsf{cert}_{\mathcal{S}}, \mathsf{Query}, \ell, \mathsf{pk}_{\mathsf{fhe}}, \mathcal{N}_1)$.
- **Server $\mathcal{S}$:** has a secret key ($sk_{\mathcal{S}}$) for a signature scheme (e.g., ECDSA), a certificate $\mathsf{cert}_{\mathcal{S}}$, and hosts and operates a public website.
- **Client:** has a private input Inp (e.g., a password).
- **Attestor:** owns a key pair $(pk_{\mathcal{A}}, sk_{\mathcal{A}})$ to a signature scheme (e.g., for ECDSA).
- **FHE Node** $\mathcal{N}_1$ has no input.

TLS Phase (Handshake, Query, and Response):

1. The client, $\mathcal{S}$, and attestor run $\Pi_{\mathsf{AuthData}}(sk_{\mathcal{S}}, \mathsf{cert}_{\mathcal{S}}, \mathsf{Inp}, \mathsf{Query})$ and retrieve the IT-MACs of the query and response as $\langle\langle \mathsf{Q} \rangle\rangle$ and $\langle\langle \mathsf{R} \rangle\rangle = [\langle\langle \mathsf{R}_1 \rangle\rangle, \ldots, \langle\langle \mathsf{R}_\ell \rangle\rangle]$. We denote by ℓ, the length of the response, i.e., $\ell \coloneqq |\mathsf{R}|$.

Share/Encryption Phase:

2. For each bit of the response R_j, $j \in [1, \ell]$:
   a. The client locally generates a binary randomness vector, $\boldsymbol{r} \xleftarrow{\$} \{0,1\}^n$; and noise samples $\boldsymbol{e} \xleftarrow{\$} \chi^n$, $e_0 \xleftarrow{\$} \chi$, where χ is a centered discrete Gaussian with standard deviation stdev over $\mathbb{Z}_q$.
   b. The client and the attestor generate IT-MACs of the secret inputs of the client as $\langle\langle \boldsymbol{r} \rangle\rangle, \langle\langle \boldsymbol{e} \rangle\rangle, \langle\langle e_0 \rangle\rangle$.
   c. The client and the attestor run an IZK using QuickSilver with witness ($\langle\langle \boldsymbol{r} \rangle\rangle, \langle\langle \mathsf{R}_j \rangle\rangle, \langle\langle \boldsymbol{e} \rangle\rangle$, $\langle\langle e_0 \rangle\rangle$) and the LWE public key $(\boldsymbol{A}, \boldsymbol{y})$ as the ZKP public statement on the following circuit:
      i. Check $0 \leq \langle\langle \mathsf{R}_j \rangle\rangle < p$ and abort in case of failure.
      ii. $\forall i \in [1, n]$ compute $\langle\langle r_i \rangle\rangle \cdot (1 - \langle\langle r_i \rangle\rangle)$, reveal its value, and abort if not zero.
      iii. $\forall i \in [0, n]$ check that $|\langle\langle e_i \rangle\rangle| \leq 6 \cdot \mathsf{stdev}$ and abort in case of failure.
      iv. Compute the ciphertext $(\boldsymbol{a}, b)$ as $(\boldsymbol{A} \cdot \langle\langle \boldsymbol{r} \rangle\rangle + \langle\langle \boldsymbol{e} \rangle\rangle, \boldsymbol{y}^\top \cdot \langle\langle \boldsymbol{r} \rangle\rangle + \langle\langle e_0 \rangle\rangle + \lfloor \langle\langle \mathsf{R}_j \rangle\rangle \cdot \Delta \rfloor)$ mod q, following Eq. (5).
      v. Reveal the ciphertext $(\boldsymbol{a}, b)$ to both client and attestor.
   d. The attestor signs the ciphertext together with its context and position, i.e., $\sigma_j \coloneqq \mathsf{Sign}(sk_{\mathcal{A}}, (\mathsf{ctx}, j, (\boldsymbol{a}, b)))$, and sends σ_j to the client.
   e. The client verifies the signature and sends it along with the ciphertext to node $\mathcal{N}_1$.
   f. The node $\mathcal{N}_1$ verifies $\mathsf{Verify}(pk_{\mathcal{A}}, (\mathsf{ctx}, j, (\boldsymbol{a}, b)), \sigma_j)$ and aborts if it fails.

Figure 6: TLShare protocol with $\mathsf{mode} = \mathsf{fhe}$ with IZK.

These are computed over IT-MAC shares of R through a conversion protocol [XYWY24, Fig. 19]. In our protocol, we avoid this conversion and directly extract the IT-MACs of R (i.e., $\langle\langle \mathsf{R} \rangle\rangle$) from GtP, as outlined in step 1 of Fig. 6. Next, we use $\langle\langle \mathsf{R} \rangle\rangle$ as part of the private witness of a circuit that proves the validity of LWE encryption in ZK.

For each bit of the response message R, we proceed as follows. In step 2.a, the client locally generates any randomness needed for Eq. (5), specifically, $\boldsymbol{r}$, $\boldsymbol{e}$, and e_0. Next, in 2.b, the client and the attestor generate authenticated shares [YSWW21] of the aforementioned

elements $\langle\langle \boldsymbol{r} \rangle\rangle, \langle\langle \boldsymbol{e} \rangle\rangle, \langle\langle e_0 \rangle\rangle$. To prove the validity of encryption $(\boldsymbol{a}, b)$ from Eq. (5), our circuit takes the tuple $(\langle\langle \boldsymbol{r} \rangle\rangle, \langle\langle \mathsf{R}_j \rangle\rangle, \langle\langle \boldsymbol{e} \rangle\rangle, \langle\langle e_0 \rangle\rangle)$ as witness and checks the following (Step 2.*c*):

- $\langle\langle \mathsf{R}_j \rangle\rangle$ lies in the valid plaintext range $[0, p)$.
- Each coordinate of $\langle\langle \boldsymbol{r} \rangle\rangle$ is binary. Notably, the expression $\langle\langle r_i \rangle\rangle \cdot (1 - \langle\langle r_i \rangle\rangle)$ yields 0 if and only if $r_i \in \{0, 1\}$.
- All noise terms e_i satisfy $|\langle\langle e_i \rangle\rangle| \leq 6 \cdot \mathsf{stdev}$ for $i = [0, n]$.
- The ciphertext $(\boldsymbol{a}, b)$ matches Eq. (5).

More formally, we show membership of $(\boldsymbol{a}, b)$ in the relation:

$$\mathcal{R} = \Big\{ \big((\boldsymbol{a}, b), (\boldsymbol{r}, \mathsf{R}_j, \boldsymbol{e}, e_0)\big) \;\Big|\; \mathsf{R}_j \in \mathbb{Z}_p \;\wedge\; \boldsymbol{r} \in \{0,1\}^n \;\wedge\; \forall i \in [0, n],\ |e_i| \leq 6 \cdot \mathsf{stdev} \;\wedge \\ \boldsymbol{a} = \boldsymbol{A}\boldsymbol{r} + \boldsymbol{e} \bmod q \;\wedge\; b = \boldsymbol{y}^\top \boldsymbol{r} + e_0 + \lfloor \mathsf{R}_j \cdot \Delta \rfloor \bmod q \Big\}.$$

Then, after verifying the validity of $(\boldsymbol{a}, b)$, the attestor signs the ciphertext and sends it to the client (step 2.*d*), which subsequently forwards it to the node $\mathcal{N}_1$ (step 2.*e*). Finally, the node verifies the signature, which guarantees that the value under the ciphertext comes from a TLS connection that the attestor verified.

5.2.1 Security of TLShare FHE Mode with IZK Proofs

As outlined in the protocol, after the GtP phase, data is secretly shared under the QuickSilver (QS) protocol. Specifically, the client and the attestor each hold the IT-MAC. The subsequent proof of a valid FHE ciphertext statement is a natural extension of QuickSilver. The security of this phase directly follows from the security guarantees of both the GtP and QS protocols. In other words, TLShare securely realizes the intended ideal functionality if the relation

$$\mathcal{R} = \Big\{ \big((\boldsymbol{a}, \boldsymbol{b}), (\boldsymbol{r}, \mu, \boldsymbol{e}, e_0)\big) \;\Big|\; \mu \in \mathbb{Z}_p \;\wedge\; \boldsymbol{r} \in \{0,1\}^n \;\wedge\; \forall i \in [0, n],\ |e_i| \leq 6 \cdot \mathsf{stdev} \;\wedge \\ \boldsymbol{a} = \boldsymbol{A}\boldsymbol{r} + \boldsymbol{e} \bmod q \;\wedge\; b = \boldsymbol{y}^\top \boldsymbol{r} + e_0 + \lfloor \mu \cdot \Delta \rfloor \bmod q \Big\}.$$

defines a valid ciphertext.

Indeed, the relation characterizes a valid LPR-style variant of Regev's LWE encryption of plaintext μ encrypted under randomness r and noise elements $(\boldsymbol{e}, e_0)$, as it enforces the standard encryption equations ($\boldsymbol{a} = \boldsymbol{A}\boldsymbol{r} + \boldsymbol{e} \bmod q$ and $b = \boldsymbol{y}^\top \boldsymbol{r} + e_0 + \lfloor \mu \cdot \Delta \rfloor \bmod q$) and bounds on errors ($\forall i \in [0, n],\ |e_i| \leq 6 \cdot \mathsf{stdev}$), thus proving encryption validity. This relation is enforced in zero-knowledge by QS through its corresponding circuit. Moreover, GtP uses information-theoretic MACs on the TLS response bits, which are carried forward into the QS protocol. By directly using these MACs inside QS, one guarantees that the plaintext μ encrypted under randomness r and noise elements $(\boldsymbol{e}, e_0)$ is exactly the same message present in the attested TLS connection.

We present the proof of Theorem 4 below.

Theorem 4. *If the LWE assumption holds, the underlying signature scheme is EUF-CMA secure, and the QuickSilver argument is knowledge-sound against malicious provers, then the FHE mode protocol with IZK from Fig. 6 securely realizes the functionality $\mathcal{F}_{TLShare}$ (Fig. 1) in the stand-alone model in the $\mathcal{F}_{AuthData}$-hybrid world.*

Proof. We need to address three separate corruption cases for the client, the attestor, and the FHE node $\mathcal{N}_1$. In each case, we construct a PPT simulator given access to functionality $\mathcal{F}_{\mathsf{TLShare}}$ which runs a PPT adversary as a subroutine, and emulates functionality $\mathcal{F}_{\mathsf{AuthData}}$.

Case 1: Malicious client. Let $\mathcal{A}$ control the client. We consider the following sequence of hybrids.

- H_0: This is the real execution of the FHE mode with IZK protocol.
- H_1: Replace Π_{AuthData} by $\mathcal{F}_{\mathsf{AuthData}}$. The simulator obtains the authenticated response R and the corresponding IT-MAC state. The hybrids H_0 and H_1 are indistinguishable by the security of Π_{AuthData}.
- H_2: Proceed as in H_1, but abort if $\mathcal{N}_1$ accepts a signed message $M'_j \neq M_j$, where $M_j = (\mathsf{ctx}, j, \mathsf{ct}_j)$ is the message signed by the honest attestor after a successful QuickSilver execution. A reduction $\mathcal{B}$, given access to the EUF-CMA signing oracle, obtains signatures on the messages M_j and runs $\mathcal{A}$. Any accepted signature on a different message is then a forgery. Hence, H_1 and H_2 are indistinguishable under EUF-CMA security.
- H_3: Replace the ciphertext delivery by the corrupt-client branch of $\mathcal{F}_{\mathsf{TLShare}}$. The simulator acts as the honest attestor in QuickSilver for the actual public ciphertext $\mathsf{ct}_j = (\boldsymbol{a}, b)$. If QuickSilver rejects, it causes the functionality to abort and does not sign. Otherwise, knowledge soundness against malicious QuickSilver provers yields a witness $\omega_j = (\boldsymbol{r}, \mu, \boldsymbol{e}, e_0)$ satisfying the encryption circuit for the same ciphertext. Since the circuit uses the authenticated sharing $\langle\!\langle \mathsf{R}_j \rangle\!\rangle$ obtained from $\mathcal{F}_{\mathsf{AuthData}}$, IT-MAC authentication implies $\mu = \mathsf{R}_j$, except with negligible probability. Thus, $\mathcal{R}_{\mathsf{Enc}}(\mathsf{pk}_{\mathsf{fhe}}, \mathsf{R}_j; \mathsf{ct}_j, \omega_j) = 1$. The simulator submits the exact pairs $(\mathsf{ct}_j, \omega_j)$ to the functionality and signs M_j only after the relation accepts. The functionality forwards the same ciphertexts to $\mathcal{N}_1$.

Thus, H_3 is the ideal execution, and the real and ideal executions are indistinguishable under the security of Π_{AuthData}, EUF-CMA security, IT-MAC authentication, and QuickSilver knowledge soundness.

Case 2: Malicious attestor. In this case the security reduces to the security of Π_{AuthData} and the zero-knowledge property of QuickSilver. Consider the following sequence of hybrids:

- H_0: This is the real execution of the FHE mode with IZK protocol in Fig. 6 with adversary controlling the attestor.
- H_1: We replace the execution of Π_{AuthData} with the ideal functionality $\mathcal{F}_{\mathsf{AuthData}}$. The simulator receives IT-MAC shares $\langle\!\langle Q \rangle\!\rangle$ and $\langle\!\langle R \rangle\!\rangle$ from $\mathcal{F}_{\mathsf{AuthData}}$.

Indistinguishability between H_0 and H_1 comes from security of Π_{AuthData}.

- H_2: The simulator replaces the real QuickSilver proof with a simulated proof using its zero-knowledge simulator.

Indistinguishability between H_1 and H_2 comes from the zero-knowledge property of QuickSilver IZK.

- H_3: The simulator replaces the real ciphertext $(\boldsymbol{a}, b)$ with an encryption of a dummy value $R'_j = 0$ under the same public key.

Indistinguishability between H_2 and H_3 comes from the IND-CPA security property of the underlying LWE encryption, which follows from the LWE assumption. Indeed, the malicious attestor cannot distinguish between encryptions of different plaintexts due to the IND-CPA security property.

The simulator finishes by receiving the signed ciphertext from the attestor.

Case 3: Malicious FHE node $\mathcal{N}_1$. Consider the following sequence of hybrids:

- H_0: This is the real execution of the FHE mode with IZK protocol in Fig. 6 with adversary controlling $\mathcal{N}_1$.
- H_1: The simulator computes a fake ciphertext $(\boldsymbol{a'}, b')$ for a fake $R' = 0$ by running the encryption algorithm with dummy values. Then, the simulator plays the role of the challenger in the EUF-CMA game and receives the signature σ corresponding to the ciphertext $(\boldsymbol{a'}, b')$. Sends both $(\boldsymbol{a'}, b')$ and σ to the FHE node.

Indistinguishability between H_0 and H_1 comes from the EUF-CMA property and the IND-CPA security property of LWE encryption. The malicious FHE node cannot distinguish between encryptions of different plaintexts under the LWE assumption. □

6 Applications

6.1 Financial Risk Assessments

We consider a use case in which an individual client wishes to utilize a third-party risk assessment service to evaluate their financial circumstances by jointly analyzing data housed within multiple institutions, such as banks, employers, and tax authorities. One major implementation challenge associated with secure computation workflows for assessing financial risk [AKL12] is determining which parties are responsible for coordinating interactions between the multiple institutions contributing data and the risk assessment service provider. The institutions may find it prohibitively expensive (or in conflict with the fact that they are competitors) to agree in advance to collaborate on such workflows. Even if these obstacles are not present, it is impractical for each participating institution to integrate with the centralized risk assessment service provider.

As illustrated in Fig. 2, our approach addresses this obstacle by enabling the client to perform this coordination independently, utilizing existing TLS servers from various institutions. With the help of an attestor, the downstream computation can be assured of the accuracy of the client-provided data. The analysis is then performed under MPC/FHE by combining the different encrypted inputs, comparing them to thresholds, and even evaluating them with machine learning models [GT00, PCR19].

6.2 Wage Equity Reporting

We consider the Boston Women's Workforce Council (BWWC) study that measures wage disparities across demographic groups (including race, gender, and level of seniority) within a metropolitan region [LJA+18, Man21]. This study is accomplished by executing an aggregation workflow on input data (provided directly by a large number of individual employer organizations) that contains wage information broken down by demographic attributes; the aggregated output is then disclosed in public reports [Cou17]. Traditional approaches, such as centralized collection of the plaintext input data by a trusted third party, are not acceptable to data contributors due to the inherent risks (to both employees and employers) of data exposure. Employing MPC to minimize data exposure risks is viable [LJA+18] but involves manual data entry that may introduce errors without careful interface design, costly user training [QLJ+19], and does not inherently provide mechanisms for verifying the accuracy of the data. These drawbacks limit the scale of such studies and can invite skepticism about data accuracy.

Our approach can mitigate these drawbacks by leveraging multiple clients, each independently collecting authenticated data in an automated and scalable manner from authoritative TLS endpoints (e.g., payroll APIs, employer HR systems, municipal databases, and so on) by working with an attestor. Finally, the sensitive data (e.g., reported compensation amounts, employment statuses, and demographic attributes) is verifiably imported

into a downstream secure computation. The downstream analysis can compute statistics (e.g., histograms of compensation ranges broken down by demographic attributes or wage ratios between demographic groups) [LJA+18], perform regressions, or even train federated models [FMM+21, GMS+23], all without revealing any party's individual records.

7 Experimental Evaluation

We implemented TLShare in C++ and Rust, choosing libraries tailored to each protocol's cryptographic requirements.[3] We modified the conversion protocol of GtP [XYWY24] and the RELIC [RT25] library to integrate both the Ed25519 curve and the pairing-friendly BLS12-446 curve, required for both the MPC mode and FHE mode with NIZK, respectively. We extended the TFHE-rs library [Zam25] to include our vector membership NIZK proofs. For the FHE mode with IZK, we relied on the QuickSilver protocol [YSWW21] and the EMP toolkit [WMK16] to implement the circuit for proving the correctness of Regev encryption. Lastly, our generated ciphertexts are compatible with the Primus FHE library [Pri25] for the downstream FHE computation, while for the downstream MPC, we used MP-SPDZ [Kel20].

TLShare supports any TLS version for which an underlying protocol securely realizes $\mathcal{F}_{\mathsf{AuthData}}$. Our implementation instantiates Π_{AuthData} with the GtP's TLS 1.2 protocol [XYWY24]. GtP also describes an extension to TLS 1.3, but does not implement or evaluate it, nor do we. Thus, the measurements reported here apply only to TLS 1.2, although TLShare can support TLS 1.3 when it is instantiated with GtP's extension or another suitable realization of $\mathcal{F}_{\mathsf{AuthData}}$.

The experiments involve the following set of parties: the client, the attestor, and secure computation nodes. The client and the attestor are placed in two different AWS `c5.xlarge` instances, each with 4 vCPUs and 8 GB of RAM. For the secure computation nodes, we used a `c5.9xlarge` instance with a total of 36 vCPUs and 72 GBs of RAM, and dedicated 4 vCPUs per node. We ran experiments over two settings: a LAN with 1 Gbps bandwidth and near-zero latency, and a WAN with 100 Mbps bandwidth and 50 ms latency.

Next, we present experimental evaluations for each mode. For all modes, the initial steps of the TLS phase take roughly 0.46 seconds under LAN and 3.19 seconds under WAN for a payload of 256 bytes (Section 7.4). From this point onward, the modes diverge and we detail the additional computation and communication overhead introduced by each in Sections 7.1 to 7.3. Finally, in Section 7.4, we combine the TLS latency with these per-mode overheads to derive the total end-to-end time for each mode (cf. Table 5).

Table 2: Latency comparisons (in ms) between TLShare in MPC mode and Garble-then-Prove over LAN and WAN for different payload sizes (bytes). We focus on the conversion protocols of the TLS Phase for both works.

Protocol	**Network**	**LAN**				**WAN**			
	Payload (B)	2^8	2^9	2^{10}	2^{11}	2^8	2^9	2^{10}	2^{11}
GtP [XYWY24]	**Client**	10.1	18.5	35.2	69.7	317.2	330.9	357.3	414.9
	Attestor	9.41	17.8	34.1	67.9	215.9	229.4	255.3	312.3
TLShare	**Client**	12.0	20.5	37.5	71.6	318.5	332.8	358.8	413.7
	Attestor	11.5	19.9	36.7	70.2	317.9	331.7	357.8	412.8

[3] https://github.com/manel1874/tlshare

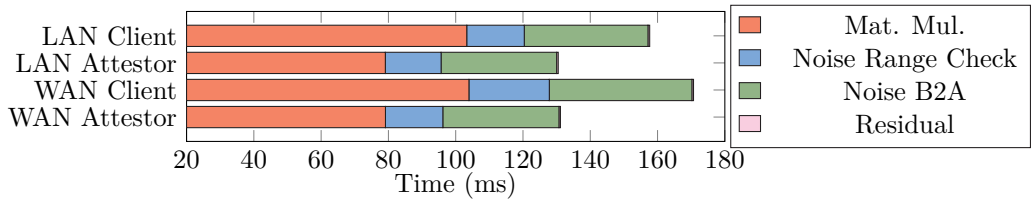

Figure 7: Latency breakdown of total time (ms) per party (client and attestor) in LAN and WAN for the Encryption phase of the FHE with IZK. "Residual" includes a check verifying that the randomness of the encryption is a binary vector, an input range check asserting that the message μ is within the plaintext modulus, and the input conversion from binary to arithmetic (B2A).

7.1 MPC Mode

In the MPC mode, the original GtP implementation uses OpenSSL's NIST P-256 curve for Pedersen commitments and lacks support for signing those commitments. We reimplemented the conversion protocol Π_{Conv} [XYWY24, Fig. 20] on the Ed25519 curve, using the RELIC library. Ed25519 provides comparable 128-bit security to P-256 while offering significantly faster operations and smaller point sizes [Ber06]. We further integrated ECDSA signatures directly into the conversion process, enabling authenticated commitments.

In Table 2, we compare the latency in milliseconds of the TLS phase of GtP and TLShare in the MPC mode for different payload sizes, ranging from 256 to 2048 bytes. We batch the IT-MACs of the payload bits, resulting in 9 and 66 batches for the 256 and 2048 byte payloads, respectively. Observe that under LAN, for both frameworks, both the client and the attestor take almost the same time. Under WAN, however, the client in GtP takes about 100 ms more than the attestor, while in TLShare, their runtime is the same. The clients of both GtP and TLShare are also equivalent. This is because the additional cost of the attestor in the TLShare over GtP (i.e., around 100ms under WAN) is introduced by signing the commitments. However, this cost is negligible when viewed in the context of the full TLS exchange, as it unlocks the ability to perform secure computation on authenticated data.

Table 3: Latency evaluations (in ms) for the Share phase of TLShare in the MPC mode under LAN and WAN for different numbers of parties n, thresholds T, and one batch.

Party	**Network**	**LAN**			**WAN**		
	(T, n)	$(2,3)$	$(3,5)$	$(4,7)$	$(2,3)$	$(3,5)$	$(4,7)$
Client	**Prep.**	0.182	0.271	0.337	0.187	0.254	0.336
	Sec. Sh.	0.00052	0.00061	0.00079	0.00052	0.00049	0.0005
MPC Nodes	**Sig. Ver.**	0.432	0.442	0.494	0.538	1.012	0.880
	Sh. Ver.	0.212	0.248	0.325	0.215	0.267	0.451
All	**Comput.**	0.826	0.961	1.155	0.940	1.533	1.667
	Total	1.245	1.501	1.730	101.62	102.57	102.47

During the Share phase, we use Pedersen verifiable secret sharing (Pedersen VSS) to distribute one secret at a time across different configurations of Shamir shares [CDM00, AFL+16]. Importantly, this phase excludes the attestor: once the client has the signed commitments, it executes a single interaction with the MPC nodes. The client first generates Pedersen VSS parameters, then computes the shares for each MPC node. On the MPC nodes' side, each party verifies the attestor's signature on the base Pedersen commitments and checks its consistency.

We present a breakdown of these costs in Table 3, including preprocessing (Pedersen VSS parameter generation), share generation time, and node-side verification for different settings of threshold T and number of nodes n. The last row of Table 3 presents the end-to-end time, including network latency, while the row above it shows the local computation timings combined (i.e., without the network latency). The latter leads to a share phase faster than 2 ms and a bandwidth of 0.2 KB across the board. While the running time under LAN is dominated by local computation, the WAN results are primarily constrained by the network latency. Finally, the computation timing is linearly dependent on the number of commitments generated and verified, which is directly related to the threshold value T and not the number of nodes n.

Table 4: Latency evaluations (in ms) of the FHE mode of TLShare with NIZK over LAN and WAN between the client ($\mathcal{C}$), the attestor ($\mathcal{A}$), and the FHE node ($\mathcal{N}_1$) for a single batch.

Phase	**Network**	**LAN**			**WAN**		
	Party	$\mathcal{C}$	$\mathcal{A}$	$\mathcal{N}_1$	$\mathcal{C}$	$\mathcal{A}$	$\mathcal{N}_1$
TLS	**Conversion**	5.06	5.23	-	307.3	307.9	-
Enc.	**CGGI Proof**	609	-	137	607	-	98.0
	VM Proof	18	-	13	18	-	12
	Sig. Verification	-	-	0.56	-	-	0.56
	Total	800	-	151	1103	-	157

7.2 FHE Mode with NIZK

The vector commitments in TFHE-rs's NIZK of correct encryption [BCD+25, Section 7.6.3] are over the second group $\hat{\mathbb{G}}$ of the BLS12-446 curve. For our vector membership proof, we reimplemented the GtP conversion protocol over $\hat{\mathbb{G}}$ to link the Pedersen commitments of the payload with these vector commitments using RELIC.

In Table 4, we report the timings for converting and encrypting a batch under the FHE mode with NIZK. Note that the NIZK-based protocol only interacts with the attestor during the TLS phase and with the FHE node during the encryption phase. In the latter, the bulk of the computation occurs within the client-side generation of the encryption correctness proof (CGGI.Prove), which dominates the overall cost, while the NIZK proof of vector membership (VM.Prove) adds minimal additional time. Indeed, the NIZK VM proof is 33× faster than the NIZK proof of valid encryption for the client (i.e., the prover) and over 8× faster for $\mathcal{N}_1$ (i.e., the verifier). Total bandwidth transferred between the client and the node is 92.4 KB and the last row of Table 4 shows the end-to-end time for the encryption phase between all parties, including network latency.

7.3 FHE Mode with IZK

Lastly, the FHE mode with IZK uses the Primus FHE parameter set [Pri25]: dimension $n = 2048$, ciphertext modulus $q = 2^{61} - 1$, message modulus $p = 4$, Gaussian noise with $\mathsf{stdev} = 3.19 \times 2^7$, and a ternary secret key. The resulting circuit comprises approximately 148 million `AND` gates for binary checks (including message and noise range validations) and an arithmetic part that performs scalar matrix–vector multiplication and randomness checks. These checks require verifying, for each random component $\langle\langle r_i \rangle\rangle$ that $\langle\langle r_i \rangle\rangle \cdot (1 - \langle\langle r_i \rangle\rangle) = 0$ to confirm $\langle\langle r_i \rangle\rangle$ is binary.

The preprocessing in QuickSilver needed for our encryption circuit requires approximately 131 thousand edaBits [EGK+20]. We measured the performance of a single batch

of 1.6 million edaBits, which takes about 40 seconds over a LAN and 118 seconds over a WAN, transferring approximately 122 MB of data. This step fully accounts for the offline work required by the IZK approach for the encryption phase. The performance of the online part of the encryption phase (see Fig. 7) is dominated by the matrix-vector multiplication (orange bar), noise range checks (blue), and conversions from binary to arithmetic form (green). This multiplication is between the public key (a matrix) and the secret randomness binary vector $\boldsymbol{r}$. As a result, it does not require any communication, as experimentally verified in Fig. 7 since the LAN and WAN orange bars are the same. The total communication is around 204 KB, resulting in a slight discrepancy between the LAN and WAN end-to-end timings.

7.4 Discussion & Applications

The three TLShare protocols presented above provide different trade-offs across performance, interaction, and trust assumptions. Notably, all protocols begin by generating IT-MACs of the TLS response, taking roughly 0.46 seconds under LAN and 3.19 seconds under WAN for a payload of 256 bytes. Next, during the TLS phase, the client and the attestor in the MPC and FHE with NIZK protocols, run our conversion protocol based on Pedersen commitments with the Ed25519 and BLS12-446 curves, respectively. In contrast, the IZK protocol bypasses this conversion with an interactive ZKP performed over IT-MACs with the attestor directly.

Table 5: TLShare total end-to-end time (sec) containing both the TLS phase and the share/encryption phase under LAN and WAN for a payload of 2^8 bytes.

Network	**MPC Mode**	**FHE with NIZK**	**FHE with IZK**
LAN	0.477	1.127	0.618
WAN	3.615	4.601	3.361

An important difference between the protocols is in the bandwidth and the elliptic curve group representation. Our MPC mode uses Ed25519 curve elements at 256 bits during the conversion step in the TLS phase, while the FHE with NIZK operates with elements in the second source group $\hat{\mathbb{G}}$ of the pairing-friendly BLS12-446 of size 892 bits. This results in roughly a 4× increase in bandwidth during the conversion protocol. This cost is also reflected empirically: in the WAN setting, the conversion of a single batch under FHE with NIZK (307 ms from Table 4) takes around as much as a conversion of a payload of size 2^8 (i.e., 9 batches) under MPC (318 ms from Table 2). In contrast, the FHE with IZK avoids the TLS conversion cost altogether, yet still requires interaction during encryption proof generation between the client and the attestor.

When comparing FHE with NIZK and FHE with IZK in terms of end-to-end time, IZK outperforms NIZK as shown in Table 5. More specifically, under LAN, IZK takes 0.6 seconds compared to 1.1 seconds, while in WAN, IZK takes 3.4 seconds as compared to the 4.6 seconds. This is due to the expensive CGGI NIZK. However, this advantage comes at the cost of preprocessing for approximately 131 thousand edaBits, taking around 3.3 and 9.6 seconds in LAN and WAN. Furthermore, the FHE mode with IZK is limited to prime-field LWE schemes and lacks support for a ring structure, not making it compatible with TFHE-rs [Zam25]. Conversely, the FHE mode with NIZK relies on a trusted setup, requiring stronger trust assumptions.

Finally, to present the capabilities of TLShare, we evaluated the application presented in Section 6.2 using malicious Shamir secret sharing [CDM00, AFL+16]. We grouped records over five industries (e.g., healthcare, finance, etc.) and five job categories (e.g., executive, administrative, etc.), and computed the wage gaps between genders for different demographics. We evaluated this application for 10^3 up to 10^6 records under LAN and

WAN, with the LAN experiments taking 0.01, 0.1, 1.3, 13.1 seconds, respectively, and WAN experiments requiring 0.6, 2.07, 15.4, 146.7 seconds, respectively. The total communication scaled linearly from 0.4 MB for 1000 records, all the way up to approximately 400 MB for 1 million records.

8 Concluding Remarks

We presented TLShare, a framework that bridges the gap between authenticated web data and privacy-preserving computation. TLShare enables clients to verifiably extract TLS payloads and input them into secure downstream MPC or FHE computations, without requiring any server-side modifications. Unlike prior work restricted to Boolean predicates and single client proofs, TLShare supports arbitrary secure computation over verifiable payloads from multiple clients and servers, enabling new applications such as privacy-preserving financial risk assessment and collaborative analytics. We introduced three protocols for TLShare: one for MPC using verifiable secret sharing, and two for FHE using interactive and non-interactive zero-knowledge proofs, each ensuring input authenticity, integrity, and end-to-end privacy.

References

[ADST21] Damiano Abram, Ivan Damgård, Peter Scholl, and Sven Trieflinger. Oblivious TLS via multi-party computation. In Kenneth G. Paterson, editor, *Topics in Cryptology – CT-RSA 2021*, volume 12704 of *Lecture Notes in Computer Science*, pages 51–74. Springer, Cham, May 2021. doi:10.1007/978-3-030-75539-3_3.

[AFL+16] Toshinori Araki, Jun Furukawa, Yehuda Lindell, Ariel Nof, and Kazuma Ohara. High-throughput semi-honest secure three-party computation with an honest majority. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016: 23rd Conference on Computer and Communications Security*, pages 805–817. ACM Press, October 2016. doi:10.1145/2976749.2978331.

[AKL12] Emmanuel A. Abbe, Amir E. Khandani, and Andrew W. Lo. Privacy-preserving methods for sharing financial risk exposures. *American Economic Review*, 102(3):65–70, May 2012. doi:10.1257/aer.102.3.65.

[B+21] Lorenz Breidenbach et al. Chainlink 2.0: Next steps in the evolution of decentralized oracle networks. *Chainlink Labs*, 1:1–136, 2021. URL: https://research.chain.link/whitepaper-v2.pdf.

[BBC+19] Dan Boneh, Elette Boyle, Henry Corrigan-Gibbs, Niv Gilboa, and Yuval Ishai. Zero-knowledge proofs on secret-shared data via fully linear PCPs. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part III*, volume 11694 of *Lecture Notes in Computer Science*, pages 67–97. Springer, Cham, August 2019. doi:10.1007/978-3-030-26954-8_3.

[BBC+21] Dan Boneh, Elette Boyle, Henry Corrigan-Gibbs, Niv Gilboa, and Yuval Ishai. Lightweight techniques for private heavy hitters. In *2021 IEEE Symposium on Security and Privacy*, pages 762–776. IEEE Computer Society Press, May 2021. doi:10.1109/SP40001.2021.00048.

[BBHR19] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Scalable zero knowledge with no trusted setup. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part III*, volume 11694 of *Lecture Notes in Computer Science*, pages 701–732. Springer, Cham, August 2019. doi:10.1007/978-3-030-26954-8_23.

[BCCT13] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. Recursive composition and bootstrapping for SNARKS and proof-carrying data. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th Annual ACM Symposium on Theory of Computing*, pages 111–120. ACM Press, June 2013. doi:10.1145/2488608.2488623.

[BCD+25] Carl Bootland, Kelong Cong, Daniel Demmler, Tore Kasper Frederiksen, Benoit Libert, Jean-Baptiste Orfila, Dragos Rotaru, Nigel P. Smart, Titouan Tanguy, Samuel Tap, and Michael Walter. Threshold (fully) homomorphic encryption. "Cryptology ePrint Archive, Report 2025/699", 2025. URL: https://eprint.iacr.org/2025/699.

[BCG+13] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, Eran Tromer, and Madars Virza. SNARKs for C: Verifying program executions succinctly and in zero knowledge. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology – CRYPTO 2013, Part II*, volume 8043 of *Lecture Notes in Computer Science*, pages 90–108. Springer, Berlin, Heidelberg, August 2013. doi:10.1007/978-3-642-40084-1_6.

[BCI+22] Nir Bitansky, Alessandro Chiesa, Yuval Ishai, Rafail Ostrovsky, and Omer Paneth. Succinct non-interactive arguments via linear interactive proofs. *Journal of Cryptology*, 35(3):15, July 2022. doi:10.1007/s00145-022-09424-4.

[BCJ+24] Foteini Baldimtsi, Konstantinos Kryptos Chalkias, Yan Ji, Jonas Lindstrøm, Deepak Maram, Ben Riva, Arnab Roy, Mahdi Sedaghat, and Joy Wang. zkLogin: Privacy-preserving blockchain authentication with existing credentials. In Bo Luo, Xiaojing Liao, Jun Xu, Engin Kirda, and David Lie, editors, *ACM CCS 2024: 31st Conference on Computer and Communications Security*, pages 3182–3196. ACM Press, October 2024. doi:10.1145/3658644.3690356.

[BDM+21] Marcus Brinkmann, Christian Dresen, Robert Merget, Damian Poddebniak, Jens Müller, Juraj Somorovsky, Jörg Schwenk, and Sebastian Schinzel. ALPACA: Application layer protocol confusion - analyzing and mitigating cracks in TLS authentication. In Michael Bailey and Rachel Greenstadt, editors, *USENIX Security 2021: 30th USENIX Security Symposium*, pages 4293–4310. USENIX Association, August 2021. URL: https://www.usenix.org/conference/usenixsecurity21/presentation/brinkmann.

[BDOZ11] Rikke Bendlin, Ivan Damgård, Claudio Orlandi, and Sarah Zakarias. Semi-homomorphic encryption and multiparty computation. In Kenneth G. Paterson, editor, *Advances in Cryptology – EUROCRYPT 2011*, volume 6632 of *Lecture Notes in Computer Science*, pages 169–188. Springer, Berlin, Heidelberg, May 2011. doi:10.1007/978-3-642-20465-4_11.

[Ber06] Daniel J. Bernstein. Curve25519: New Diffie-Hellman speed records. In Moti Yung, Yevgeniy Dodis, Aggelos Kiayias, and Tal Malkin, editors, *PKC 2006: 9th International Conference on Theory and Practice of Public Key Cryptography*, volume 3958 of *Lecture Notes in Computer Science*, pages 207–228. Springer, Berlin, Heidelberg, April 2006. doi:10.1007/11745853_14.

[BMRS21] Carsten Baum, Alex J. Malozemoff, Marc B. Rosen, and Peter Scholl. Mac'n'cheese: Zero-knowledge proofs for boolean and arithmetic circuits with nested disjunctions. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology – CRYPTO 2021, Part IV*, volume 12828 of *Lecture Notes in Computer Science*, pages 92–122, Virtual Event, August 2021. Springer, Cham. doi:10.1007/978-3-030-84259-8_4.

[Can00] Ran Canetti. Security and composition of multiparty cryptographic protocols. *Journal of Cryptology*, 13(1):143–202, January 2000. doi:10.1007/s00145 9910006.

[CCX+19] Guoxing Chen, Sanchuan Chen, Yuan Xiao, Yinqian Zhang, Zhiqiang Lin, and Ten-Hwang Lai. SgxPectre: Stealing intel secrets from SGX enclaves via speculative execution. In *2019 IEEE European Symposium on Security and Privacy*, pages 142–157. IEEE Computer Society Press, June 2019. doi: 10.1109/EuroSP.2019.00020.

[CCY23] Kwan Yin Chan, Handong Cui, and Tsz Hon Yuen. DIDO: Data Provenance from Restricted TLS 1.3 Websites. In Weizhi Meng, Zheng Yan, and Vincenzo Piuri, editors, *International Conference on Information Security Practice and Experience*, volume 14341 of *Lecture Notes in Computer Science*, pages 154–169, Singapore, 2023. Springer. doi:10.1007/978-981-99-7032-2_10.

[CDH+25] Sofía Celi, Alex Davidson, Hamed Haddadi, Gonçalo Pestana, and Joe Rowell. DiStefano: Decentralized infrastructure for sharing trusted encrypted facts and nothing more. In *32nd Annual Network and Distributed System Security Symposium (NDSS 2025)*, pages 1–18, San Diego, CA, USA, February 2025. The Internet Society. doi:10.14722/ndss.2025.241367.

[CDM00] Ronald Cramer, Ivan Damgård, and Ueli M. Maurer. General secure multi-party computation from any linear secret-sharing scheme. In Bart Preneel, editor, *Advances in Cryptology – EUROCRYPT 2000*, volume 1807 of *Lecture Notes in Computer Science*, pages 316–334. Springer, Berlin, Heidelberg, May 2000. doi:10.1007/3-540-45539-6_22.

[CF13] Dario Catalano and Dario Fiore. Vector commitments and their applications. In Kaoru Kurosawa and Goichiro Hanaoka, editors, *PKC 2013: 16th International Conference on Theory and Practice of Public Key Cryptography*, volume 7778 of *Lecture Notes in Computer Science*, pages 55–72. Springer, Berlin, Heidelberg, February / March 2013. doi:10.1007/978-3-642-36362-7_5.

[CGB17] Henry Corrigan-Gibbs and Dan Boneh. Prio: Private, robust, and scalable computation of aggregate statistics. In Aditya Akella and Jon Howell, editors, *Proceedings of the 14th USENIX Conference on Networked Systems Design and Implementation*, NSDI'17, pages 259–282, USA, 2017. USENIX Association. URL: https://www.usenix.org/conference/nsdi17/technical-sessions/presentation/corrigan-gibbs.

[CGGI20] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. TFHE: Fast fully homomorphic encryption over the torus. *Journal of Cryptology*, 33(1):34–91, January 2020. doi:10.1007/s00145-019-09319-x.

[CLR24] Sofía Celi, Shai Levin, and Joe Rowell. CDLS: Proving knowledge of committed discrete logarithms with soundness. In Serge Vaudenay and Christophe Petit, editors, *AFRICACRYPT 24: 15th International Conference on Cryptology in Africa*, volume 14861 of *Lecture Notes in Computer Science*, pages 69–93. Springer, Cham, July 2024. doi:10.1007/978-3-031-64381-1_4.

[Cou17] Boston Women's Workforce Council. Boston women's workforce council 2017 wage gap report. Technical report, City of Boston, 2017.

[DGPS24] Moumita Dutta, Chaya Ganesh, Sikhar Patranabis, and Nitin Singh. Compute, but verify: Efficient multiparty computation over authenticated inputs. In Kai-Min Chung and Yu Sasaki, editors, *Advances in Cryptology – ASIACRYPT 2024, Part VI*, volume 15489 of *Lecture Notes in Computer Science*, pages 133–166. Springer, Singapore, December 2024. doi:10.1007/978-981-96-0938-3_5.

[DILO22] Samuel Dittmer, Yuval Ishai, Steve Lu, and Rafail Ostrovsky. Authenticated garbling from simple correlations. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022, Part IV*, volume 13510 of *Lecture Notes in Computer Science*, pages 57–87. Springer, Cham, August 2022. doi:10.1007/978-3-031-15985-5_3.

[DM15] Léo Ducas and Daniele Micciancio. FHEW: Bootstrapping homomorphic encryption in less than a second. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology – EUROCRYPT 2015, Part I*, volume 9056 of *Lecture Notes in Computer Science*, pages 617–640. Springer, Berlin, Heidelberg, April 2015. doi:10.1007/978-3-662-46800-5_24.

[DPRS23] Hannah Davis, Christopher Patton, Mike Rosulek, and Phillipp Schoppmann. Verifiable distributed aggregation functions. *Proceedings on Privacy Enhancing Technologies*, 2023(4):578–592, October 2023. doi:10.56553/popets-2023-0126.

[DPSZ12] Ivan Damgård, Valerio Pastro, Nigel P. Smart, and Sarah Zakarias. Multiparty computation from somewhat homomorphic encryption. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology – CRYPTO 2012*, volume 7417 of *Lecture Notes in Computer Science*, pages 643–662. Springer, Berlin, Heidelberg, August 2012. doi:10.1007/978-3-642-32009-5_38.

[DR08] T. Dierks and E. Rescorla. The Transport Layer Security (TLS) Protocol Version 1.2. RFC 5246, August 2008. Obsoleted by RFC 8446. URL: https://www.rfc-editor.org/rfc/rfc5246.

[EFM+19] Úlfar Erlingsson, Vitaly Feldman, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Abhradeep Thakurta. Amplification by shuffling: From local to central differential privacy via anonymity. In Timothy M. Chan, editor, *30th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2468–2479. ACM-SIAM, January 2019. doi:10.1137/1.9781611975482.151.

[EGK+20] Daniel Escudero, Satrajit Ghosh, Marcel Keller, Rahul Rachuri, and Peter Scholl. Improved primitives for MPC over mixed arithmetic-binary circuits. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology – CRYPTO 2020, Part II*, volume 12171 of *Lecture Notes in Computer Science*, pages 823–852. Springer, Cham, August 2020. doi:10.1007/978-3-030-56880-1_29.

[ELW+25] Jens Ernstberger, Jan Lauinger, Yinnan Wu, Arthur Gervais, and Sebastian Steinhorst. ORIGO: Proving Provenance of Sensitive Data with Constant Communication. *Proceedings on Privacy Enhancing Technologies*, 2025(2):416–433, July 2025. doi:10.56553/popets-2025-0069.

[FKL18] Georg Fuchsbauer, Eike Kiltz, and Julian Loss. The algebraic group model and its applications. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018, Part II*, volume 10992 of *Lecture Notes in Computer Science*, pages 33–62. Springer, Cham, August 2018. doi:10.1007/978-3-319-96881-0_2.

[FMM+21] Hossein Fereidooni, Samuel Marchal, Markus Miettinen, Azalia Mirhoseini, Helen Möllering, Thien Duc Nguyen, Phillip Rieger, Ahmad-Reza Sadeghi, Thomas Schneider, Hossein Yalame, and Shaza Zeitouni. SAFELearn: Secure Aggregation for private FEderated Learning. In *2021 IEEE Security and Privacy Workshops (SPW)*, pages 56–62, 2021. doi:10.1109/SPW53761.2021.00017.

[FMRV22] Thibauld Feneuil, Jules Maire, Matthieu Rivain, and Damien Vergnaud. Zero-knowledge protocols for the subset sum problem from MPC-in-the-head with rejection. In Shweta Agrawal and Dongdai Lin, editors, *Advances in Cryptology – ASIACRYPT 2022, Part II*, volume 13792 of *Lecture Notes in Computer Science*, pages 371–402. Springer, Cham, December 2022. doi:10.1007/978-3-031-22966-4_13.

[Gen09] Craig Gentry. *A fully homomorphic encryption scheme*. PhD thesis, Stanford University, 2009.

[GGPR13] Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic span programs and succinct NIZKs without PCPs. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology – EUROCRYPT 2013*, volume 7881 of *Lecture Notes in Computer Science*, pages 626–645. Springer, Berlin, Heidelberg, May 2013. doi:10.1007/978-3-642-38348-9_37.

[GMS+23] Till Gehlhar, Felix Marx, Thomas Schneider, Ajith Suresh, Tobias Wehrle, and Hossein Yalame. SafeFL: MPC-friendly Framework for Private and Robust Federated Learning. In *2023 IEEE Security and Privacy Workshops (SPW)*, pages 69–76, Los Alamitos, CA, USA, May 2023. IEEE Computer Society. doi:10.1109/SPW59333.2023.00012.

[GMT23] Charles Gouert, Dimitris Mouris, and Nektarios Georgios Tsoutsos. SoK: New insights into fully homomorphic encryption libraries via standardized benchmarks. *Proceedings on Privacy Enhancing Technologies*, 2023(3):154–172, July 2023. doi:10.56553/popets-2023-0075.

[GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In Alfred Aho, editor, *19th Annual ACM Symposium on Theory of Computing*, pages 218–229. ACM Press, May 1987. doi:10.1145/28395.28420.

[Goo20] Google/Apple. Exposure notification, cryptography specification. https://covid19-static.cdn-apple.com/applications/covid19/current/static/contact-tracing/pdf/ExposureNotification-CryptographySpecificationv1.2.pdf, 2020.

[GT00] J. Galindo and P. Tamayo. Credit Risk Assessment Using Statistical and Machine Learning: Basic Methodology and Risk Modeling Applications. *Computational Economics*, 15(1):107–143, Apr 2000. doi:10.1023/A:1008699112516.

[GUM+24] Charles Gouert, Mehmet Ugurbil, Dimitris Mouris, Miguel de Vega, and Nektarios Georgios Tsoutsos. Ripple: Accelerating programmable bootstraps for FHE with wavelet approximations. In Nicky Mouha and Nick Nikiforakis, editors, *ISC 2024: 27th International Conference on Information Security, Part I*, volume 15257 of *Lecture Notes in Computer Science*, pages 273–293. Springer, Cham, October 2024. doi:10.1007/978-3-031-75757-0_14.

[JLP03] Kose John, Anthony W Lynch, and Manju Puri. Credit ratings, collateral, and loan characteristics: Implications for yield. *The Journal of Business*, 76(3):371–409, 2003. doi:10.1086/375252.

[Joy21] Marc Joye. Balanced non-adjacent forms. In Mehdi Tibouchi and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2021, Part III*, volume 13092 of *Lecture Notes in Computer Science*, pages 553–576. Springer, Cham, December 2021. doi:10.1007/978-3-030-92078-4_19.

[Joy22] Marc Joye. SoK: Fully homomorphic encryption over the [discretized] torus. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2022(4):661–692, 2022. doi:10.46586/tches.v2022.i4.661-692.

[Joy24] Marc Joye. TFHE public-key encryption revisited. In Elisabeth Oswald, editor, *Topics in Cryptology – CT-RSA 2024*, volume 14643 of *Lecture Notes in Computer Science*, pages 277–291. Springer, Cham, May 2024. doi:10.1007/978-3-031-58868-6_11.

[Kel20] Marcel Keller. MP-SPDZ: A versatile framework for multi-party computation. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020: 27th Conference on Computer and Communications Security*, pages 1575–1590. ACM Press, November 2020. doi:10.1145/3372297.3417872.

[KOS16] Marcel Keller, Emmanuela Orsini, and Peter Scholl. MASCOT: Faster malicious arithmetic secure computation with oblivious transfer. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016: 23rd Conference on Computer and Communications Security*, pages 830–842. ACM Press, October 2016. doi:10.1145/2976749.2978357.

[KS08] Vladimir Kolesnikov and Thomas Schneider. Improved garbled circuit: Free XOR gates and applications. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfsdóttir, and Igor Walukiewicz, editors, *ICALP 2008: 35th International Colloquium on Automata, Languages and Programming, Part II*, volume 5126 of *Lecture Notes in Computer Science*, pages 486–498. Springer, Berlin, Heidelberg, July 2008. doi:10.1007/978-3-540-70583-3_40.

[LEFS25] Jan Lauinger, Jens Ernstberger, Andreas Finkenzeller, and Sebastian Steinhorst. Janus: Fast Privacy-Preserving Data Provenance For TLS 1.3. *Proceedings on Privacy Enhancing Technologies*, 2025(1):511–530, July 2025. doi:10.56553/popets-2025-0028.

[Lib24] Benoît Libert. Vector commitments with proofs of smallness: Short range proofs and more. In Qiang Tang and Vanessa Teague, editors, *PKC 2024: 27th International Conference on Theory and Practice of Public Key Cryptography, Part II*, volume 14602 of *Lecture Notes in Computer Science*, pages 36–67. Springer, Cham, April 2024. doi:10.1007/978-3-031-57722-2_2.

[LJA+18] Andrei Lapets, Frederick Jansen, Kinan Dak Albab, Rawane Issa, Lucy Qin, Mayank Varia, and Azer Bestavros. Accessible privacy-preserving web-based data analysis for assessing and addressing economic inequalities. In Ellen W. Zegura, editor, *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies*, COMPASS '18, New York, NY, USA, 2018. Association for Computing Machinery. doi:10.1145/3209811.3212701.

[LK23] Iuon-Chang Lin and Chun-Wei Kuo. Trustworthy blockchain oracles for smart contracts. In George A. Tsihrintzis, Shiuh-Jeng Wang, and Iuon-Chang Lin, editors, *2021 International Conference on Security and Information Technologies with AI, Internet Computing and Big-data Applications*, pages 379–389, Cham, 2023. Springer International Publishing. doi:10.1007/978-3-031-05491-4_38.

[LMP22] Zeyu Liu, Daniele Micciancio, and Yuriy Polyakov. Large-precision homomorphic sign evaluation using FHEW/TFHE bootstrapping. In Shweta Agrawal and Dongdai Lin, editors, *Advances in Cryptology – ASIACRYPT 2022, Part II*, volume 13792 of *Lecture Notes in Computer Science*, pages 130–160. Springer, Cham, December 2022. doi:10.1007/978-3-031-22966-4_5.

[LP11] Richard Lindner and Chris Peikert. Better key sizes (and attacks) for LWE-based encryption. In Aggelos Kiayias, editor, *Topics in Cryptology – CT-RSA 2011*, volume 6558 of *Lecture Notes in Computer Science*, pages 319–339. Springer, Berlin, Heidelberg, February 2011. doi:10.1007/978-3-642-19074-2_21.

[LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *Advances in Cryptology – EUROCRYPT 2010*, volume 6110 of *Lecture Notes in Computer Science*, pages 1–23. Springer, Berlin, Heidelberg, May / June 2010. doi:10.1007/978-3-642-13190-5_1.

[LY10] Benoît Libert and Moti Yung. Concise mercurial vector commitments and independent zero-knowledge sets with short proofs. In Daniele Micciancio, editor, *TCC 2010: 7th Theory of Cryptography Conference*, volume 5978 of *Lecture Notes in Computer Science*, pages 499–517. Springer, Berlin, Heidelberg, February 2010. doi:10.1007/978-3-642-11799-2_30.

[Man21] Gina Mantica. Using Data Science to Address the Gender and Racial Wage Gap. https://www.bu.edu/articles/2021/using-data-science-to-address-the-gender-and-racial-wage-gap, 2021.

[Mic00] Silvio Micali. Computationally sound proofs. *SIAM Journal on Computing*, 30(4):1253–1298, 2000. doi:10.1137/S0097539795284959.

[MMS16] Christian Mainka, Vladislav Mladenov, and Jörg Schwenk. Do not trust me: Using malicious IdPs for analyzing and attacking single sign-on. In *2016 IEEE European Symposium on Security and Privacy*, pages 321–336. IEEE Computer Society Press, March 2016. doi:10.1109/EuroSP.2016.33.

[MMT+24] Dimitris Mouris, Daniel Masny, Ni Trieu, Shubho Sengupta, Prasad Buddhavarapu, and Benjamin M. Case. Delegated private matching for compute. *Proceedings on Privacy Enhancing Technologies*, 2024(2):49–72, April 2024. doi:10.56553/popets-2024-0040.

[MP24] Nikolaos Makriyannis and Udi Peled. A note on the security of GG18. https://info.fireblocks.com/hubfs/A_Note_on_the_Security_of_GG.pdf, 2024.

[MPD+25] Dimitris Mouris, Christopher Patton, Hannah Davis, Pratik Sarkar, and Nektarios Georgios Tsoutsos. Mastic: Private Weighted Heavy-Hitters and Attribute-Based Metrics. *Proceedings on Privacy Enhancing Technologies*, 2025(1):290–319, July 2025. doi:10.56553/popets-2025-0017.

[MST24] Dimitris Mouris, Pratik Sarkar, and Nektarios Georgios Tsoutsos. PLASMA: Private, lightweight aggregated statistics against malicious adversaries. *Proceedings on Privacy Enhancing Technologies*, 2024(3):4–24, July 2024. doi:10.56553/popets-2024-0064.

[MT21] Dimitris Mouris and Nektarios Georgios Tsoutsos. Zilch: A Framework for Deploying Transparent Zero-Knowledge Proofs. *IEEE Transactions on Information Forensics and Security*, 16:3269–3284, 2021. doi:10.1109/TIFS.2021.3074869.

[NNOB12] Jesper Buus Nielsen, Peter Sebastian Nordholt, Claudio Orlandi, and Sai Sheshank Burra. A new approach to practical active-secure two-party computation. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology – CRYPTO 2012*, volume 7417 of *Lecture Notes in Computer Science*, pages 681–700. Springer, Berlin, Heidelberg, August 2012. doi:10.1007/978-3-642-32009-5_40.

[PCR19] Nicola Paltrinieri, Louise Comfort, and Genserik Reniers. Learning about risk: Machine learning for risk assessment. *Safety Science*, 118:475–486, October 2019. doi:10.1016/j.ssci.2019.06.001.

[Ped92] Torben P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In Joan Feigenbaum, editor, *Advances in Cryptology – CRYPTO'91*, volume 576 of *Lecture Notes in Computer Science*, pages 129–140. Springer, Berlin, Heidelberg, August 1992. doi:10.1007/3-540-46766-1_9.

[Pri25] Primus Labs. Primus FHE. https://github.com/primus-labs/primus-fhe, 2025.

[PSE22] Privacy and Ethereum Foundation Scaling Explorations. TLSNotary, 2022. Accessed: 2014. URL: https://tlsnotary.org.

[QLJ+19] Lucy Qin, Andrei Lapets, Frederick Jansen, Peter Flockhart, Kinan Dak Albab, Ira Globus-Harris, Shannon Roberts, and Mayank Varia. From usability to secure computing and back again. In Heather Richter Lipford, editor, *Proceedings of the Fifteenth USENIX Conference on Usable Privacy and Security*, SOUPS '19, pages 191–210, USA, August 2019. USENIX Association. URL: https://dl.acm.org/doi/10.5555/3361476.3361490.

[Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th Annual ACM Symposium on Theory of Computing*, pages 84–93. ACM Press, May 2005. doi:10.1145/1060590.1060603.

[Res18] E. Rescorla. The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446, August 2018. URL: https://www.rfc-editor.org/rfc/rfc8446.

[Rot11] Ron Rothblum. Homomorphic encryption: From private-key to public-key. In Yuval Ishai, editor, *TCC 2011: 8th Theory of Cryptography Conference*, volume 6597 of *Lecture Notes in Computer Science*, pages 219–234. Springer, Berlin, Heidelberg, March 2011. doi:10.1007/978-3-642-19571-6_14.

[RT25] Relic-Toolkit. RELIC. https://github.com/relic-toolkit/relic, 2025.

[RWG+18] Hubert Ritzdorf, Karl Wüst, Arthur Gervais, Guillaume Felley, and Srdjan Capkun. TLS-N: Non-repudiation over TLS enablign ubiquitous content signing. In *ISOC Network and Distributed System Security Symposium – NDSS 2018*. The Internet Society, February 2018. doi:10.14722/ndss.2018.23272.

[RWGM23] Michael Rosenberg, Jacob D. White, Christina Garman, and Ian Miers. zkcreds: Flexible anonymous credentials from zkSNARKs and existing identity infrastructure. In *2023 IEEE Symposium on Security and Privacy*, pages 790–808. IEEE Computer Society Press, May 2023. doi:10.1109/SP46215.2023.10179430.

[San22] Manuel B. Santos. PECO: methods to enhance the privacy of DECO protocol. Cryptology ePrint Archive, Report 2022/1774, 2022. URL: https://eprint.iacr.org/2022/1774.

[Sha79] Adi Shamir. How to share a secret. *Communications of the Association for Computing Machinery*, 22(11):612–613, November 1979. doi:10.1145/359168.359176.

[TCDP23] Sijun Tan, Weikeng Chen, Ryan Deng, and Raluca Ada Popa. MPCAuth: Multi-factor authentication for distributed-trust systems. In *2023 IEEE Symposium on Security and Privacy*, pages 829–847. IEEE Computer Society Press, May 2023. doi:10.1109/SP46215.2023.10179481.

[TS21] Dmytro Tymokhanov and Omer Shlomovits. Alpha-rays: Key extraction attacks on threshold ECDSA implementations. Cryptology ePrint Archive, Report 2021/1621, 2021. URL: https://eprint.iacr.org/2021/1621.

[VJH21] Alexander Viand, Patrick Jattke, and Anwar Hithnawi. SoK: Fully homomorphic encryption compilers. In *2021 IEEE Symposium on Security and Privacy*, pages 1092–1108. IEEE Computer Society Press, May 2021. doi:10.1109/SP40001.2021.00068.

[VMW+18] Jo Van Bulck, Marina Minkin, Ofir Weisse, Daniel Genkin, Baris Kasikci, Frank Piessens, Mark Silberstein, Thomas F. Wenisch, Yuval Yarom, and Raoul Strackx. Foreshadow: Extracting the keys to the intel SGX kingdom with transient out-of-order execution. In William Enck and Adrienne Porter Felt, editors, *USENIX Security 2018: 27th USENIX Security Symposium*, pages 991–1008. USENIX Association, August 2018. URL: https://www.usenix.org/conference/usenixsecurity18/presentation/bulck.

[WMK16] Xiao Wang, Alex J. Malozemoff, and Jonathan Katz. EMP-toolkit: Efficient MultiParty computation toolkit. https://github.com/emp-toolkit, 2016.

[WRK17a] Xiao Wang, Samuel Ranellucci, and Jonathan Katz. Authenticated garbling and efficient maliciously secure two-party computation. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017: 24th Conference on Computer and Communications Security*, pages 21–37. ACM Press, October / November 2017. doi:10.1145/3133956.3134053.

[WRK17b] Xiao Wang, Samuel Ranellucci, and Jonathan Katz. Global-scale secure multiparty computation. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017: 24th Conference on Computer and Communications Security*, pages 39–56. ACM Press, October / November 2017. doi:10.1145/3133956.3133979.

[WYKW21] Chenkai Weng, Kang Yang, Jonathan Katz, and Xiao Wang. Wolverine: Fast, scalable, and communication-efficient zero-knowledge proofs for boolean and arithmetic circuits. In *2021 IEEE Symposium on Security and Privacy*, pages 1074–1091. IEEE Computer Society Press, May 2021. doi:10.1109/SP40001.2021.00056.

[XYWY24] Xiang Xie, Kang Yang, Xiao Wang, and Yu Yu. Lightweight authentication of web data via garble-then-prove. In Davide Balzarotti and Wenyuan Xu, editors, *USENIX Security 2024: 33rd USENIX Security Symposium*. USENIX Association, August 2024. URL: https://www.usenix.org/conference/usenixsecurity24/presentation/xie-xiang.

[Yao82] Andrew Chi-Chih Yao. Protocols for secure computations (extended abstract). In *23rd Annual Symposium on Foundations of Computer Science*, pages 160–164. IEEE Computer Society Press, November 1982. doi:10.1109/SFCS.1982.38.

[Yao86] Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *27th Annual Symposium on Foundations of Computer Science*, pages 162–167. IEEE Computer Society Press, October 1986. doi:10.1109/SFCS.1986.25.

[YF14] Yuval Yarom and Katrina Falkner. FLUSH+RELOAD: A high resolution, low noise, L3 cache side-channel attack. In Kevin Fu and Jaeyeon Jung, editors, *USENIX Security 2014: 23rd USENIX Security Symposium*, pages 719–732. USENIX Association, August 2014. URL: https://www.usenix.org/conference/usenixsecurity14/technical-sessions/presentation/yarom.

[YSWW21] Kang Yang, Pratik Sarkar, Chenkai Weng, and Xiao Wang. QuickSilver: Efficient and affordable zero-knowledge proofs for circuits and polynomials over any field. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021: 28th Conference on Computer and Communications Security*, pages 2986–3001. ACM Press, November 2021. doi:10.1145/3460120.3484556.

[YWZ20] Kang Yang, Xiao Wang, and Jiang Zhang. More efficient MPC from improved triple generation and authenticated garbling. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020: 27th Conference on Computer and Communications Security*, pages 1627–1646. ACM Press, November 2020. doi:10.1145/3372297.3417285.

[Zam25] Zama. TFHE-rs. https://github.com/zama-ai/tfhe-rs, 2025.

[ZCC+16] Fan Zhang, Ethan Cecchetti, Kyle Croman, Ari Juels, and Elaine Shi. Town crier: An authenticated data feed for smart contracts. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016: 23rd Conference on Computer and Communications Security*, pages 270–282. ACM Press, October 2016. doi:10.1145/2976749.2978326.

[ZMM+20] Fan Zhang, Deepak Maram, Harjasleen Malvai, Steven Goldfeder, and Ari Juels. DECO: Liberating web data using decentralized oracles for TLS. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020: 27th Conference on Computer and Communications Security*, pages 1919–1938. ACM Press, November 2020. doi:10.1145/3372297.3417239.