Private Certifier Intersection

Bishakh Chandra Ghosh
Indian Institute of Technology Kharagpur
ghoshbishakh@gmail.com

Sikhar Patranabis
IBM Research, India
sikhar.patranabis@ibm.com

Dhinakaran Vinayagamurthy
IBM Research, India
dvinaya1@in.ibm.com

Venkatraman Ramakrishna
IBM Research, India
vramakr2@in.ibm.com

Krishnasuri Narayanam
IBM Research, India
knaraya3@in.ibm.com

Sandip Chakraborty
Indian Institute of Technology Kharagpur
sandipc@cse.iitkgp.ac.in

Abstract*—We initiate the study of Private Certifier Intersection (PCI), which allows mutually distrusting parties to establish a trust basis for cross-validation of claims if they have one or more trust authorities (certifiers) in common. This is one of the essential requirements for verifiable presentations in Web 3.0, since it provides additional privacy without compromising on decentralization. A PCI protocol allows two or more parties holding certificates to identify a common set of certifiers while additionally validating the certificates issued by such certifiers, without leaking any information about the certifiers not in the output intersection. In this paper, we formally define the notion of multi-party PCI in the Simplified-UC framework for two different settings depending on whether certificates are required for any of the claims (called PCI-Any) or all of the claims (called PCI-All). We then design and implement two provably secure and practically efficient PCI protocols supporting validation of digital signature-based certificates: a PCI-Any protocol for ECDSA-based certificates and a PCI-All protocol for BLS-based certificates. The technical centerpiece of our proposals is the first secret-sharing-based MPC framework supporting efficient computation of elliptic curve-based arithmetic operations, including elliptic curve pairings, in a black-box way. We implement this framework by building on top of the well-known MP-SPDZ library using OpenSSL and RELIC for elliptic curve operations, and use this implementation to benchmark our proposed PCI protocols in the LAN and WAN settings. In an intercontinental WAN setup with parties located in different continents, our protocols execute in less than a minute on input sets of size 40, which demonstrates the practicality of our proposed solutions.

I. Introduction

In the traditional web (Web 2.0), users are dependent on a limited set of identity and service providers and public Certificate Authorities (CAs) [2] to initiate trusted interactions. Recent trends in decentralization towards Web 3.0 aim to remove such dependencies on centralized service providers. A prominent problem in the decentralized web revolves around identity and trust. Decentralized identifiers (DIDs) [50] and Verifiable Credentials (VCs) [53] enable parties to own and control their identities. This implies a self-sovereign ability to create, update, and selectively share identity records. Importantly, one can prove properties (or *claims*) about themselves without relying on centralized/federated identity providers or a canonical trusted set of CAs [8], [50], [54], as long as the VC issuer (also called a trust anchor [3]) is trusted by both the prover and the verifier of a claim. In a nutshell, existing DID and VC recommendations give users the ability to control their privacy while engaging in a trusted decentralized interaction. But, there are scenarios where these recommendations cannot adequately safeguard user privacy unless we introduce new privacy-preserving mechanisms. In its most general form, the scenario we are concerned about involves two parties wishing to establish a trust basis for future interactions. Service providers in the Semantic Web have encountered such situations, and mechanisms for trust negotiation [56] were proposed to minimize privacy compromise without sacrificing decentralization, albeit for a specific model of service provider-consumer interaction. In grid computing, service-level agreements (SLAs) [28] followed a similar template. This challenge has returned to salience in today's Web3 world, where private and independent blockchain systems have business imperatives to interoperate [9]. The interaction model common to these scenarios involves no a priori trust between the interacting parties, though they may, unbeknownst to each other, possess VCs (or more generally *certificates*) from common trust anchors (or more generally *certifiers*) attesting to different claims.

A trust basis for interoperation can be established between two parties if they can determine that they both possess valid certificates attesting to certain claims, and that these certificates are issued by one or more certifiers that they both trust. But this is hard to do in the absence of a priori trust or knowledge of the counterparty's intentions, or without compromising one's privacy. We can see why this is so by applying the standard VC recommendation, whereby one party makes a Verifiable Presentation (VP) [53] to another, to our scenario. In a typical VC use case, the relationship between credential presenter and verifier is asymmetric, as the verifier is typically a well-known entity from whom the presenter seeks service or approval. The presenter knows at least one certifier that is trusted by it and the verifier. Typically, this requires the verifier to publish its complete list of certifiers so the presenter can determine ones that are commonly trusted by both parties [22]. But in our interaction model, the relationship between parties is symmetrical, as they are both trying to simultaneously prove something to the other. In a standard VP, the presenter is willing to share credentials (albeit selectively) with the verifier. But, if we use this asymmetric VP-based solution in our scenario where neither party knows anything about the other a

Network and Distributed System Security (NDSS) Symposium 2023
27 February - 3 March 2023, San Diego, CA, USA
ISBN 1-891562-83-5
https://dx.doi.org/10.14722/ndss.2023.23563
www.ndss-symposium.org

priori, the revelation of credentials by the party that presents first will automatically give more leverage to the counterparty (verifier), which learns more about the presenter than it reveals.

A naïve adaptation of an asymmetric solution (such as [22]) to our symmetric setting would require both parties to reveal to each other the list of certifiers from which they have valid certificates, and then identify if there is a mutually trusted certifier. This entails complete loss of privacy for both parties, but especially for an honest party if the other behaves maliciously. There are strong reasons why revealing one's complete list of certifiers might not be in one's interest. A business-oriented certifier, for instance, might not like its clientele to be visible to its market competitors. Consider a blockchain interoperability scenario, where shipment carriers on different trade networks certify their respective networks' participants, e.g., Maersk Shipping Company (on the TradeLens network [7]) and the American Bureau of Shipping (ABS). But as Maersk and ABS are market competitors, they may not necessarily want their clients (the certificate holders) to reveal their respective associations [37]. Knowing the clientele of Maersk may benefit ABS, and vice versa; hence there is a privacy cost to revealing certifier lists in a symmetrical interaction unless those certifier lists are identical.

The other privacy violation aspect is from the perspective of the certificate holder. Every certificate possessed indicates an affiliation with some real world entity, often a well-known one; this could include government agencies, political organizations, NGOs, etc., and such affiliations might be sensitive information that could potentially be misused. And here lies the biggest hazard in the naïve trust basis establishment solution: one of the two interacting parties could be malicious and is trying to fish for information about its counterparty's affiliations. A simple attack would be for the malicious party to offer a long list of certifiers, regardless of whether it possesses valid certificates from them, and have the honest counterparty reveal its true certifier list. Now the malicious party knows, and can misuse, the honest party's affiliations, without revealing its own true affiliations. In the context of trust anchors (TAs) in the DID & VC world, where any entity can issue a VC and there does not exist a canonical list or registry of global TAs, it would not be a hard task for a malicious counterparty to list as many of them as possible to mount the attack we just described. Therefore, we can identify a compelling need to maintain certifier privacy and authenticity, which are not addressed by the naïve solution for determining common certifiers. This motivates us to ask the following question:

Can parties owning certificates efficiently identify a common set of certifiers without leaking anything else?

In particular, the parties should not learn any information about certifiers that may be in the lists of other parties but are not in the intersection.

A. Our Contributions

Private Certifier Intersection (PCI). In this paper, we initiate the study of Private Certifier Intersection (PCI) – a cryptographic primitive that aims to answer the above question in the affirmative. Informally speaking, a PCI protocol allows a set of mutually distrusting certificate-holding parties to achieve

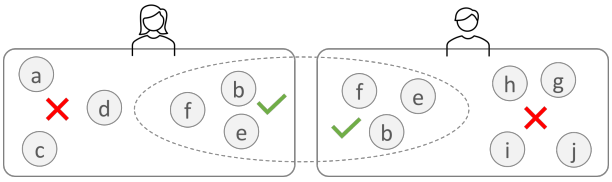

Fig. 1: Private Set Intersection (PSI): Match Values

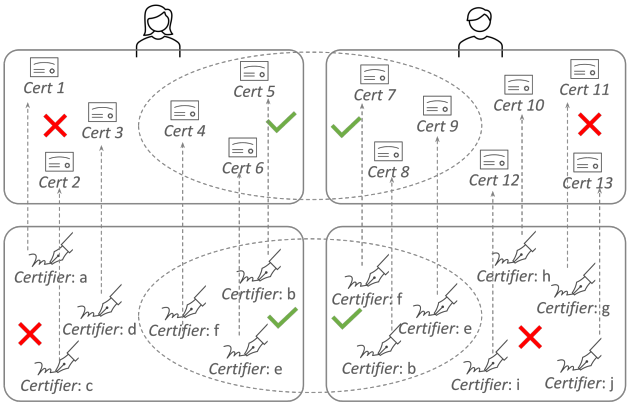

Fig. 2: Private Certifier Intersection (PCI): Match Certificates with Common Issuers

a privacy-preserving trust negotiation with the following objectives: (i) find an intersection among the set of certifiers across the parties, (ii) ensure that the certificates issued by these certifiers are valid, and (iii) reveal no information about the certifiers that may be in the lists of individual parties but are not in the intersection.

Comparison with Private Set Intersection. At a first glance, the classic *Private Set Intersection* (PSI) problem [26], [49], where the intersection of two private sets must be determined without a trusted mediator, bears a strong resemblance to PCI (also see Figure 1). In both PCI and PSI, a set of mutually distrusting parties holding private sets of entities aim to compute the intersection between their sets without revealing any additional information about the elements in their individual sets that are not in the intersection. However, the non-triviality of PCI arises from the need to additionally validate the certificates issued by the certifiers in the intersection. In this sense, one can think of PCI as a form of "predicated" PSI, where the inclusion of a common certifier in the final output set is predicated on the certificates issued by this certifier to each of the parties being valid (see Figure 2 for an illustration). We argue in this paper that realizing an efficient PCI protocol with ideal security guarantees requires novel techniques beyond simply using PSI as a building block. Consider the hazard we encountered earlier in the naïve solution to establish a trust basis. Using standard PSI, a malicious party could simply supply a long (or universal) list of certifiers as input and determine the list of certifiers of the other (honest) party. To avoid this hazard, we need to enforce the ability of participants to prove that they possess genuine certificates issued by thir claimed certifiers. There is no obvious way to do this using standard PSI, and therefore PCI requires novel mechanisms that are not congruent to PSI's mechanisms. We refer the reader to Section I-C for additional related work.

Achieving Semi-Honest PCI. It turns out that in the setting of semi-honest corruptions (i.e., when the participating parties behave honestly as prescribed in the protocol), one can easily achieve a secure PCI protocol by using any secure PSI protocol in a black-box way. Consider the following simple construction: each party first locally "filters" its private list of certifiers based on the validity of the certificates issued by such certifiers, and then uses this filtered list of certifiers as its input to an execution of a PSI protocol to securely identify their intersection. Correctness is immediate, since, assuming honest behavior, the filtered list for each party only contains certifiers issuing valid certificates. Security follows from the security of the underlying PSI protocol.

Upgrading to Malicious Security. Unfortunately, in the setting of malicious corruptions (i.e., when the participating parties can deviate arbitrarily from the protocol), it is seemingly hard to achieve a secure PCI protocol by simply using certification validation and a (maliciously secure) PSI as individual black-boxes. To begin with, we cannot rely on the parties to filter the local sets of certifiers correctly; in fact, the parties can prepare arbitrary sets of certifiers, including those for which it does not have valid certificates.

For example, in the setting of two-party PCI, if one party (say Alice) provides a "universal set" of certifiers as input to a PSI protocol, it can learn the complete set of certifiers of the other party (say Bob). This attack may not be feasible in a general PSI setting where listing the entire range of values in an input set may be infeasible or prohibitively expensive, but is quite feasible in a PCI setting where the range of certifiers (trusted authorities) is limited. Therefore, it is crucial for both Alice and Bob to verify that the other is not faking its input set, and so the validity of certificates and the signatures within must be proven by both parties during the protocol. This is challenging because neither Alice nor Bob knows a priori which set of certifiers it needs to supply proof for (indeed, this is the objective of PCI), and providing more proof than strictly required (i.e., revealing certifiers outside the intersection) would violate privacy goals. Therefore, we must somehow intertwine certificate validation with a PSI-like protocol to achieve PCI. In other words, a maliciously secure PCI protocol cannot be achieved securely without a mechanism that somehow intertwines certificate validation with the subsequent PSI protocol.

Theoretically, a maliciously secure PCI protocol can be achieved as follows: run a maliciously secure multi-party computation (MPC) protocol for the functionality that: (i) filters the certifier list for each party to identify the certifiers issuing valid certificates attesting to the relevant claims, and (ii) computes the intersection between these filtered sets. This solution is highly inefficient in practice for essentially all widely used cryptographically secure certification mechanisms. For example, the most common method of generating certificates is to sign the claim using a digital signature algorithm. In this case, claim validation would require us to perform signature verifications inside the MPC protocol, which is prohibitively expensive for popular digital signature schemes such as ECDSA [10], [41] and BLS [18], [19], [21], that rely on elliptic curve-based finite-field arithmetic operations. Implementing such a verification algorithm inside a maliciously secure MPC protocol would involve non-black-box usage of the various elliptic-curve (EC) operations, i.e., we would have to express these operations as (potentially complicated) binary/arithmetic circuits with gate operations over $\{0, 1\}$ or over some finite field F_p. Such a maliciously secure MPC protocol is likely to incur huge computational and communication overheads in practice.

Need for Efficient Protocols. The above discussion motivates specialized PCI protocols that efficiently enable computing the intersection of certifier-sets while: (i) achieving the desired security guarantees in the setting where a majority of the parties could be maliciously corrupt, and (ii) minimizing non-black-box usage of the operations in the certificate validation algorithm. In this paper, we design and implement two concrete PCI protocols – based on the ECDSA signature scheme and the BLS signature scheme – that achieve the above goal while supporting different variations of claim validation (we expand on this later). While our protocols broadly follow the generic approach outlined above, the main novelty lies in how we validate signatures while using the underlying elliptic curve-based operations in a black-box manner. For an (informal) comparison, the generic MPC-based solution is expected to incur $O(xd)$ computation/communication cost, where x is the corresponding cost of our protocols, and d is the average depth of the arithmetic circuits representing EC operations (e.g., $d = 256$ for constant-time scalar multiplication over curve-ED25519 and curve-BLS12-381).

B. Overview of Contributions

In this section, we provide an informal overview of our key technical contributions.

Defining PCI. We formalize the security guarantees expected of a (multi-party) PCI protocol using the simplified universal composability (SUC) framework due to Canetti, Cohen, and Lindell [24] in the real/ideal world paradigm. We consider two variations of PCI protocols in this paper:

- Validate-Any PCI: A PCI-Any protocol outputs the set of common certifiers for which each party has at least one valid certificate attesting to *any* one of its (publicly known) claims.
- Validate-All PCI: A PCI-All protocol outputs the set of common certifiers for which each party has valid certificates attesting to *all* of its (publicly known) claims.

We also consider a variant of validate-any PCI which we call validate-any PCI with disclosed claims (abbreviated as PCI-Any-DC) where, for each common certifier in the output set, the parties additionally learn the set of claims attested by the certifier. We refer to Section II for a formal description.

MPC for Elliptic Curve Pairings. As a fundamental building block of our proposed PCI protocols, we introduce a new secret-sharing based MPC framework that is tuned for elliptic curve pairings. Our overall approach is to design a secret-sharing based MPC protocol that efficiently supports basic elliptic curve operations (i.e., point addition and scalar multiplication) and elliptic curve bilinear pairing operations as fundamental building blocks. We build upon the SPDZ secret-sharing based MPC protocol [31], [32] to achieve the first secret-sharing based MPC framework that seamlessly supports elliptic curve pairing operations as fundamental gate-level

building-blocks with malicious security against a dishonest majority of adversarial parties. A technical cornerstone of our framework is the round-preserving upgradation of SPDZ from basic field operations to the significantly more complicated elliptic curve operations, including pairings. Our framework allows us to directly use standardized and open-source implementations of elliptic curve libraries [4], [11], [45], thereby leveraging both the performance improvements/optimizations as well as the protections against evolving implementation-level attacks that such libraries usually offer. We believe that this is a contribution of independent interest.

Efficient Two-Party PCI. We use our proposed MPC framework to design the following provably secure yet practically efficient two-party PCI protocols:

- A two-party PCI-Any-DC protocol using the ECDSA signature scheme [10] – an elliptic-curve-based digital signature scheme which is standardized and widely adopted in multiple real-world applications including X.509 public key infrastructure in the Internet, TLS [46], DNSSEC [39], etc. Moreover, ECDSA is a candidate signature scheme in verifiable credentials [53] which is one of the target applications of PCI. Choosing ECDSA also allows us to use its standard implementation in the OpenSSL [4] library for EC group operations. This naturally motivates designing a PCI protocol supporting ECDSA-based certification of claims.
- A two-party PCI-All protocol using the BLS signature scheme [18], [19], [21]– an elliptic-curve pairing-based digital signature that is popularly used in blockchain applications and is in the process of being standardized [20]. We design a PCI-All protocol supporting BLS-based certification of claims that exploits the signature-aggregation capabilities of BLS to perform efficient validation of certificates over all of the public claims of each party.

The starting point of our protocols is the generic maliciously secure PCI protocol outlined earlier, with several optimizations to obviate or minimize expensive elliptic curve operations inside the MPC protocol. In our ECDSA-based PCI-Any-DC protocol, we develop techniques that enable securely yet efficiently performing the expensive algebraic operations (such as field inversion) and non-algebraic operations (such finding the x-coordinate of an elliptic curve point) required by the ECDSA verification algorithm *outside* the MPC protocol. The protocol is then implemented using our proposed MPC framework, which allows performing ECDSA signature validations while using all elliptic curve operations in a black-box manner. We also discuss how to upgrade this protocol to full-fledged PCI-Any where the claims are no longer disclosed publicly (see Section IV for details).

Trivially extending the approach used in our ECDSA-based PCI-Any-DC protocol to design a PCI-All protocol would require iterating through all of the public claims, and validating the signatures on these claims by a specific certifier. This results in a claim validation complexity that grows with the number of claims. We overcome this challenge by designing a PCI-All protocol using BLS-based signature-aggregation that only requires a single (aggregate-)signature verification per certifier inside the MPC protocol. We introduce additional optimizations that exploit the deterministic nature of the BLS signature to further reduce the number of elliptic curve pairing operations inside MPC to just one per certifier, which is then implemented in a black-box manner using our proposed MPC framework over pairings.

Implementation and Evaluation. We extend MP-SPDZ [42] to implement our proposed secret-sharing framework supporting elliptic curve operations including bilinear pairings. For the black-box operations on elliptic curves we use OpenSSL [4] and RELIC [11] libraries. We then implement ECDSA-based PCI-Any-DC and BLS-based PCI-All protocols. We make the source code of our implementation available at https://github.com/irondeveloper321/pci for independent benchmarking (with the repository anonymized for double-blind review). We provide a detailed analysis of the performance of the individual components of our MPC framework, followed by the end-to-end performance evaluation of the protocols in realistic setups by placing parties in three geographic regions across two continents. In an intercontinental WAN setup with parties located in different continents, our PCI-Any-DC and PCI-All protocols execute in less than a minute on input sets of size 40. This demonstrates the practicality of our proposed solutions. We refer to Section VI for details.

C. Related Work

Private Set Intersection (PSI). Private set intersection (PSI) [49] has been extensively studied, with a wide range of solutions based on garbled circuits [40], homomorphic encryption [26], oblivious transfer [49], and other techniques [25], [27], [29], [35], [44], [48], [51]. However, as outlined earlier, there is no straightforward way of using PSI as a black-box to achieve PCI, particularly in the face of malicious adversarial corruptions, due to the additional requirement of certificate validation.

PSI over Certified Sets. Private intersection of "certified sets", introduced in [23], is an augmentation of PSI with the additional requirement that the input claim-sets are certified by some certification authority (CA). However, this primitive has fundamentally different privacy goals as compared to PCI; it assumes that the information of the CAs is public and that the two parties agree apriori on which CAs they mutually trust. Conversely, in the case of PCI, the CAs (certifiers) are, in fact, the input to the protocol (and thus cannot be made public apriori) while the claims are public. We could also have a variant of PCI where the claims are additionally private; we leave it as an open question to investigate this variant further.

HIAC. Hidden-issuer anonymous credentials (HIAC), introduced very recently in [22], is an elegant cryptographic primitive that allows a credential holder to prove its claim(s) to a verifier without disclosing the identity of the credential issuer (i.e., the certifier). However, HIAC inherently requires the set of certifiers trusted by the verifier to be published as an "aggregator", thereby revealing the identity of each such certifier. Hence, while one could use HIAC to solve the same problem at PCI, such an adaptation would only achieve *one-sided privacy* since of the parties would have to make its list of certifiers publicly available. On the other hand, PCI aims to enable *two-sided privacy* by allowing the two parties to negotiate their common certifiers while preserving the privacy

of *both* individual lists, and while simultaneously validating the certificates issued to *both* the parties.

IHABC. Issuer-Hidden Attribute-Based Credential [17] is another related system in which a user can prove a credential issued to it without revealing which issuer among a set of issuers acceptable to the verifier issued that credential. Similar to HIAC, this system also provides one-sided privacy while revealing the certifier set of the verifier (PCI, on the other hand, ensures privacy of both the parties' list of certifiers). Moreover, the concrete solution presented in [17] uses a trusted setup, which is costly in practice and is not a requirement for any of our PCI solutions.

Secret Handshake. The "secret handshake" family of protocols [12], [13] enable (role-based) authenticated key exchange between parties without revealing any information beyond the common group memberships shared by the parties. These protocols, however, differ fundamentally from PCI in the sense that: (a) they do not capture the notion of validating certificates and claims (which is one of the core requirements addressed by PCI), and (b) the process of issuing membership credentials is part of the protocol itself (in PCI, the process of issuing credentials/certificates is not considered part of the primitive).

II. Private Certifier Intersection (PCI)

In this section, we formally define Private Certifier Intersection (PCI). We begin by introducing some notations and background material. We subsequently formalize the functionality and security guarantees that a PCI protocol should satisfy.

General Notations. We write $x \leftarrow \chi$ to represent that an element x is sampled uniformly at random from a set/distribution $\mathcal{X}$. The output x of a deterministic algorithm $\mathcal{A}$ is denoted by $x = \mathcal{A}$ and the output x' of a randomized algorithm $\mathcal{A}'$ is denoted by $x' \leftarrow \mathcal{A}'$. For $a, b \in \mathbb{N}$ such that $a, b \geq 1$, we denote by $[a, b]$ the set of integers lying between a and b (both inclusive). We refer to $\lambda \in \mathbb{N}$ as the security parameter, and denote by $\mathsf{poly}(\lambda)$ and $\mathsf{negl}(\lambda)$ any generic (unspecified) polynomial function and negligible function in λ, respectively.[1]

PCI Notations. Let $\mathcal{ID}$ be a set of identities corresponding to the certifiers. Given a claim $\mathsf{m} \in \mathcal{M}$ by a party P, a certifier with identity id can issue a certificate $\sigma \in \mathcal{C}$, such that there exists a relation $\mathbf{R}$ that satisfies the following:

$$\mathbf{R}(\mathsf{id}, \sigma, \mathsf{m}) = 1 \text{ iff } \sigma \text{ is a valid certificate by } \mathsf{id} \text{ on } \mathsf{m}$$

A natural instantiation of the certification process outlined above is a digital signature, where the certificate issuance corresponds to the signing algorithm and the relation $\mathbf{R}$ corresponds to the verification algorithm, with σ being the signature on a claim m under the signing key corresponding to id. Looking ahead, our proposed realizations of PCI protocols in this paper will use this digital signature-based instantiation of the certification process.

We now introduce some additional notations for ease of exposition, these notations will be useful in understanding our definitions for PCI. Let S be a set of (identity, certificate, claim) tuples of the form

$$S = \{(\mathsf{id}_j, \sigma_j, \mathsf{m}_j) \in \mathcal{ID} \times \mathcal{C} \times \mathcal{M}\}_{j \in [1,n]}$$

where N is the number of tuples in the set S. We define the following projection functions on the set S:

$$\mathsf{id}(S) := \{\mathsf{id} : \exists \sigma, \mathsf{m} \text{ s.t. } (\mathsf{id}, \sigma, \mathsf{m}) \in S\}$$
$$\mathsf{m}(S) := \{\mathsf{m} : \exists \mathsf{id}, \sigma \text{ s.t. } (\mathsf{id}, \sigma, \mathsf{m}) \in S\}$$
$$\overline{\mathsf{m}}(S) := (\mathsf{m}_j)_{(\mathsf{id}_j, \sigma_j, \mathsf{m}_j) \in S}$$

Here, $\overline{\mathsf{m}}(S)$ is a list/multiset of the claims corresponding to each tuple in the set S.

A. Defining a PCI Protocol

We now formally define a PCI protocol in the two-party setting, which is the focus of this paper. Our definitions naturally extend to multiple parties, as discussed subsequently.

Two-Party PCI. A two-party PCI protocol Π involves parties P_1 and P_2, where each party P_i for $i \in \{1, 2\}$ inputs a tuple of the form $\mathsf{inp}_i = \left(\mathsf{inp}_{i,1}, \mathsf{inp}_{i,2}\right)$, where:

- The *private* input $\mathsf{inp}_{i,1}$ is a set of (identity, certificate, claim) tuples of the form
$$\mathsf{inp}_{i,1} = \{(\mathsf{id}_{i,j}, \sigma_{i,j}, \mathsf{m}_{i,j}) \in \mathcal{ID} \times \mathcal{C} \times \mathcal{M}\}_{j \in [1, N_i]}$$
where N_i is the number of tuples in $\mathsf{inp}_{i,1}$ from party P_i.
- The *public* input $\mathsf{inp}_{i,2}$ is a set of claims of the form $\{\widehat{\mathsf{m}}_{i,j} \in \mathcal{M}\}_{j \in [1, N_i']}$, where N_i' is the number of tuples in $\mathsf{inp}_{i,2}$ from party P_i.

Note that a party P_i can produce multiple certificates from the same certifier on same or different claims. Additionally, a party P_i can also request certifications on the same claim from multiple certifiers. Hence, in the most general setting, a party's input could have multiple tuples with the common id or a common m. Also note that the public input for P_1 is known to P_2 at the start of the protocol and vice versa.[2]

Remark. A couple of remarks on the definition follow:

1) One could have a variant of PCI with the claims being private. This work considers the above defined variant with the claims being public. We leave it to future work for instantiating PCI with private claims.
2) Our definition lets a (corrupt) party provide claims in the public input that are different from those in the tuple in the private input. One could also restrict the public input $\mathsf{inp}_{i,2}$ to be $\mathsf{m}(\mathsf{inp}_{i,1})$, which is the expected behaviour of the honest parties.

At the end of the protocol Π, each party P_i receives as output a set of certifiers. In this paper, we consider different variations of (two-party) PCI protocols that produce different kinds of output sets, that we outline below:

- **Validate-Any:** In this flavor of PCI protocol, denoted by PCI-Any, both parties P_1 and P_2 receive as output the set of certifiers $\mathsf{out}_{\mathsf{PCI\text{-}Any}}$, such that an identity $\mathsf{id} \in \mathsf{out}$ if

[1] Note that a function $f : \mathbb{N} \to \mathbb{N}$ is said to be negligible in λ if for every positive polynomial p, $f(\lambda) < 1/p(\lambda)$ when λ is sufficiently large.

[2] We assume that these sets are shared between P_1 and P_2 via some apriori mechanism that is not within the purview of the PCI protocol itself.

and only if both P_1 and P_2 have valid certificates on some $\mathsf{m}_1 \in \mathsf{inp}_{i,2}$ and $\mathsf{m}_2 \in \mathsf{inp}_{i,2}$, respectively, such that both the certificates are issued by id. More formally, for each $i \in \{1,2\}$, we define the following Boolean predicate:

$$\mathbf{R}_{\mathsf{PCI\text{-}Any},\mathsf{inp}_i}(\mathsf{id}) = 1 \text{ if and only if } \exists \mathsf{m} \in \mathsf{inp}_{i,2} : \exists(\mathsf{id},\mathsf{m},\sigma) \in \mathsf{inp}_{i,1} \text{ s.t. } \mathbf{R}(\mathsf{id},\mathsf{m},\sigma) = 1$$

Then we have

$$\mathsf{out}_{\mathsf{PCI\text{-}Any}}(\mathsf{inp}_1,\mathsf{inp}_2) = \big\{\mathsf{id} \in \mathsf{id}(\mathsf{inp}_{1,1}) \cap \mathsf{id}(\mathsf{inp}_{2,1}) : \mathbf{R}_{\mathsf{PCI\text{-}Any},\mathsf{inp}_1}(\mathsf{id}) = \mathbf{R}_{\mathsf{PCI\text{-}Any},\mathsf{inp}_2}(\mathsf{id}) = 1\big\}$$

- **Validate-Any with Disclosed Claims:** We also consider a weaker variant of the aforementioned validate-any PCI protocol (denoted by PCI-Any-DC), where the parties additionally learn the following: (i) the claim $\mathsf{m}_{i,j}$ corresponding to each tuple $(\mathsf{id}_{i,j},\sigma_{i,j},\mathsf{m}_{i,j}) \in \mathsf{inp}_{i,1}$ for each party P_i, (ii) for each id in the output set of certifiers $\mathsf{out}_{\mathsf{PCI\text{-}Any}}$, each party learns the set of (public) claims on which the other party has a valid certificate issued by id. Note that no information is revealed about any (valid/invalid) certificates that the parties might have that are issued by some $\mathsf{id}' \notin \mathsf{out}_{\mathsf{PCI\text{-}Any}}$. Formally, for each $i \in \{1,2\}$, we define the function

$$\mathsf{m}_{\mathsf{inp}_i}(\mathsf{id}) = \big\{\mathsf{m} : \exists(\mathsf{id},\mathsf{m},\sigma) \in \mathsf{inp}_{i,1} \text{ s.t. } \mathbf{R}(\mathsf{id},\mathsf{m},\sigma) = 1\big\}$$

Then the output set $\mathsf{out}_{\mathsf{PCI\text{-}Any\text{-}DC}}$ is described formally as follows

$$\mathsf{out}_{\mathsf{PCI\text{-}Any\text{-}DC}}(\mathsf{inp}_1,\mathsf{inp}_2) = \Big(\big\{\overline{\mathsf{m}}(\mathsf{inp}_{i,1})\big\}_{i\in[1,2]}, \big\{\big(\mathsf{id},\{\mathsf{m}_{\mathsf{inp}_i}(\mathsf{id})\}_{i\in\{1,2\}}\big) : \mathsf{id} \in \mathsf{out}_{\mathsf{PCI\text{-}Any}}(\mathsf{inp}_1,\mathsf{inp}_2)\big\}\Big)$$

PCI-Any-DC is relevant in most real-world scenarios since the parties would know the claims of the counterparty that they want to validate, and vice versa. Moreover, traditional VC interactions also work on disclosed claims (see Section I).
- **Validate-All:** In this flavor of PCI protocol, denoted by PCI-All, both parties P_1 and P_2 receive as output the set of certifiers $\mathsf{out}_{\mathsf{PCI\text{-}All}}$, such that for each $\mathsf{id} \in \mathsf{out}_{\mathsf{PCI\text{-}All}}$, P_1 and P_2 have valid certificates issued by id on *all* of the (public) claims in their input sets $\mathsf{inp}_{1,2}$ and $\mathsf{inp}_{2,2}$, respectively. More formally, for each $i \in \{1,2\}$, we define the following Boolean predicate:

$$\mathbf{R}_{\mathsf{PCI\text{-}All},\mathsf{inp}_i}(\mathsf{id}) = 1 \text{ if and only if } \forall \mathsf{m} \in \mathsf{inp}_{i,2} : \exists(\mathsf{id},\mathsf{m},\sigma) \in \mathsf{inp}_{i,1} \text{ s.t. } \mathbf{R}(\mathsf{id},\mathsf{m},\sigma) = 1$$

Then we have

$$\mathsf{out}_{\mathsf{PCI\text{-}All}}(\mathsf{inp}_1,\mathsf{inp}_2) = \big\{\mathsf{id} \in \mathsf{id}(\mathsf{inp}_{1,1}) \cap \mathsf{id}(\mathsf{inp}_{2,1}) : \mathbf{R}_{\mathsf{PCI\text{-}All},\mathsf{inp}_1}(\mathsf{id}) = \mathbf{R}_{\mathsf{PCI\text{-}All},\mathsf{inp}_2}(\mathsf{id}) = 1\big\}$$

B. Security of PCI

We now define the security guarantees expected of a PCI protocol in the two-party setting. Informally, we require that in any PCI protocol Π, party P_1 (resp. party P_2) learns nothing about the inputs of party P_2 (resp. party P_1) except what is revealed by the output out of the protocol Π, and the sizes N_1 and N_2 of the input sets of P_1 and P_2. In the rest of this section, we formalize this security guarantee using the simplified universal composability (SUC) framework due to Canetti, Cohen, and Lindell [24] in the real/ideal world paradigm. We consider a *dishonest majority* in our definitions,

$\underline{\mathcal{F}_{\mathsf{PCI}}(\mathsf{mode} \in \{\mathsf{Any},\mathsf{Any\text{-}DC},\mathsf{All}\})}$

- For $i \in \{1,2\}$, let the input of party P_i be $\mathsf{inp}_i = (\mathsf{inp}_{i,1},\mathsf{inp}_{i,2})$, where
$$\mathsf{inp}_{i,1} = \{(\mathsf{id}_{i,j},\sigma_{i,j},\mathsf{m}_{i,j}) \in \mathcal{ID}\times\mathcal{C}\times\mathcal{M}\}_{j\in[1,N_i]}$$
$$\mathsf{inp}_{i,2} = \{\widehat{\mathsf{m}}_{i,j} \in \mathcal{M}\}_{j\in[1,N_i']}$$
The honest party P_2 provides its input directly to $\mathcal{F}_{\mathsf{PCI}}$, while the input of the corrupt party P_1 is provided to $\mathcal{F}_{\mathsf{PCI}}$ by the simulator $\mathcal{S}$.
- $\mathcal{F}_{\mathsf{PCI}}$ computes $\mathsf{out}_{\mathsf{PCI\text{-}mode}}$, where for $\mathsf{mode} \in \{\mathsf{Any},\mathsf{Any\text{-}DC},\mathsf{All}\}$, we have
$$\mathsf{out}_{\mathsf{PCI\text{-}Any}}(\mathsf{inp}_1,\mathsf{inp}_2) = \big\{\mathsf{id} \in \mathsf{id}(\mathsf{inp}_{1,1}) \cap \mathsf{id}(\mathsf{inp}_{2,1}) : \mathbf{R}_{\mathsf{PCI\text{-}Any},\mathsf{inp}_1}(\mathsf{id}) = \mathbf{R}_{\mathsf{PCI\text{-}Any},\mathsf{inp}_2}(\mathsf{id}) = 1\big\}$$
$$\mathsf{out}_{\mathsf{PCI\text{-}Any\text{-}DC}}(\mathsf{inp}_1,\mathsf{inp}_2) = \Big(\big\{\overline{\mathsf{m}}(\mathsf{inp}_{i,1})\big\}_{i\in[1,2]}, \big\{\big(\mathsf{id},\{\mathsf{m}_{\mathsf{inp}_i}(\mathsf{id})\}_{i\in\{1,2\}}\big) : \mathsf{id} \in \mathsf{out}_{\mathsf{PCI\text{-}Any}}(\mathsf{inp}_1,\mathsf{inp}_2)\big\}\Big)$$
$$\mathsf{out}_{\mathsf{PCI\text{-}All}}(\mathsf{inp}_1,\mathsf{inp}_2) = \big\{\mathsf{id} \in \mathsf{id}(\mathsf{inp}_{1,1}) \cap \mathsf{id}(\mathsf{inp}_{2,1}) : \mathbf{R}_{\mathsf{PCI\text{-}All},\mathsf{inp}_1}(\mathsf{id}) = \mathbf{R}_{\mathsf{PCI\text{-}All},\mathsf{inp}_2}(\mathsf{id}) = 1\big\}$$
- $\mathcal{F}_{\mathsf{PCI}}$ sends $\big(\mathsf{out}_{\mathsf{PCI\text{-}mode}},N_1,\mathsf{inp}_{1,2}\big)$ to the simulator $\mathcal{S}$.
- If $\mathcal{S}$ responds with an abort, $\mathcal{F}_{\mathsf{PCI}}$ aborts.
- Otherwise, $\mathcal{F}_{\mathsf{PCI}}$ sends to P_1 and P_2 the tuple
$$\big(\mathsf{out}_{\mathsf{PCI\text{-}mode}},N_1,N_2,\mathsf{inp}_{1,2},\mathsf{inp}_{2,2}\big)$$

Fig. 3: Ideal functionality $\mathcal{F}_{\mathsf{PCI}}$ in the two-party setting

wherein the adversary can corrupt one of the two participating parties. For ease of exposition, we assume without loss of generality that P_1 and P_2 are the corrupt party and the honest party, respectively.

Ideal Functionality for PCI. We begin by formally defining the first component of our simulation-based security definition, namely the ideal functionality $\mathcal{F}_{\mathsf{PCI}}$, as described in Figure 3. This functionality $\mathcal{F}_{\mathsf{PCI}}$ formally defines what each party is meant to learn at the completion of the protocol.

The Real World. In the real world, the following participants engage in the protocol Π:

- The honest party P_2 that receives its input from the environment $\mathcal{Z}$ and honestly follows the protocol Π.
- A real-world adversary $\mathcal{A}$ that controls the corrupt party P_1, and interacts with P_2 and the environment $\mathcal{Z}$.
- The environment $\mathcal{Z}$ that provides P_2 with its input, and interacts with the real-world adversary $\mathcal{A}$. The environment $\mathcal{Z}$ also receives the output of P_2, and eventually outputs a bit $b \in \{0,1\}$.

The Ideal World. In the ideal world, the following participants interact with the ideal functionality $\mathcal{F}_{\mathsf{PCI}}$ described in Figure 3.

- The honest party P_2 that receives its input from the environment $\mathcal{Z}$ and directly forwards this input to $\mathcal{F}_{\mathsf{PCI}}$.
- An ideal-world simulator $\mathcal{S}$ that sends inputs to $\mathcal{F}_{\mathsf{PCI}}$ on behalf of the corrupt party P_1 and receives back the corresponding output from $\mathcal{F}_{\mathsf{PCI}}$. $\mathcal{S}$ also interacts with the environment $\mathcal{Z}$, with the aim of making this interaction indistinguishable from the interaction between the real world $\mathcal{A}$ and the environment $\mathcal{Z}$.

- The environment $\mathcal{Z}$ that provides P_2 with its input, and interacts with the simulator $\mathcal{S}$. As in the real world, $\mathcal{Z}$ also receives the output of P_2, and eventually outputs a bit $b \in \{0,1\}$.

For any two-party PCI protocol Π, any adversary $\mathcal{A}$, any simulator $\mathcal{S}$, and any environment $\mathcal{Z}$, define the following random variables:

- $\mathsf{real}_{\Pi,\mathcal{A},\mathcal{Z}}$: a random variable that denotes the output of the environment $\mathcal{Z}$ after interacting with the adversary $\mathcal{A}$ during an execution of the real-world protocol Π.
- $\mathsf{ideal}_{\mathcal{F}_{\mathsf{PCI}},\mathcal{S},\mathcal{Z}}$: a random variable that denotes the output of the environment $\mathcal{Z}$ after interacting with the simulator $\mathcal{S}$ in the ideal world.

Definition 1 (Secure Two-Party PCI)**.** A PCI protocol Π securely emulates the ideal functionality $\mathcal{F}_{\mathsf{PCI}}$ described in Figure 3 if for any security parameter $\lambda \in \mathbb{N}$ and any probabilistic polynomial time (PPT) adversary $\mathcal{A}$, there exists a PPT simulator $\mathcal{S}$ such that, for any PPT environment $\mathcal{Z}$,

$$|\Pr[\mathsf{real}_{\Pi,\mathcal{A},\mathcal{Z}} = 1] - \Pr[\mathsf{ideal}_{\mathcal{F}_{\mathsf{PCI}},\mathcal{S},\mathcal{Z}} = 1]| \leq \mathsf{negl}(\lambda)$$

C. Extensions and Generic Constructions

Multi-Party PCI. Our definition of two-party PCI naturally extends to the more general setting of multi-party PCI involving n parties $P_1, \ldots, P_n$. We defer a formal treatment of multi-party PCI to the full version of our paper [36].

Generic Construction of PCI. In the full version of our paper [36], we describe a generic approach to achieving a semi-honest secure PCI-mode protocol for mode $\in \{\mathsf{Any}, \mathsf{Any\text{-}DC}, \mathsf{All}\}$ given any semi-honest secure private set intersection (PSI) protocol. We also discuss in [36] how to upgrade this construction to malicious security and the challenges thereof, and present a discussion on why the generic construction is practically infeasible and why we need concretely efficient PCI protocols in practice.

III. MPC for Elliptic Curve Pairings

As a fundamental building block of our proposed PCI protocols, we introduce a new secret-sharing based MPC framework that is tuned for elliptic curve pairings. In this section, we describe this framework. Our framework is based on the SPDZ family of secret-sharing based MPC protocols. In the full version of our paper [36], we present a detailed discussion on why we choose secret-sharing based MPC, and more specifically SPDZ, for our proposed framework and for our PCI protocols.

Our Framework for MPC over EC Pairings. We now detail our framework for designing secret-sharing based MPC protocols over EC pairings. Our framework can be broadly divided into three-tiers, where each tier builds upon the preceding one.

- **Tier-1:** This tier of our framework supports the basic operations over F_p for some prime p.
- **Tier-2:** This tier of our framework supports group operations over any generic group $\mathcal{G}$ with order p. We use this tier to implement basic EC operations over the source groups of an EC pairing (i.e., point addition and scalar multiplication), as well as the group operations over the output group of the EC pairing (i.e., multiplication and exponentiation).
- **Tier-3:** This tier of our framework supports EC pairing operations, subject to the restriction that the pairing map e takes its inputs from two source groups $\mathcal{G}_1$ and $\mathcal{G}_2$, both of which have order p, and produces an output in a target group $\mathcal{G}_T$, also of order p.

$\underline{\mathcal{F}[F_p]}$

Init-F: On input (init, F_p) from all parties, the functionality stores $(domain, F_p)$. A list of identifiers is established for F_p, if not already done before.
Input-F: On input $(\mathsf{inp}F, P_i, \mathsf{varid}, x)$ with $x \in F_p$ from P_i and $(\mathsf{inp}F, P_i, \mathsf{varid}, \phi_{F_p})$ from all other parties, with varid a fresh identifier, the functionality stores (varid, x) in the list of field identifiers.
Rand-F: On input $(\mathsf{rand}, \mathsf{varid})$ from all parties (if varid is not stored in memory), the functionality generates a uniformly random $a \in F_p$ and stores (varid, a) in the list of field identifiers.
Triple-F: On input $(\mathsf{triple}, \mathsf{varid}_1, \mathsf{varid}_2, \mathsf{varid}_3)$ from all parties (if none of the varid_i are stored in memory), the functionality generates a uniformly random $a, b \in F_p$ and computes $c = a \cdot b$ and then stores (varid_1, a), (varid_2, b) and (varid_3, c) in the list of field identifiers.
Add-F: On command $(\mathsf{add}F, \mathsf{varid}_1, \mathsf{varid}_2, \mathsf{varid}_3)$ from all parties where $\mathsf{varid}1$, varid_2 are in the list of field identifiers and varid_3 is not, the functionality retrieves (varid_1, x), (varid_2, y) from the list of field identifiers and stores $(\mathsf{varid}_3, x + y)$ in the list of field identifiers.
Mult-F: On command $(\mathsf{mult}F, \mathsf{varid}_1, \mathsf{varid}_2, \mathsf{varid}_3)$ from all parties where varid_1, varid_2 are in the list of field identifiers and varid_3 is not, the functionality retrieves (varid_1, x), (varid_2, y) from the list of field identifiers and stores $(\mathsf{varid}_3, x \cdot y)$ in the list of field identifiers.
Output-F: On input $(\mathsf{out}F, \mathsf{varid}, i)$ from all honest parties (if varid is present in the list of field identifiers), the functionality retrieves (varid, y) from the set of field identifiers and outputs it to the environment. The functionality waits for an input from the environment. If this input is Deliver then y is output to all parties if $i = 0$, or y is output to party P_i if $i \neq 0$. If the adversarial input is not equal to Deliver then ϕ is output to all parties.

Fig. 4: Ideal functionality for MPC over field operations in F_p

While each tier supports a different set of operations, we exploit the fact that each tier shares a common algebraic structure (up to group homomorphisms), and we can manoeuvre over this structure to progressively support more complicated operations. We now describe each of these tiers in greater details below.

A. Tier-1: MPC for Basic F_p Operations

Our starting point is a secret-sharing based MPC engine for operating over secret-shared inputs in some field F_p that implements the ideal functionality $\mathcal{F}[F_p]$ as described in Figure 4. This engine can be realized directly using SPDZ (the SPDZ-based realization ensures security against both semi-honest and malicious corruption of parties by using an additional authentication mechanism to enforce honesty of operations over secret-shared values). We use the representation $[x]$ for any $x \in F_p$ to denote that the value x is secret-shared, i.e., no individual party has access to x, but each party has access to some share of x (for simplicity, we will assume that this notation incorporates the additional authentication components required to ensure malicious security).

Linearity-Preservation. Fundamentally, we require that the secret-shared representation $[x]$ is "linearity-preserving", i.e., for any $x, y, z, \alpha, \beta \in F_p$ such that $u = \alpha \cdot x + \beta \cdot y + z$, given

the secret shares $[x]$ and $[y]$ and the public values z, α, β, the parties can compute a secret-sharing of u "for free" as

$$[u] = \alpha \cdot [x] + \beta \cdot [y] + z$$

Note that, in the case of malicious security, we also need this property to be preserved for the authentication components.

Additional Functionalities. We additionally require two deterministic functionalities to be supported by the MPC engine:

1) A functionality that "opens" a secret shared value $[x]$, i.e., reconstructs and distributes the value x to all or a subset of the parties.
2) A functionality that "multiplies" secret shared inputs, i.e., given two secret-shared inputs $[x]$ and $[y]$, produces a secret-shared output $[z]$ such that $z = x \cdot y$.

Finally, we require two randomized functionalities to be supported by the MPC engine:

1) A functionality that generates a secret-shared representation $[a]$ for a randomly sampled value $a \leftarrow F_p$.
2) A functionality that generates secret-shared representations of uniformly random multiplicative "triples", i.e., it generates $[a]$, $[b]$ and $[c]$ for $a, b \leftarrow F_p$ and $c = a \cdot b$.

We refer to $\mathcal{F}[F_p]$ described in Figure 4 for a formal description of these functionalities. Note that, for malicious security, we would need each of the above functionalities to also preserve (or, in the case of opening, validate) the authentication components of the output appropriately.

SPDZ-based Realization. While we can use any secret-sharing-based MPC engine that securely realizes $\mathcal{F}[F_p]$, we choose to use SPDZ as a concrete realization, with security against a malicious corruption of the majority of the parties. We briefly recall here that, in addition to securely implementing $\mathcal{F}[F_p]$, SPDZ also implements a MAC-check based authentication mechanism for secret-shared values $[x]$ to achieve active security against malicious corruption of parties. We recall the details of this mechanism at a very high level; the low-level details are not important for understanding our proposed framework. Informally, in SPDZ, each party P_i for $i \in [1, n]$ holds a sharing of a global MAC-key $\alpha \in F_p$ (this sharing follows a slightly different mechanism; we omit the details as our framework is oblivious to the same). Any value $x \in F_p$ is shared as

$$[x] = (\delta, (x_1, \ldots, x_n), (\gamma_1(x), \ldots, \gamma_n(x))),$$

where for each $i \in [n]$, party P_i holds the tuple $(x_i, \gamma_i(x), \delta)$ and where the following invariant holds:

$$x = \sum_{i \in [n]} x_i, \quad \alpha \cdot (x + \delta) = \sum_{i \in [n]} \gamma_i(x).$$

The SPDZ Opening Protocol. we briefly recall how the "opening" protocol in SPDZ allows the parties to authenticate, via a MAC-check mechanism, that a secret-shared value has been opened correctly. The opening protocol for a secret-shared value $[x]$ involves the following steps:

- Each party P_i, upon receiving a reconstructed value x', uses its share α_i of the global MAC-key α, as well as $\gamma_i(x)$ and δ, to compute $\sigma_i = \gamma_i(x) - \alpha_i \cdot (x' + \delta)$.
- Each party P_i then broadcasts a commitment $\mathsf{Com}(\sigma_i)$ to all the other parties.
- Finally, each party P_i opens the commitments $\{\mathsf{Com}(\sigma_j)\}$ received from $\{P_j\}_{j \neq i}$, computes $\mathsf{chk} = \sum_{j \in [n]} \sigma_j$, and aborts if $\mathsf{chk} \neq 0$.

We use the term *partial opening* to refer the procedure that just publicly reconstructs the value x without going through the subsequent MAC-check procedure.

Suppose that a malicious adversary $\mathcal{A}$ manages to add an error ϵ during the reconstruction phase, i.e., we have $x' = x + \epsilon$. Suppose also that the adversary $\mathcal{A}$ commits to a subset of false $\{\sigma'_j\}_{j \in \mathcal{C}}$ values corresponding to the subset $\mathcal{C} \subset [n]$ of parties it corrupts. In order to bypass the MAC-check, the adversary $\mathcal{A}$ must ensure that

$$\sum_{j \in \mathcal{C}} (\sigma'_j - \sigma_j) = \alpha\epsilon.$$

However, this happens with probability no greater than $1/p$, since the global MAC value α is uniformly random in F_p and (information-theoretically) unknown to $\mathcal{A}$, and hence, $\mathcal{A}$ cannot bypass the MAC-check protocol except with negligible probability.

Additional Functionalities in SPDZ. We note that the randomized functionalities for generating secret-shared representations of singleton values or multiplicative triples are implemented by the offline phase of SPDZ [43]. We omit the low-level details of these functionalities because they are not necessary to understand our framework and proposed protocols; it suffices to state that our framework uses the native implementations of these functionalities directly from SPDZ. We also directly use SPDZ's implementation of the functionality for multiplying secret-shared values, which is based on generating a random multiplicative triple and then using Beaver's re-randomization technique. We refer to [31], [32] for the details.

B. Tier-2: MPC over any Generic Group

In Tier-2, we aim to realize an MPC protocol over any generic group $\mathcal{G}$ with prime order p. More concretely, we require the MPC protocol to implement the ideal functionality $\mathcal{F}[\mathcal{G}]$ as described in Figure 5. Such a protocol would allow us to support basic EC operations (i.e., point addition and scalar multiplication) over the source groups of an EC pairing, as well as the operations over the target group of the EC pairing (i.e., group multiplication and exponentiation). Similar to Tier-1, we use a linearity-preserving representation $[\cdot]_{\mathcal{G}}$ for elements in $\mathcal{G}$ such that for any $g_1, g_2, g_3 \in \mathcal{G}$ and any $\alpha, \beta \in Z_p$ such that $h = g_1^{\alpha} \cdot g_2^{\beta} \cdot g_3$, given the secret shares $[g_1]_{\mathcal{G}}$ and $[g_2]_{\mathcal{G}}$ and the public values g_3, α, β, the parties can locally compute

$$[h]_{\mathcal{G}} = {[g_1]_{\mathcal{G}}}^{\alpha} \cdot {[g_2]_{\mathcal{G}}}^{\beta} \cdot g_3$$

Once again, in the case of malicious security, we need this property to be preserved for the authentication components.

Homomorphic Relation with Tier-1. We note that the aforementioned linearity-preservation property in $\mathcal{G}$ shares a similar algebraic structure with the tier-1 linearity-preservation property in F_p described earlier. Let $F_p = Z_p$, and let

$$g_1 = g^x, \quad g_2 = g^y, \quad g_3 = g^z, \quad h = g^u.$$

$\underline{\mathcal{F}[\mathcal{G}]}$

Init-G: On input $(\mathsf{init}, \mathcal{G})$ from all parties, the functionality stores $(domain, \mathcal{G})$. A list of identifiers is established for $\mathcal{G}$, if not already done before.

Input-G: On input $(\mathsf{inp}\mathcal{G}, P_i, \mathsf{varid}, g)$ with $g \in \mathcal{G}$ from P_i and $(\mathsf{inp}\mathcal{G}, P_i, \mathsf{varid}, \phi_{\mathcal{G}})$ from all other parties, with varid a fresh identifier, the functionality stores (varid, g) in the list of field identifiers.

Op-G: On command $(\mathsf{op}\mathcal{G}, \mathsf{varid}_1, \mathsf{varid}_2, \mathsf{varid}_3)$ from all parties where $\mathsf{varid1}$, varid_2 are in the list of group identifiers and varid_3 is not, the functionality retrieves (varid_1, g), (varid_2, h) from the list of group identifiers and stores $(\mathsf{varid}_3, g \cdot h)$ in the list of group identifiers, where $\cdot$ is the group operation.

Exp-G-P: On command $(\exp \mathcal{G}P, \mathsf{varid}_1, g, \mathsf{varid}_2)$ from all parties where varid_1 is in the list of field identifiers, $g \in \mathcal{G}$, and varid_2 is a fresh identifier in the list of group identifiers, the functionality retrieves (varid_1, x) from the list of field identifiers and stores (varid_2, g^x).

Exp-G-S: On command $(\exp \mathcal{G}S, \mathsf{varid}_1, \mathsf{varid}_2, \mathsf{varid}_3)$ from all parties where varid_1 is in the list of field identifiers, varid_2 is in the list of group identifiers, and varid_3 is a fresh identifier in the list of group identifiers, the functionality retrieves (varid_1, x) from the list of field identifiers and (varid_2, h) from the list of group identifiers and stores (varid_2, h^x).

Output-G: On input $(\mathsf{out}\mathcal{G}, \mathsf{varid}, i)$ from all honest parties (if varid is present in the list of group identifiers), the functionality retrieves (varid, g) from the set of group identifiers and outputs it to the environment. The functionality waits for an input from the environment. If this input is Deliver then g is output to all parties if $i = 0$, or g is output to party P_i if $i \neq 0$. If the adversarial input is not equal to Deliver then ϕ is output to all parties.

Fig. 5: Ideal functionality for MPC over the group operations in $\mathcal{G}$, which includes basic EC operations and the operations over the output group of a pairing. We assume that $\mathcal{F}[\mathcal{G}]$ also includes all Tier-1 sub-functionalities in $\mathcal{F}[F_p]$, but we avoid re-writing them for modularity.

Then observe that the linearity-preservation property in Z_p with u

Additional Functionalities. We additionally require three deterministic functionalities to be supported by the MPC engine:

1) A functionality that "opens" a secret shared value $[g]_{\mathcal{G}}$, i.e., reconstructs and distributes the group element g to all or a subset of the parties.
2) A functionality that "exponentiates" a publicly available group element in $\mathcal{G}$ using a secret-shared value in Z_p, i.e., given a public $g \in \mathcal{G}$ and a secret-shared value $[x]$ for $x \in Z_p$, produces a secret-shared output $[h]_{\mathcal{G}}$ such that $h = g^x$.
3) A functionality that "exponentiates" a secret-shared group element in $\mathcal{G}$ using a secret-shared value in Z_p, i.e., given a secret-shared element $[g]_{\mathcal{G}}$ for $g \in \mathcal{G}$ and a secret-shared value $[x]$ for $x \in Z_p$, produces a secret-shared output $[h]_{\mathcal{G}}$ such that $h = g^x$.

We refer to $\mathcal{F}[\mathcal{G}]$ described in Figure 5 for a formal description of these functionalities. Once again, for malicious security, we would need each of the above functionalities to preserve (or, in the case of opening, validate) the authentication components of the output appropriately.

Tier-2 Extension of SPDZ. As a concrete instantiation of $\mathcal{F}[\mathcal{G}]$, we generalize the extensions to SPDZ for basic EC operations proposed in [30], [52] to any generic group of order p. We briefly recall the details of the approach, albeit in its generalized form. At a high level, we exploit the homomorphic relationship between the additive group over Z_p and the group $\mathcal{G}$, which yields a natural way to map the linearity-preserving property of SPDZ over Z_p to its extension over $\mathcal{G}$. Informally speaking, for $h = g^x$ for some publicly available generator g of $\mathcal{G}$, let $[h]_{\mathcal{G}} := g^{[x]}$. Then, observe that the linearity-preservation property in $\mathcal{G}$ follows from the linearity-preservation property in Z_p, albeit implicitly in the exponent of the public group element g.

Concretely, any group element $g \in \mathcal{G}$ is shared as

$$[g]_{\mathcal{G}} = (\delta_{\mathcal{G}}, (g_1, \ldots, g_n), (\gamma_1(g), \ldots, \gamma_n(g))),$$

where for each $i \in [n]$, party P_i holds the tuple $(g_i, \gamma_i(x), \delta_{\mathcal{G}}) \in \mathcal{G} \times \mathcal{G} \times \mathcal{G}$, and where the following invariant holds:

$$g = \prod_{i \in [n]} g_i, \quad (g \cdot \delta_{\mathcal{G}})^{\alpha} = \prod_{i \in [n]} \gamma_i(g),$$

where α is the same global MAC-key as used in Tier-1.

Opening and MAC-Check in $\mathcal{G}$. The opening protocol for a secret-shared group element $[g]_{\mathcal{G}}$ is also analogous to the corresponding protocol for F_p where each party P_i does the following: (a) upon receiving a reconstructed value x', computes $\sigma_i = \gamma_i(g)/(g' \cdot \delta_{\mathcal{G}})^{\alpha_i}$, (b) broadcasts a commitment $\mathsf{Com}(\sigma_i)$ to all the other parties, and (c) opens the commitments $\{\mathsf{Com}(\sigma_j)\}$ received from $\{P_j\}_{j \neq i}$, computes $\mathsf{chk} = \prod_{j \in [n]} \sigma_j$, and aborts if $\mathsf{chk} \neq \mathrm{id}_{\mathcal{G}}$, where $\mathrm{id}_{\mathcal{G}}$ is the additive identity for the group $\mathcal{G}$. We can use a very similar argument as that in Tier-1 to prove that an adversary $\mathcal{A}$ cannot bypass this extended MAC-check protocol over $\mathcal{G}$, except with negligible probability.

Exponentiating a Public Element in $\mathcal{G}$. As mentioned in prior works [52], exponentiating a publicly available group element in $\mathcal{G}$ using a secret-shared value in Z_p is immediate; given a public group element g and a secret-sharing of x of the form

$$[x] = (\delta, (x_1, \ldots, x_n), (\gamma_1(x), \ldots, \gamma_n(x))),$$

one can easily compute a secret-sharing of $h = g^x$ as

$$[h]_{\mathcal{G}} = g^{[x]} := \left(g^{\delta}, (g^{x_1}, \ldots, g^{x_n}), \left(g^{\gamma_1(x)}, \ldots, g^{\gamma_n(x)}\right)\right).$$

Exponentiating a Secret-Shared Element in $\mathcal{G}$. In order to exponentiate a secret-shared group element $[g]_{\mathcal{G}}$ using a secret-shared value $[x]$, the parties use a protocol that naturally extends SPDZ's implementation of the functionality for multiplying secret-shared values (based on generating a random multiplicative triple and then using Beaver's re-randomization technique). Concretely, the parties follows the following steps:

- Generate $[a]$, $[b]$ and $[c]$ for $a, b \leftarrow Z_p$ and $c = a \cdot b$ using the triple-generation functionality in Tier-1
- Locally compute $[h_1]_{\mathcal{G}} = g^{[b]}$ and $[h_2]_{\mathcal{G}} = g^{[c]}$ using the exponentiation algorithm outlined above.
- Partially open the values $\epsilon = (x - a)$ and $h_3 = g/h_1$.
- Locally compute $[h_4]_{\mathcal{G}} = h_3^{[a]}$ (using the exponentiation algorithm outlined above) and $h_5 = h_3^{\epsilon}$.
- Locally compute $[h]_{\mathcal{G}} = [h_2]_{\mathcal{G}} \cdot \left([h_1]_{\mathcal{G}}\right)^{\epsilon} \cdot [h_4]_{\mathcal{G}} \cdot h_5$.

Note that the final local computation is allowed by the linearity-preserving property of the secret-sharing over $\mathcal{G}$; we omit the explicit details for simplicity.

$\mathcal{F}[\mathsf{Pair}]$

Pair-G1-P: On command $(\mathsf{pair}\mathcal{G}P, g_1, \mathsf{varid}_1, \mathsf{varid}_2)$ from all parties where $g_1 \in \mathcal{G}_1$, varid_1 is in the list of group $\mathcal{G}_2$ identifiers, and varid_2 is a fresh identifier in the list of group $\mathcal{G}_T$ identifiers, the functionality retrieves (varid_1, g_2) from the list of $\mathcal{G}_2$ identifiers and stores $(\mathsf{varid}_2, e(g_1, g_2))$, where e is the pairing function.

Pair-G2-P: On command $(\mathsf{pair}\mathcal{G}P, \mathsf{varid}_1, g_2, \mathsf{varid}_2)$ from all parties where varid_1 is in the list of group $\mathcal{G}_1$ identifiers, $g_2 \in \mathcal{G}_2$, and varid_2 is a fresh identifier in the list of group $\mathcal{G}_T$ identifiers, the functionality retrieves (varid_1, g_1) from the list of $\mathcal{G}_1$ identifiers and stores $(\mathsf{varid}_2, e(g_1, g_2))$, where e is the pairing function.

Pair-S: On command $(\mathsf{pair}S, \mathsf{varid}_1, \mathsf{varid}_2, \mathsf{varid}_3)$ from all parties where varid_1 is in the list of group $\mathcal{G}_1$ identifiers, varid_2 is in the list of group $\mathcal{G}_2$ identifiers, and varid_3 is a fresh identifier in the list of group $\mathcal{G}_T$ identifiers, the functionality retrieves (varid_1, g_1) from the list of $\mathcal{G}_1$ identifiers, (varid_2, g_2) from the list of $\mathcal{G}_2$ identifiers and stores $(\mathsf{varid}_3, e(g_1, g_2))$.

Fig. 6: Ideal functionality for MPC over the EC pairing operation with $\mathcal{G}_1$ and $\mathcal{G}_2$ as the input groups and $\mathcal{G}_T$ as the target group. We assume that $\mathcal{F}[\mathsf{Pair}]$ also includes all Tier-1 and Tier-2 sub-functionalities in $\mathcal{F}[F_p]$ and $\mathcal{F}[\mathcal{G}]$, but we avoid re-writing them for modularity.

Remark. We note here that while the aforementioned extension of SPDZ was proposed theoretically in prior works [30], [52], it was done specifically for EC groups (in particular, [52] is the only prior work to propose protocols for scalar-multiplying public/secret-shared EC points with secret-shared scalars, and their treatment is entirely specific to plain EC groups). As already mentioned, our generalized approach allows us to make this engine usable for pairing-friendly EC curves, since we can instantiate this engine not only for the source groups of an EC pairing (which are both elliptic curve groups), but also for the target group of the EC pairing, which is not an EC group but a multiplicative group over some extension field of F_p. As it turns out, this is an important building block that eventually allows us to support EC pairing operations in Tier-3 of our framework.

C. Tier-3: MPC over EC Pairings

We now build upon the infrastructure set up in Tier-1 and Tier-2 and design the MPC engine to support EC pairing operations. In particular, for a bilinear pairing $e : \mathcal{G}_1 \times \mathcal{G}_2 \to \mathcal{G}_T$, we start with Tier-2 instances for each of the groups $\mathcal{G}_1$, $\mathcal{G}_2$ and $\mathcal{G}_T$ (all of which satisfy linearity-preserving and support the operations outlined earlier), and realize the following three deterministic functionalities for EC pairings:

1) An EC pairing functionality that pairs a publicly available group element in $\mathcal{G}_1$ with a secret-shared group element in $\mathcal{G}_2$, i.e., given a public $g_1 \in \mathcal{G}_1$ and a secret-shared group element $[g_2]_{\mathcal{G}_2}$ for $g_2 \in \mathcal{G}_2$, outputs a secret-shared output $[g_T]_{\mathcal{G}_T}$ such that $g_T = e(g_1, g_2)$.
2) An EC pairing functionality that pairs a secret-shared group element in $\mathcal{G}_1$ with a publicly available group element in $\mathcal{G}_2$, i.e., given a secret-shared group element $[g_1]_{\mathcal{G}_1}$ for $g_1 \in \mathcal{G}_1$ and a public $g_2 \in \mathcal{G}_2$, produces a secret-shared output $[g_T]_{\mathcal{G}_T}$ such that $g_T = e(g_1, g_2)$.
3) An EC pairing functionality that pairs a secret-shared group element in $\mathcal{G}_1$ with a secret-shared group element in $\mathcal{G}_2$, i.e., given a secret-shared element $[g_1]_{\mathcal{G}_1}$ for $g_1 \in \mathcal{G}_1$ and a secret-shared group element $[g_2]_{\mathcal{G}_2}$ for $g_2 \in \mathcal{G}_2$, outputs a secret-shared output $[g_T]_{\mathcal{G}_T}$ such that $g_T = e(g_1, g_2)$.

We refer to $\mathcal{F}[\mathsf{Pair}]$ described in Figure 6 for a formal description of these functionalities. Once again, for malicious security, we would need each of the above functionalities to preserve the authentication components of the output appropriately. We note here that this functionality supports both symmetric and asymmetric pairings (in the symmetric case, we simply instantiate the framework with $\mathcal{G}_1 = \mathcal{G}_2 = \mathcal{G}$).

Tier-3 Extension of SPDZ. One of our technical contributions is an extension of the SPDZ framework to support MPC protocols realizing $\mathcal{F}[\mathsf{Pair}]$, which we describe here.

Pairing with One Secret-Shared Input. We begin by describing how to compute an EC pairing when one of the input group elements is secret-shared and the other input group element is public. We realize this by exploiting the bilinear property of the EC pairing. Recall that if $e : \mathcal{G}_1 \times \mathcal{G}_2 \to \mathcal{G}_T$ is a bilinear pairing, then for any $g_1, h_1 \in \mathcal{G}_1$ and any $g_2, h_2 \in \mathcal{G}_2$, we have

$$e(g_1 \cdot h_1, g_2) = e(g_1, g_2) \cdot e(h_1, g_2),$$
$$e(g_1, g_2 \cdot h_2) = e(g_1, g_2) \cdot e(g_1, h_2).$$

Now, observe that to pair a publicly available group element in $\mathcal{G}_1$ with a secret-shared group element in $\mathcal{G}_2$, each party can just locally compute

$$[h_T]_{\mathcal{G}_T} = e\left(h_1, [h_2]_{\mathcal{G}_2}\right),$$

and this yields a valid secret-sharing of pairing output h_T because of: (a) the bilinearity property of e as described above, and (b) the linearity-preservation property of the secret-sharing mechanism over $\mathcal{G}_2$. Pairing a publicly available group element in $\mathcal{G}_2$ with a secret-shared group element in $\mathcal{G}_1$ is analogously straightforward, wherein each party locally computes

$$[h_T]_{\mathcal{G}_T} = e\left([h_1]_{\mathcal{G}_1}, h_2\right).$$

Pairing with Two Secret-Shared Inputs. We now propose a protocol that allows the parties to pair a secret-shared group element $[h_1]_{\mathcal{G}_1}$ with a secret-shared group element $[h_2]_{\mathcal{G}_2}$, the parties follows the following steps. The protocol is inspired by SPDZ's implementation of the functionality for multiplying secret-shared values (based on generating a random multiplicative triple and then using Beaver's re-randomization technique), but needs to be carefully adapted to the setting of EC pairings. Concretely, in our proposed protocol, the parties proceed as follows:

- Generate $[a]$, $[b]$ and $[c]$ for $a, b \leftarrow Z_p$ and $c = a \cdot b$ using the triple-generation functionality in Tier-1.
- Locally compute

  $$[u_1]_{\mathcal{G}_1} = g_1^{[a]}, \quad [u_2]_{\mathcal{G}_2} = g_2^{[b]}, \quad [u_3]_{\mathcal{G}_1} = g_1^{[c]},$$

  using the exponentiation algorithm for public group elements in the Tier-2 MPC engine for $\mathcal{G}_1$ and $\mathcal{G}_2$.
- Partially open the values $h_3 = h_1/u_1 \in \mathcal{G}_1$ and $h_4 = h_2/u_2 \in \mathcal{G}_2$.
- Locally compute

  $$[v_1]_{\mathcal{G}_T} = e\left([u_3]_{\mathcal{G}_1}, g_2\right), \qquad [v_2]_{\mathcal{G}_T} = e\left(h_3, [u_2]_{\mathcal{G}_2}\right)$$
  $$[v_3]_{\mathcal{G}_T} = e\left([u_1]_{\mathcal{G}_1}, h_4\right), \qquad v_4 = e(h_3, h_4).$$

- Locally compute $[h_T]_{\mathcal{G}_T} = [v_1]_{\mathcal{G}_T} \cdot [v_2]_{\mathcal{G}_T} \cdot [v_3]_{\mathcal{G}_T} \cdot v_4$.

Note that the final local computation is allowed by the linearity-preserving property of the secret-sharing over $\mathcal{G}_T$; we omit the explicit details for simplicity. To prove correctness, it suffices to prove that $h_T = v_1 \cdot v_2 \cdot v_3 \cdot v_4$; correctness of the sharing again follows immediately from: (a) the bilinearity property of e described above, and (b) the linearity-preservation property of the secret-sharing mechanism over $\mathcal{G}_1$ and $\mathcal{G}_2$. Observe that

$$
\begin{aligned}
& v_1 \cdot v_2 \cdot v_3 \cdot v_4 \\
& \quad = e\left(u_3, g_2\right) \cdot e\left(h_3, u_2\right) \cdot e\left(u_1, h_4\right) \cdot e(h_3, h_4) \\
& \quad = e\left(g_1^c, g_2\right) \cdot e\left(h_1 \cdot g_1^{-a}, g_2^b\right) \cdot e\left(g_1^a, h_2 \cdot g_2^{-b}\right) \\
& \qquad\qquad \cdot e\left(h_1 \cdot g_1^{-a}, h_2 \cdot g_2^{-b}\right) \\
& \quad = e\left(g_1, g_2\right)^c \cdot e\left(h_1, g_2\right)^b \cdot e\left(g_1, g_2\right)^{-ab} \cdot e\left(g_1, g_2\right)^{-ab} \cdot e\left(g_1, h_2\right)^a \\
& \qquad \cdot e\left(h_1, h_2\right) \cdot e\left(h_1, g_2\right)^{-b} \cdot e\left(g_1, h_2\right)^{-a} \cdot e\left(g_1, g_2\right)^{ab} \\
& \quad = e(h_1, h_2) = h_T
\end{aligned}
$$

We highlight here that our solution uses the group operations and the pairing operations of the pairing-friendly EC group as a black-box. This enables us to use the state-of-the-art libraries such as RELIC [11] for implementing the pairing operations on top of the MP-SPDZ framework. To the best of our knowledge, this is first proposal and implementation of an MPC protocol that efficiently supports EC pairings, and is likely to have applications beyond PCI.

IV. PCI-Any-DC USING ECDSA SIGNATURE SCHEME

In this section, we describe a concrete instantiation of two-party PCI-Any-DC using the ECDSA signature scheme. We subsequently discuss how to extend this scheme to support PCI-Any and PCI-All.

Notations. Let the elliptic curve group $\mathcal{G}$ of prime order p be defined over a field F_p as a set of points $(x, y) \in F_p \times F_p$. Though the EC group $\mathcal{G}$ is an additive group of points over the elliptic curve, we will continue to use the multiplicative notation to ensure uniformity throughout the paper. Hence, we will denote point addition between two points Q_1 and Q_2 as $Q_1 \cdot Q_2$, and the scalar multiplication between a point Q and $x \in Z_p$ as Q^x. Let $Q \in \mathcal{G}$ be the generator of the group $\mathcal{G}$ (base point in standard EC parlance), and therefore we have $Q^p = \mathcal{O}$, where $\mathcal{O}$ is the point at infinity (the identity element). For any $Q' \in \mathcal{G}$, we use $[Q']_{\mathcal{G}}$ to denote the linearity preserving secret-sharing of Q'.

The ECDSA Signature Scheme. We briefly recall the key generation, signing, and verification equations for ECDSA.

$\mathsf{KeyGen}(\lambda)$: On input a security parameter λ, the key generation algorithm samples a private signing key $x \leftarrow [1, p-1]$, and computes the public verification key $Y := Q^x$. The algorithm outputs the pair (x, Y).

$\mathsf{Sign}(x, m)$: On input a signing key x and a message $m \in \{0,1\}^*$, the signing algorithm does the following: (i) samples a random $k \leftarrow [1, p-1]$, (ii) computes $R = (x, y) := Q^k$ (a random point on the curve), (iii) computes $r = x \mod p$ and $s = k^{-1}(H(m) + r \cdot x) \mod p$, where $\mathcal{H} : \{0,1\}^* \rightarrow [0, p-1]$ denotes a hash function, (iv) repeats (i)-(iii) until $r \neq 0$ and $s \neq 0$. The algorithm finally outputs the signature $\sigma = (r, s)$.

$\mathsf{Verify}(Y, \sigma, m)$: On input a verification key Y, a signature σ and a message m, the verification algorithm computes $u_1 = H(m) \cdot s^{-1} \mod p$, $u_2 = r \cdot s^{-1} \mod p$ and computes $R := (x', y') = Q^{u_1} \cdot Y^{u_2}$. The algorithm outputs 1 if $(x', y') \neq \mathcal{O}$ and $x' = r$, and outputs 0 otherwise.

Algorithm 1: PCI-Any-DC **using ECDSA**

Private inputs from P_1: $\mathsf{inp}_{1,1} = [(Y_{1,\ell}, s_{1,\ell}^{-1}, \mathsf{m}_{1,\ell})]_{\ell \in [1, N_1]}$
Each $Y_{1,\ell}$ is shared as $\left[Y_{1,\ell}\right]_{\mathcal{G}_2}$ using Input-G, and each $s_{1,\ell}^{-1}$ is shared as $\left[s_{1,\ell}^{-1}\right]$ using Input-F.
Public inputs from P_1: $\mathsf{inp}_{1,2} = [(r_{1,\ell}, R_{1,\ell}, \mathsf{m}_{1,\ell})]_{\ell \in [1, N_1]}$
Private inputs from P_2: $\mathsf{inp}_{2,1} = [(Y_{2,\ell}, s_{2,\ell}^{-1}, \mathsf{m}_{2,\ell})]_{\ell \in [1, N_2]}$
Each $Y_{2,\ell}$ is shared as $\left[Y_{2,\ell}\right]_{\mathcal{G}_2}$ using Input-G, and each $s_{2,\ell}^{-1}$ is shared as $\left[s_{2,\ell}^{-1}\right]$ using Input-F.
Public inputs from P_2: $\mathsf{inp}_{2,2} = [(r_{2,\ell}, R_{2,\ell}, \mathsf{m}_{2,\ell})]_{\ell \in [1, N_2]}$
P_1 validates each $R_{2,\ell} \neq \mathcal{O}$ and has-x coordinate $r_{2,\ell}$.
P_2 validates each $R_{1,\ell} \neq \mathcal{O}$ and has x-coordinate $r_{1,\ell}$.
▷ **Validate P_1's input signatures**
for $\ell := 1 \ldots N_1$ **do**
  $\left[u_{1,\ell}\right] := H\left(\mathsf{m}_{1,\ell}\right) \cdot \left[s_{1,\ell}^{-1}\right]$
  $\left[v_{1,\ell}\right] := r_{1,\ell} \cdot \left[s_{1,\ell}^{-1}\right]$
  $\left[C_\ell^1\right]_{\mathcal{G}} :=$ Exp-G-P($\left[u_{1,\ell}\right], Q$) $\cdot$ Exp-G-S($\left[v_{1,\ell}\right], \left[Y_{1,\ell}\right]_{\mathcal{G}_2}$)$/R_{1,\ell}$
▷ **Validate P_2's input signatures**
for $\ell' := 1 \ldots N_2$ **do**
  $\left[u_{2,\ell'}\right] := H\left(\mathsf{m}_{2,\ell'}\right) \cdot \left[s_{2,\ell'}^{-1}\right]$
  $\left[v_{2,\ell'}\right] := r_{2,\ell'} \cdot \left[s_{2,\ell'}^{-1}\right]$
  $\left[C_{\ell'}^2\right]_{\mathcal{G}} :=$ Exp-G-P($\left[u_{2,\ell'}\right], Q$) $\cdot$ Exp-G-S($\left[v_{2,\ell'}\right], \left[Y_{2,\ell'}\right]_{\mathcal{G}_2}$)$/R_{2,\ell'}$
▷ **Match certifier**
The parties agree on public random values $\mathsf{rnd}_1, \mathsf{rnd}_2 \leftarrow Z_p$.
for $\ell := 1 \ldots N_1$ **do**
  for $\ell' := 1 \ldots N_2$ **do**
    Generate secret-shared randomness $\left[\mathsf{rnd}_{\ell,\ell'}\right] \leftarrow$ Rand-F.
    $[C]_{\mathcal{G}} := \left[Y_{1,\ell}\right]_{\mathcal{G}} / \left[Y_{2,\ell'}\right]_{\mathcal{G}}$
    $[C']_{\mathcal{G}} := \left[C_\ell^1\right]_{\mathcal{G}} \cdot \left[C_{\ell'}^2\right]_{\mathcal{G}}^{\mathsf{rnd}_1} \cdot [C]_{\mathcal{G}}^{\mathsf{rnd}_2}$
    $\left[C''_{\ell,\ell'}\right]_{\mathcal{G}} :=$ Exp-G-S($\left[\mathsf{rnd}_{\ell,\ell'}\right], [C']_{\mathcal{G}}$)
    Output-G($\left[C''_{\ell,\ell'}\right]_{\mathcal{G}}$)
    If $C''_{\ell,\ell'} == \mathcal{O}$, then Output-G($\left[Y_{1,\ell}\right]_{\mathcal{G}}$)

Protocol overview. The starting point of our protocol is the generic maliciously secure protocol outlined in the introduction where we have the certificate validation and creation of the filtered sets of identities followed by the intersection of the sets from the two parties. We note here that we could have a single certifier issue multiple certificates on multiple different claims, or multiple certificates some of the same claims. However, we prescribe the parties to select only one certificate from a single certifier on one claim, i.e., there is a single (certificate, claim) pair for each certifier input to the protocol. We also expect an honest party to only input valid certificates on its set of public claims (although this is not a strict requirement for our protocol).

Optimizing Verify*:* Our main effort here is to reduce or obviate the non-algebraic operations in the Verify algorithm. In addition to the additions and multiplications, Verify requires an inverse operation in F_p and the extraction of the x-coordinate of an EC point from the point description (which is a trivial task to do in the plaintext world but not so inside an MPC). To do this, we make two observations. First, we note that

the unforgeability of the signature scheme is retained if s^{-1} is input instead of s; given a signature (r, s), it is trivial to compute (r, s^{-1}) and hence the unforgeability guarantees are equivalent for (r, s) and (r, s^{-1}). This way the inverse can be done outside the MPC and the parties can provide the corresponding s^{-1} as their secret inputs.

Second, in addition to r, we input the point $R = (r, y)$ by calculating the y-coordinate, and check that the signature verification procedure actually yields the point R (recall that the original ECDSA signature verification algorithm first reconstructs the point R and then extracts its x-coordinate r). If r and R were to be private inputs, the MPC algorithm would have to check that the r is the valid x-coordinate of R to prevent maliciously constructed inputs. We obviate this by making r and R public. Observe that, in the ECDSA signing algorithm, the point R is a uniformly random point in the group $\mathcal{G}$, thus R and its x-coordinate r are statistically independent of the corresponding public key. In other words, the public key is not revealed when r and R are provided, even if the universal set of public keys is available to the adversary. We also note that a malicious adversary cannot forge signatures by inputting an invalid point R' since, given the x-coordinate r and the public description of the elliptic curve group $\mathcal{G}$, one can efficiently compute the two possible EC points the form (r, y) in the group $\mathcal{G}$, and either of these would match the point R reconstructed by the verification algorithm if and only if the original signature (r, s) was valid. At this point, we can perform certificate verification inside MPC using the operations in Tier-2 of our proposed MPC engine.

Computing the intersection: We now perform the intersection of the sets of public keys by subtracting the corresponding elliptic curve points (dividing in the multiplicative notation) and checking if it opens to the identity element (point at infinity). It is important to hide the difference value if it is not the identity; otherwise we leak information about the public keys which are not part of the output set, which is not an allowed leakage according to our definition. So, we randomize the difference before opening while retaining the identity value. Another optimization in our protocol is that we store the information on the validity of the certificates in $\left[C_l^1\right]_{\mathcal{G}}$s and $\left[C_{l'}^2\right]_{\mathcal{G}}$s and open them along with the variable $[C]_{\mathcal{G}}$ storing the equality of public keys, as a random linear combination of three variables corresponding to the validity of P_1's certificate, validity of P_2's certificate and the equality of the public keys of the certifiers. This opens to the identity element if and only if all of the three requirements are satisfied.

The detailed description of our PCI-Any-DC protocol for ECDSA is provided in Algorithm 1. Here, each party inputs tuples of (identifier, certificate, claim) with the above discussed modifications as its private input, and the corresponding claim and (r, R) for each tuple as its public input. Note that the validation of P_1's certificates and P_2's certificates will be executed in parallel by the MPC algorithm. We describe the protocol in the $\mathcal{F}[\mathcal{G}]$-hybrid model, i.e., we assume that each sub-functionality in $\mathcal{F}[\mathcal{G}]$ has a secure instantiation. This allows us to define and prove the protocols in a modular way. A concrete instance of the protocol would use the SPDZ-based instantiation described in Section III to perform ECDSA signature validations while using all operations over the EC group $\mathcal{G}$ in a black-box way.

Correctness and Security. Correctness of the protocol follows immediately. We state the following theorem for the security of the protocol:

Theorem 1. *Our proposed* PCI-Any-DC *protocol for ECDSA signatures as described in Algorithm 1 securely emulates* $\mathcal{F}_{\mathsf{PCI}}(\mathsf{PCI\text{-}Any\text{-}DC})$ *(for the two-party setting).*

Proof Overview. We defer a detailed formal proof of this theorem to the full version of our paper [36]. We provide a brief proof overview here. Informally, we construct a PPT simulator $\mathcal{S}$ that simulates the view of a PPT environment $\mathcal{Z}$, such that that this simulated view is computationally indistinguishable from the real view of $\mathcal{Z}$. The crux of the proof is the following observation: prior to the output stage in Line 25 of Algorithm 1, the entire computation of the protocol is local. Thus, the environment's view, up to this point, will not leak whether inputs used by the honest player P_2 are dummy inputs or the ones that the environment actually provided (this guarantee follows immediately from the security of the underlying MPC framework in the $\mathcal{F}[\mathcal{G}]$ hybrid-model). Hence, the simulator $\mathcal{S}$ can *assume* entirely dummy inputs on behalf of the honest party P_2, and proceed with the simulation exactly as in the protocol.

To handle openings of the $C''_{\ell,\ell'}$ values (Line 25 of Algorithm 1), the simulator $\mathcal{S}$ invokes the ideal functionality $\mathcal{F}_{\mathsf{PCI}}(\mathsf{PCI\text{-}Any\text{-}DC})$ using the inputs of the corrupt party P_1 and obtains the output of the protocol $\mathsf{out}_{\mathsf{PCI\text{-}Any\text{-}DC}}(\mathsf{inp}_1, \mathsf{inp}_2)$. From the output, the $\mathcal{S}$ knows precisely which (ℓ, ℓ') tuples result in the opening of a $C''_{\ell,\ell'}$ value that is equal to $0_{\mathcal{G}}$, since this corresponds to an intersecting public key Y. Based on this information, $\mathcal{S}$ ensures consistent openings by suitably modifying the simulated share of $C''_{\ell,\ell'}$ corresponding to the honest party P_2 by exploiting the algebraic structure of the EC group and its knowledge of the MAC key α used in the simulation. Finally, to handle openings of $Y_{1,\ell}$ values (Line 26 of Algorithm 1), it suffices for the simulator $\mathcal{S}$ to proceed exactly as in the real protocol. This is because the public keys in the input of the corrupted party P_2 are available to the simulator $\mathcal{S}$ in the clear, and were shared by $\mathcal{S}$ exactly as in the real protocol. We refer to the full version of our paper [36] for a detailed description of the simulation strategy.

Extension to PCI-Any**.** One can naturally upgrade the above PCI-Any-DC protocol to a PCI-Any protocol that additionally guarantees privacy of the input claims for each party. More concretely, the claims would be secret-shared across the participating parties instead of being publicly available, and all operations on the input claims would have to be performed inside the MPC protocol. While the extension is conceptually simple, it incurs some additional costs. For instance, we can no longer directly use our proposed optimizations to reduce or obviate the non-algebraic operations in the Verify algorithm, and we would incur the additional cost of performing these operations *inside* the underlying MPC protocol. We would also incur the additional cost of hashing the claims *inside* the MPC protocol (since the claims would now be secret-shared as opposed to being publicly available). One could use an MPC-friendly family of hash functions [38], but this would be non-compliant with standardized implementations of ECDSA that typically do not use such hash function families. We leave it

as an interesting future direction to investigate optimization strategies that would allow performing the above operations efficiently (i.e., outside the MPC protocol) while ensuring privacy of the input claims *and* maintaining compliance with standardized ECDSA implementations.

Extension to PCI-All**.** The above PCI-Any-DC protocol can also be extended naturally to PCI-All by iterating through all the claims to validate the certificates on these claims by a specific certifier. To enable this, the private inputs will be ordered in a 2-D grid, where each row corresponds to the certificates by a certifier on all the claims in $\mathsf{inp}_{i,1}$, and the protocol needs to validate $|\mathsf{inp}_{i,1}|$ certificates per certifier inside the MPC protocol. The complexity grows with the number of claims which seems unavoidable since the ECDSA signatures cannot be aggregated across different claims. Therefore in the next section, we introduce an optimized PCI-All protocol using the BLS signature scheme [21] that only requires a single signature verification per certifier inside the MPC protocol.

Extension to Multi-Party PCI-Any-DC**.** Finally, we refer to the full version of our paper [36] for a discussion on how to extend the above PCI-Any-DC protocol (and its upgradation to PCI-Any) from the two-party to the multi-party setting.

V. PCI-All USING BLS SIGNATURE

This section provides a concrete instantiation of the PCI-All protocol using the BLS signature scheme [18], [19], [21]. At a high level, we use the aggregatable feature of BLS signatures over different claims to minimize the number of signature verifications inside the PCI-All protocol. Note however that BLS signature verification involves EC pairings, which we handle in a black-box way using **Tier-3** (Section III) of our proposed MPC engine.

Notations. Let $e : \mathcal{G}_1 \times \mathcal{G}_2 \rightarrow \mathcal{G}_T$ be a non-degenerate, efficiently computable bilinear pairing, where $\mathcal{G}_1, \mathcal{G}_2$ are elliptic curve groups and $\mathcal{G}_T$ is a multiplicative group, all of prime order p. Let Q_1 and Q_2 be generators of $\mathcal{G}_1$ and $\mathcal{G}_2$ respectively, and hence $g_T = e(Q_1, Q_2)$ is a generator of $\mathcal{G}_T$.

The BLS Signature Scheme. We briefly describe the key generation, signing and verification algorithms of the BLS signature scheme, followed by the algorithms for signature aggregation (over multiple messages signed under the same verification key) and the verification of aggregate signatures.

$\mathsf{KeyGen}(\lambda)$: On input a security parameter λ, the key generation algorithm samples a private signing key $x \leftarrow [1, p-1]$ and computes the public verification key as $Y = Q_2^x \in \mathcal{G}_2$. The algorithms outputs the key pair (x, Y).

$\mathsf{Sign}(x, m)$: On input a signing key x and message m, the signing algorithm first computes $M = H(m) \in \mathcal{G}_1$ where $H : \{0,1\}^* \rightarrow \mathcal{G}_1$. The algorithm then computes and outputs the signature $\sigma = M^x \in \mathcal{G}_1$.

$\mathsf{Verify}(Y, \sigma, m)$: On input a verification key Y, a signature σ and a message m, the verification algorithm outputs 1 if $e(\sigma, Q_2) = e(M, Y)$, and 0 otherwise.

Signature aggregation: On input signature-message pairs $\{\sigma_i, m_i\}_{i \in [1,N]}$, the signature aggregation algorithm produces an aggregated signature $\sigma_{(m_1,\ldots,m_N)} = \prod_{i \in [1,N]} \sigma_i$.

Aggregated signature verification: On input a verification key Y, an aggregated signature $\sigma_{(m_1,\ldots,m_N)}$ and a list/multiset of messages $(m_1, \ldots, m_N)$, the aggregated signature verification algorithm outputs 1 if $e(\sigma_{(m_1,\ldots,m_N)}, Q_2) = \prod_{i \in [1,N]} e(M_i, Y)$ where $M_i = H(m_i)$. The algorithm outputs 0 otherwise.

Remark. We note here that BLS signature aggregation is susceptible to a rogue public key attack when aggregating signatures on the same message under different verification keys. However, the attack is not applicable when aggregating signatures over multiple messages signed under the same public verification key, and hence does not impact the security of our proposed protocol.

Algorithm 2: PCI-All **using BLS**

P_1 has
$\mathsf{inp}_{1,1} = \left[\left(Y_{1,\ell_1}, \sigma_{1,\ell_1,\ell_2}, \mathsf{m}_{1,\ell_2}\right)\right]_{\ell_1 \in [1,N_{1,1}], \ell_2 \in [1,N_{1,2}]}$ and
$\mathsf{inp}_{1,2} = \{\mathsf{m}_{1,\ell_2}\}_{\ell_2 \in [1,N_{1,2}]}$
P_2 has
$\mathsf{inp}_{2,1} = \left[\left(Y_{2,\ell_1}, \sigma_{2,\ell_1,\ell_2}, \mathsf{m}_{2,\ell_2}\right)\right]_{\ell_1 \in [1,N_{2,1}], \ell_2 \in [1,N_{2,2}]}$ and
$\mathsf{inp}_{2,2} = \{\mathsf{m}_{2,\ell_2}\}_{\ell_2 \in [1,N_{2,2}]}$
Private inputs from P_1: the aggregated tuples and the set of preempted pairings
(i) $\overline{\mathsf{inp}}_{1,1} = \left[\left(Y_{1,\ell}, \overline{\sigma}_{1,\ell}, \overline{M}_1\right)\right]_{\ell \in [1,N_{1,1}]}$
(ii) $\{z_{1,\ell} = e(\overline{M}_2, Y_{1,\ell})\}_{\ell \in [1,N_{1,1}]}$
where $\overline{\sigma}_{i,\ell} = \prod_{\ell_2 \in [1,N_{i,2}]} \sigma_{i,\ell,\ell_2}$ and
$\overline{M}_i = \prod_{\ell \in [1,N_{i,2}]} H\left(\mathsf{m}_{i,\ell}\right)$. Note that each $Y_{1,\ell}$ is secret-shared as $\left[Y_{1,\ell}\right]_{\mathcal{G}_2}$, each $\overline{\sigma}_{1,\ell}$ is secret-shared as $\left[\overline{\sigma}_{1,\ell}\right]_{\mathcal{G}_1}$, and each $z_{1,\ell}$ is secret-shared as $\left[z_{1,\ell}\right]_{\mathcal{G}_T}$.
Public inputs from P_1: $\mathsf{inp}_{1,2}$.
Private inputs from P_2: the aggregated tuples and the set of preempted pairings
(i) $\overline{\mathsf{inp}}_{2,1} = \left[\left(Y_{2,\ell}, \overline{\sigma}_{2,\ell}, \overline{M}_2\right)\right]_{\ell \in [1,N_{2,1}]}$
(ii) $\{z_{2,\ell} = e(\overline{M}_1, Y_{2,\ell})\}_{\ell \in [1,N_{2,1}]}$
Note that each $Y_{2,\ell}$ is secret-shared as $\left[Y_{2,\ell}\right]_{\mathcal{G}_2}$, each $\overline{\sigma}_{2,\ell}$ is secret-shared as $\left[\overline{\sigma}_{2,\ell}\right]_{\mathcal{G}_1}$, and each $z_{2,\ell}$ is secret-shared as $\left[z_{2,\ell}\right]_{\mathcal{G}_T}$.
Public inputs from P_2: $\mathsf{inp}_{2,2}$.
for $\ell := 1 \ldots N_{1,1}$ **do**
  $\left[z'_{1,\ell}\right]_{\mathcal{G}_T} :=$ Pair-G2-P($\left[\overline{\sigma}_{1,\ell}\right]_{\mathcal{G}_1}, Q_2$)
for $\ell' := 1 \ldots N_{2,1}$ **do**
  $\left[z'_{2,\ell'}\right]_{\mathcal{G}_T} :=$ Pair-G2-P($\left[\overline{\sigma}_{2,\ell'}\right]_{\mathcal{G}_1}, Q_2$)
The parties agree on public random $r \leftarrow Z_p$.
for $\ell := 1 \ldots N_{1,1}$ **do**
  for $\ell' := 1 \ldots N_{2,1}$ **do**
    Generate secret-shared randomness $\left[r_{\ell,\ell'}\right] \leftarrow$ Rand-F.
    Each party locally computes:
    $\left[c_{\ell,\ell'}\right]_{\mathcal{G}_T} :=$
    $\left(\left[z_{1,\ell}\right]_{\mathcal{G}_T} / \left[z'_{2,\ell'}\right]_{\mathcal{G}_T}\right) \cdot \left(\left[z_{2,\ell'}\right]_{\mathcal{G}_T} / \left[z'_{1,\ell}\right]_{\mathcal{G}_T}\right)^r$
    $\left[c'_{\ell,\ell'}\right]_{\mathcal{G}_T} :=$ Exp-G-S$\left(\left[r_{\ell,\ell'}\right], \left[c_{\ell,\ell'}\right]_{\mathcal{G}_T}\right)$
    Output-G$\left(\left[c'_{\ell,\ell'}\right]_{\mathcal{G}_T}\right)$
    if $c'_{\ell,\ell'} == 1_T$ **then**
      Output-G$\left(\left[Y_{1,\ell}\right]_{\mathcal{G}_2}\right)$

Protocol overview. We follow the same generic approach as in our ECDSA-based protocol, with some optimizations to reduce BLS signature verifications inside the MPC protocol.

We note here that we could have a single certifier issue multiple certificates on the same claim for some of the claims. However, we prescribe the parties to select only one certificate from a single certifier on each claim, i.e., there is a single (certificate, claim) pair for each certifier per claim input to the protocol. We also expect an honest party to only input valid certificates on its set of public claims.

Reducing Claim Validation: As mentioned earlier, trivially extending the approach used in our ECDSA-based PCI-Any-DC protocol to design a PCI-All protocol would require iterating through all of the public claims, and validate the certificates on these claims by a specific certifier. This results in a claim validation complexity that grows with the number of claims, which is undesirable because the straightforward way of claim validation using BLS signatures would require computing two bilinear pairings inside the MPC protocol per validation, which is prohibitively expensive. Our main effort here is to reduce the number of pairing operations inside the MPC protocol as far as possible. To do this, we first use BLS signature aggregation over multiple claims signed under the same public verification key. Concretely, suppose that the private input $\mathsf{inp}_{i,1}$ for each (honest) party P_i is ordered in a 2-D grid of tuples of the form

$$\mathsf{inp}_{i,1} = [(Y_{i,\ell_1}, \sigma_{i,\ell_1,\ell_2}, \mathsf{m}_{i,\ell_2})]_{\ell_1 \in [1,N_{i,1}], \ell_2 \in [1,N_{i,2}]}$$

with $N_{i,1}$ certifiers and $N_{i,2}$ claims to be validated, where row-ℓ_1 contains certificates of the form σ_{i,ℓ_1,ℓ_2} on the claim m_{i,ℓ_2}, signed by the certifier associated with the verification key Y_{i,ℓ_1}. The party P_i performs some pre-processing to aggregate the certificates in each row using the BLS signature aggregation algorithm as:

$$\overline{\sigma}_{i,\ell_1} = \prod_{\ell_2 \in [1,N_{i,1}]} \sigma_{i,\ell_1,\ell_2}, \quad \overline{M}_i = \prod_{\ell_2 \in [1,N_{i,1}]} H\left(\mathsf{m}_{i,\ell_2}\right)$$

and uses an *aggregated private input* of the form

$$\overline{\mathsf{inp}}_{i,1} = \left[\left(Y_{i,\ell_1}, \overline{\sigma}_{i,\ell_1}, \overline{M}_i\right)\right]_{\ell_1 \in [1,N_{i,1}]}$$

for the MPC protocol. This now reduces the number of pairing computations inside the MPC protocol to two per certifier (required to verify each aggregated certificate); in particular, the complexity no longer grows with the number of public claims to be validated.

The next optimization involves further reducing the number of pairing computations inside the MPC to one per certifier. Note that we could avoid the pairing computation that requires pairing the public key with the aggregated claim-hash by having each party pre-compute this and directly input it to the MPC protocol. Note, however, that doing this naïvely would break the "unforgeability" guarantee of our protocol because a malicious party could simply input the pairing of a (potentially) invalid signature with the group generator Q_2 to trivially satisfy the verification check. To counter this, we exploit the uniqueness of BLS signatures for a given (key, claim) pair as follows: each party preempts the output of pairing its own verification keys with the aggregated claim-hashes of the other party (this is possible since the claims are public), which in the case of an intersecting certifier (i.e. when the verification keys are the same), is identical to the pairing of the aggregated public claim-hashes with the other party's verification key. This enables performing certificate verification for one party by using the preempted pairing values computed by the other party. This obviates the need for computing one of the pairings inside the MPC protocol (since the preempted pairing computation is done outside the MPC), while also preserving security of the end-to-end protocol.

Computing the intersection: In addition to certificate verification, the above step also enables computing the intersection of the identity sets between the two parties. In particular, we perform an equality check in $\mathcal{G}_T$ by simply dividing the corresponding group elements, and checking that the result opens to the identity element in $\mathcal{G}_T$. As in our ECDSA-based protocol, it is important to hide the output of this computation if it is not the identity; otherwise we leak information about the public keys which are not part of the output set, which is not an allowed leakage according to our definition. So, we randomize the difference before opening while retaining the identity value.

The detailed description of our PCI-All protocol for BLS signatures is provided in Algorithm 2. Here, each party P_i inputs tuples of (identifier, aggregated certificate, aggregated claim-hash) as its private input $\mathsf{inp}_{i,1}$, and the corresponding claims for each tuple as part of its public input $\mathsf{inp}_{i,2}$ (for the honest parties, $\mathsf{inp}_{i,2}$ is expected to be simply the set of public claims as in the definition of PCI-All in Section II). Each party also inputs the preempted pairing outputs as described earlier. We describe the protocol in the $(\mathcal{F}[\mathsf{Pair}])$-hybrid model, i.e., we assume that each sub-functionality in $\mathcal{F}[\mathsf{Pair}]$ has a secure instantiation. A concrete instance of the protocol would use the SPDZ-based instantiation described in Section III to perform BLS signature validations while using all operations over the EC groups $\mathcal{G}_1, \mathcal{G}_2$ and the target group $\mathcal{G}_T$ and the bilinear pairing e in a black-box way.

Correctness and Security. Correctness of the protocol follows immediately. We state the following theorem for the security of the protocol:

Theorem 2. *Our proposed* PCI-All *protocol for BLS signatures as described in Algorithm 2 securely emulates* $\mathcal{F}_{\mathsf{PCI}}(\mathsf{PCI\text{-}All})$ *(for the two-party setting).*

We defer a formal proof of this theorem to the full version of our paper [36].

Extension to Multi-Party PCI-All**.** We refer to the full version of our paper [36] for a discussion on how to extend the above PCI-All protocol from the two-party to the multi-party setting.

VI. Evaluation

This section details our implementation of the EC building blocks, the ECDSA-based PCI-Any-DC protocol, and the BLS-based PCI-All protocol. We independently benchmark the individual components of our protocols (including the protocols for EC operations) in a local server. We then evaluate the end-to-end performance of our PCI-Any-DC and PCI-All protocols in a LAN, an intra-continental WAN and an inter-continental WAN by spawning parties over three geographic regions across two continents.

A. Implementation Details

Our implementation builds on the MP-SPDZ [42] framework to support the EC operations, including pairing described

TABLE I: Throughput (operations per second) for Local EC Operations using RELIC and OpenSSL

	RELIC - Ed25519	OpenSSL - Secp256k1
Op-G	2,254,758	459,801
Exp-G-P	7,281	2,175

TABLE II: Throughput (operations per second) for Local EC Operations on Pairing-friendly Curves using RELIC

	BLS12-381	BLS12-446	BN-254	BLS12-638
Op-G : $\mathcal{G}_1$	1,079,688	834,877	687,906	435,223
Exp-G-P : $\mathcal{G}_1$	523,529	404,051	296,905	217,412
Op-G : $\mathcal{G}_2$	6,453	4,535	4,228	1,782
Exp-G-P : $\mathcal{G}_2$	3,684	2,683	1,990	1,019
Pair-G-P : $\mathcal{G}_1$, $\mathcal{G}_2$	960	689	508	307

in Section III. To the best of our knowledge, this is the first implementation of an MPC protocol that supports all the EC group operations as basic gates. In particular, we implement all the functionalities described in $\mathcal{F}[F_p]$, $\mathcal{F}[\mathcal{G}]$, and $\mathcal{F}[\mathsf{Pair}]$. The closest prior work [30] had implemented only two selected operations – Output-G and Exp-G-P. Our implementation of ECDSA PCI-Any-DC variant uses the standard OpenSSL (3.0) [4] library for EC operations. For the BLS PCI-All variant, we use the RELIC toolkit [11] to compute pairings and the EC operations on the corresponding groups. Both variants protect against malicious adversaries. As described earlier, our implementation builds on the SPDZ protocol with MASCOT [43] pre-processing. Analyzing the single-threaded CPU bottlenecks of the protocols, we have incorporated multi-threading to parallelize parts that individual parties locally execute without involving any communication (such as steps 9, 10, 14, 15, 22, & 23 in Algorithm 1, and 8, 10, & 16 in Algorithm 2). The source code of the implementation is made available here – https://github.com/ghoshbishakh/pci[3].

B. Component wise performance analysis

In this section we benchmark the individual operations of our proposed MPC framework for elliptic curve pairings. The different types of operations involved in the protocols can be categorized into (i) offline *pre-processing*, (ii) *input sharing*, (iii) *local operations* – performed by a party without any communication involved, e.g., Exp-G-P, (iv) *communication dependent operations* – which require inter-party communication, e.g., Exp-G-S, (v) *output* – which includes MAC-check. We perform experiments to analyze the performance of these different operations in terms of throughput (operations per second) and the impact of network latency on them. We separately compare the performance of local operations, followed by communication dependent operations including pre-processing, input sharing and output.

Platform Used. We used a workstation with dual Intel Xeon Gold 5118 2.30GHz CPUs, with 24 cores, and having 128 GB RAM. The system runs Ubuntu 18.04 operating system with Linux kernel version 4.15.

[3]We have also made some engineering contributions to the RELIC framework that makes progress in an easier integration of RELIC to other applications.

TABLE III: Throughput (operations per second) for Operations Requiring Communication

	RTT 1ms		RTT 100ms	
Pre-processing	967		267	
Input	261		245	
Output	457		363	
	Single Threaded	Multi Threaded	Single Threaded	Multi Threaded
Exp-G-S : $\mathcal{G}_1$	547	1,280	473	1,121
Exp-G-S : $\mathcal{G}_2$	277	554	257	554
Exp-G-S : $\mathcal{G}_T$	166	322	164	314
Pair-S: $\mathcal{G}_1$, $\mathcal{G}_2$	80	417	78	409

Local Operations. We start by benchmarking the local EC operations namely Op-G (point addition) and Exp-G-P (scalar multiplication with a point) separately for OpenSSL and RELIC. The throughput values (using a single thread) depicted in Table I make it evident that the performance of RELIC with Ed25519 [16] curve is significantly better than that of OpenSSL with Secp256k1 [5] curve. Nevertheless, we use OpenSSL for our ECDSA-based implementation of PCI-Any-DC since it one of the most widely-used libraries implementing the ECDSA algorithm [34], [47]. Following this, we evaluate the performance of EC operations on pairing-friendly curves with RELIC and carry out the experiments on four different curves, namely BLS12-381 [14], [55], [57], BN-254 [15], BLS12-446 [33], and BLS12-638 [57]. Table II summarizes the throughput for Op-G, Exp-G-P, and Pair-G-P for the above four curves. We observe that Op-G and Exp-G-P operations on $\mathcal{G}_2$ are much slower compared to that on $\mathcal{G}_1$, with Pair-G-P being the slowest operation by far. Among the curves benchmarked, BLS12-381 performs the best, and therefore we select this for the end-to-end experiments in Section VI-C.

Operations Requiring Communication. Moving to the more interesting benchmarks of the operations involving inter-party communication, namely *Pre-processing*, *Input*, *Output*, and *EC operations* Exp-G-S and Pair-S, we use two different setups – (a) a LAN setup with RTT between two parties being about 1ms, and (b) an emulated WAN setup with RTT of 100ms. In order to vary the link latency, we use the `tc` tool [6] to manipulate the loopback interface. Table III shows the throughput observed in the single threaded and multi-threaded implementation for Exp-G-S and Pair-S. We observe that *Pre-processing* slows down significantly with increasing latency, so is *Output* but to a lesser extent. The throughput values of Exp-G-S and Pair-S operations slightly drop with increasing latency but, even with a high RTT of 100ms, multithreading significantly increases the throughput, indicating that CPU is a major bottleneck for these operations. This validates the expectation since Exp-G-S and Pair-S are performed in batches and involve only one round of communication in which a batch of tuples are partially opened (see Sections III-B and III-C), thereby limiting the impact of network latency. However, if the batches are split (when a single batch becomes too large to handle), the impact of the communication latency will increase. Note that we perform this in a setup where bandwidth is sufficient enough to not be a bottleneck, and therefore, does not impact the benchmarks.

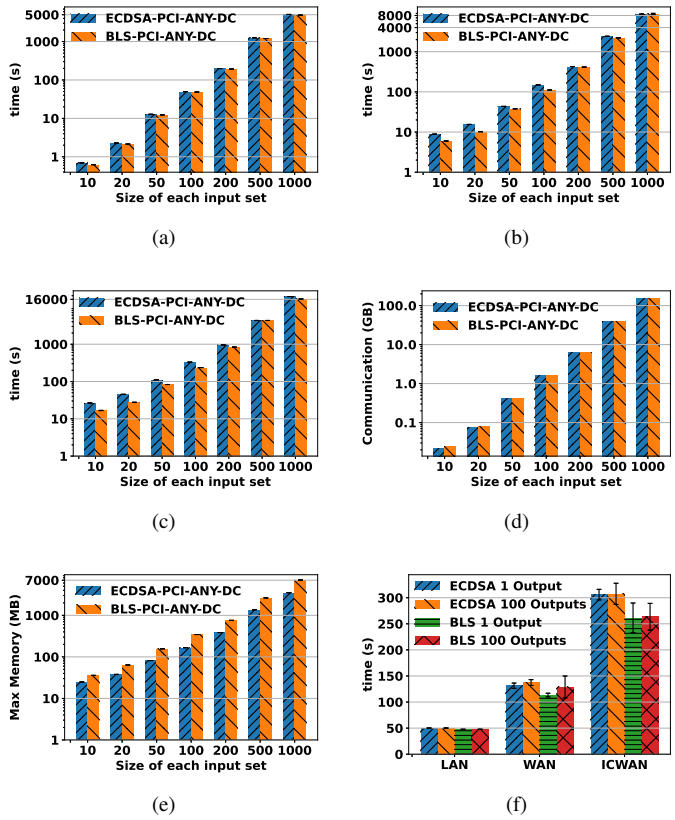

Fig. 7: (a), (b) and (c) depict latency (in logarithmic scale) of ECDSA PCI-Any-DC vs BLS PCI-Any-DC in LAN, WAN and ICWAN setups respectively. (d) and (e) represents total communication and maximum memory used respectively (in logarithmic scale) . (f) presents the latency with different output intersection sizes.

C. End-to-end performance analysis

In order to get real world performance metrics, we evaluate our implementations by placing the parties in the (a) same region – LAN, (b) different regions in the same continent – Continental WAN (WAN), and (c) different continents – Intercontinental WAN (ICWAN).

Platform Used. To gauge the practical performance of PCI on consumer hardware, we carried out the experiments on AWS EC2 `c6i.xlarge` virtual machine instances with only 4 vCPUs and 8 GB RAM. The instances were running the Ubuntu 22.04 operating system and were connected with a network having up to 12.5 Gbps bandwidth [1]. For the ICWAN setup, we use instances located in Asia (ap-south-1) and North America (us-east-1), with an RTT latency of about 186ms. For WAN, we use two instances in the USA, one in east coast (us-east-1), and another in the west coast (us-west-1) with an RTT latency of about 62ms. For the LAN setup, we spawned the two parties in two separate VMs in the same data center (ap-south-1). We also performed experiments on more powerful hardware (48 vCPUs, 96 GB RAM), the results of which are reported in the full version of our paper [36].

Overall Latency of PCI-Any-DC. We evaluate the end-to-end ECDSA and BLS-based PCI-Any-DC protocols, with each party's input set sizes varying from 10 to 1000. Here, the BLS PCI-Any-DC refers to the BLS PCI-All (Algorithm 2) with the parties using a single claim and its corresponding signature instead of the aggregated claim and signature. Figures Fig.7a, Fig.7b, and Fig.7c show the mean and standard deviations of the latency in LAN, WAN, and ICWAN setups, respectively, taken over multiple runs. The y-axis shows the time taken in seconds in a logarithmic scale. For the input sets of size 10 from each party, the mean time taken is about 0.69 seconds, 8.8 seconds, and 26.4 seconds for the ECDSA PCI-Any-DC protocol in LAN, WAN, and ICWAN, respectively. In such a setting, the BLS PCI-Any-DC protocol takes 0.62 seconds, 5.9 seconds, and 16.6 seconds respectively. This is better than the ECDSA variant, albeit by a small margin because the ECDSA protocol requires additional Exp-G-S operations in the signature validation steps (lines 11 and 16 of Algorithm 1), which is not required in the BLS variant. Exp-G-S operation requires communication and hence is significantly expensive as analyzed in detail in Appendix VI-B. For 1000 inputs, both ECDSA and BLS PCI-Any-DC takes less than 84 minutes, 149 minutes, and 316 minutes in LAN, WAN, and ICWAN, respectively. Notably, in practice, the size of the centralized trusted set of all CAs on the web is around 200 [2]; therefore, we expect the plausible set of certifiers for a party to be less than 200. Here the number of certifiers do not imply the global set of all possible certifiers, instead it is the number of certifiers that have issued certificates for a given claim to a user. For 200 inputs, both ECDSA and BLS PCI-Any-DC takes less than 3.5 minutes, 7 minutes, and 15 minutes in LAN, WAN, and ICWAN, respectively. This is improved further by using more powerful hardware, which we report in the full version of our paper [36].

Communication and Memory Overhead of PCI-Any-DC. We observe that the volume of data communication across parties is deterministic and is defined by the size of their input sets as expected. Hence, there are no variations across the different runs and across LAN, WAN, and ICWAN. We report the communication bandwidth required for different input sizes in Fig.7d. With input size of 10 from each party, the total volume of data communicated is 22 MB for ECDSA and 25 MB with BLS PCI-Any-DC. With input sizes of 1000, the total communication goes up to 152.8 GB and 153.4 GB for ECDSA and BLS PCI-Any-DC, respectively. Unlike data communication overhead where ECDSA and BLS variants are close, the memory consumption of BLS is consistently higher as depicted in Fig.7e. For 1000 inputs, ECDSA PCI-Any-DC requires around 3.4 GB memory (maximum usage during the runtime), whereas the BLS variant uses around 6.8 GB.

Latency of PCI-Any-DC with varying output size. We evaluate the impact of varying overlap in the input certifier sets of the parties implying varying size of output intersection set. Fig.7f represents the end-to-end latency of both ECDSA and BLS PCI-Any-DC while keeping the number of input from each party constant at 100, and varying the output size from 1 to 100. We observe that compared to the output size 1, the end-to-end latency for 100 outputs is higher by a very small margin on an average in all the settings, namely, LAN, WAN, and ICWAN. This is because of the differences in the number of outputs from the protocol that has to be opened (line 26 of Algorithm. 1, and line 20 of Algorithm. 2). We note, however, that no additional information is leaked outside what is permitted by the definition of PCI-Any-DC (Section II) from the difference in the latency, since the intersection set

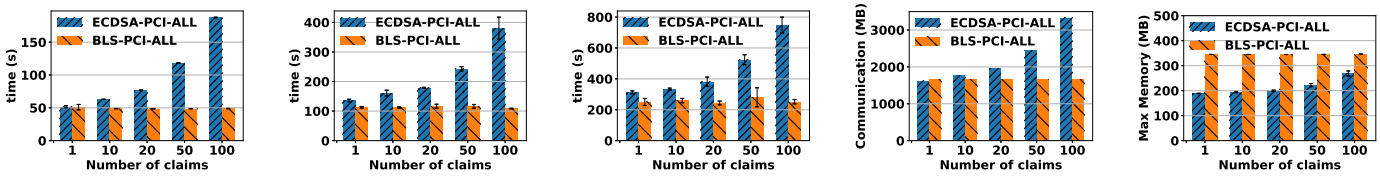

(a) (b) (c) (d) (e)

Fig. 8: (a), (b) and (c) depict latency of ECDSA PCI-All vs BLS PCI-All with 100 certifiers and 1 to 100 claims as input from each party in (a) LAN (b) Continental WAN (c) Inter-continental WAN setups respectively. (d) and (e) presents the total data communicated and maximum memory consumption of PCI-All respectively.

is already known to the parties one step prior to this opening phase (line 25 of Algorithm. 1, and line 18 of Algorithm. 2).

Comparing Latency of BLS PCI-All and ECDSA PCI-All. In order to evaluate the gains of using BLS signature aggregation for PCI-All over the ECDSA implementation, we use a (somewhat artificial) construction of ECDSA-based PCI-All which iterates through all the claims to validate the certificates on them (see Section IV). We evaluate the end-to-end latency by keeping the input set size of each party constant at 100, and increasing the number of claims from 1 to 100. The results in Fig.8a, Fig.8b, and Fig.8c depict the mean and the standard deviation of the overall latency in LAN, WAN and ICWAN setups, respectively, taken over multiple runs. While the BLS PCI-All consistently takes about 50 seconds, 115 seconds and 250 seconds for any number of claims (from 1 to 100) in LAN, WAN, and ICWAN setups, respectively, the time taken by ECDSA PCI-All gradually increases with the increase in the number of claims. ECDSA PCI-All takes on an average 188 seconds, 380 seconds, and 748 seconds for 100 claims in LAN, WAN and ICWAN, respectively. This clearly highlights the gains of using BLS construction of PCI-All.

Communication and Memory Overhead of PCI-All. The volume of data communicated between the two parties for the above scenario is depicted in Fig.8d. With increasing number of claims, the communication overhead increases for ECDSA PCI-All, whereas it stays constant for BLS PCI-All which is the expected outcome. For 100 claims, the volume of data communicated by ECDSA PCI-All is 3333 MB, and by BLS PCI-All it is 1658 MB. Memory consumption of ECDSA PCI-All also increases with the increasing number of claims as represented by Fig.8e. For 100 certifiers, with 100 claims for each party, the memory usage by ECDSA PCI-All is about 268 MB, and the same by the BLS variant is 345 MB. Overall, the memory consumption overhead of the BLS implementation is more than the ECDSA implementation for up to a reasonable number of claims such as 100.

VII. Future Directions

Our work gives rise to many interesting open questions. We leave it open to study PCI in the setting where claims are private, as well as to define and realize variants of PCI that outputs a priority list of certifiers. Designing PCI protocols supporting other signature schemes, including quantum-safe schemes, is another challenging direction of research.Our MPC framework over EC pairings can plausibly be leveraged for building MPC-based PCI supporting other EC-based signature schemes. While our PCI constructions based on ECDSA and BLS cannot be be immediately/trivially extended to other signature schemes, we expect that carefully designed and specifically optimized PCI constructions supporting other EC-based signature schemes can plausibly be realized by using our proposed MPC framework over EC pairings as a building block.

Acknowledgments

We thank the anonymous reviewers of NDSS 2023 for their valuable feedback, comments and suggestions.

References

[1] "Amazon ec2 c6i instances," (Last accessed: August 23, 2022). [Online]. Available: https://aws.amazon.com/ec2/instance-types/c6i/

[2] "Ca certificates in firefox," (Last accessed: August 23, 2022). [Online]. Available: https://ccadb-public.secure.force.com/mozilla/CACertificatesInFirefoxReport

[3] "Hyperledger indy," (Last accessed: May 13, 2022). [Online]. Available: https://www.hyperledger.org/use/hyperledger-indy

[4] "Openssl," (Last accessed: May 13, 2022). [Online]. Available: https://www.openssl.org/

[5] "Sec 2: Recommended elliptic curve domain parameters," (Last accessed: May 13, 2022). [Online]. Available: https://www.secg.org/sec2-v2.pdf

[6] "tc - show / manipulate traffic control settings," (Last accessed: May 13, 2022). [Online]. Available: https://man7.org/linux/man-pages/man8/tc.8.html

[7] "Tradelens," (Last accessed: August 29, 2022). [Online]. Available: https://www.tradelens.com/

[8] "Did specification registries," 2020, (Last accessed: May 13, 2022). [Online]. Available: https://www.w3.org/TR/did-spec-registries

[9] E. Abebe, D. Behl, C. Govindarajan, Y. Hu, D. Karunamoorthy, P. Novotny, V. Pandit, V. Ramakrishna, and C. Vecchiola, "Enabling enterprise blockchain interoperability with trusted data transfer (industry track)," in *Proceedings of the 20th International Middleware Conference Industrial Track*, 2019, pp. 29–35.

[10] X. ANSI, "62: public key cryptography for the financial services industry: the elliptic curve digital signature algorithm (ecdsa)," *Am. Nat'l Standards Inst*, 1999.

[11] D. F. Aranha, C. P. L. Gouvêa, T. Markmann, R. S. Wahby, and K. Liao, "RELIC is an Efficient LIbrary for Cryptography," https://github.com/relic-toolkit/relic.

[12] G. Ateniese, J. Kirsch, and M. Blanton, "Secret handshakes with dynamic and fuzzy matching." in *NDSS*, vol. 7, no. 24, 2007, pp. 43–54.

[13] D. Balfanz, G. Durfee, N. Shankar, D. Smetters, J. Staddon, and H.-C. Wong, "Secret handshakes from pairing-based key agreements," in *2003 Symposium on Security and Privacy, 2003.* IEEE, 2003, pp. 180–196.

[14] P. S. Barreto, B. Lynn, and M. Scott, "Constructing elliptic curves with prescribed embedding degrees," in *International conference on security in communication networks*. Springer, 2002, pp. 257–267.

[15] P. S. Barreto and M. Naehrig, "Pairing-friendly elliptic curves of prime order," in *International workshop on selected areas in cryptography*. Springer, 2005, pp. 319–331.

[16] D. J. Bernstein, N. Duif, T. Lange, P. Schwabe, and B.-Y. Yang, "High-speed high-security signatures," *Journal of cryptographic engineering*, vol. 2, no. 2, pp. 77–89, 2012.

[17] J. Bobolz, F. Eidens, S. Krenn, S. Ramacher, and K. Samelin, "Issuer-hiding attribute-based credentials," in *International Conference on Cryptology and Network Security*. Springer, 2021, pp. 158–178.

[18] D. Boneh, M. Drijvers, and G. Neven, "Compact multi-signatures for smaller blockchains," in *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 2018, pp. 435–464.

[19] D. Boneh, C. Gentry, B. Lynn, and H. Shacham, "Aggregate and verifiably encrypted signatures from bilinear maps," in *International conference on the theory and applications of cryptographic techniques*. Springer, 2003, pp. 416–432.

[20] D. Boneh, S. Gorbunov, R. S. Wahby, H. Wee, and Z. Zhang, "BLS Signatures," Internet Engineering Task Force, Internet-Draft draft-irtf-cfrg-bls-signature-04, Sep. 2020, work in Progress. [Online]. Available: https://datatracker.ietf.org/doc/html/draft-irtf-cfrg-bls-signature-04

[21] D. Boneh, B. Lynn, and H. Shacham, "Short signatures from the weil pairing," in *International conference on the theory and application of cryptology and information security*. Springer, 2001, pp. 514–532.

[22] D. Bosk, D. Frey, M. Gestin, and G. Piolle, "Hidden issuer anonymous credential," *Proceedings on Privacy Enhancing Technologies*, vol. 4, pp. 571–607, 2022.

[23] J. Camenisch and G. M. Zaverucha, "Private intersection of certified sets," in *International Conference on Financial Cryptography and Data Security*. Springer, 2009, pp. 108–127.

[24] R. Canetti, A. Cohen, and Y. Lindell, "A simpler variant of universally composable security for standard multiparty computation," in *Annual Cryptology Conference*. Springer, 2015, pp. 3–22.

[25] M. Chase and P. Miao, "Private set intersection in the internet setting from lightweight oblivious PRF," in *Advances in Cryptology - CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020*. Springer, 2020.

[26] H. Chen, K. Laine, and P. Rindal, "Fast private set intersection from homomorphic encryption," in *ACM CCS*, 2017, pp. 1243–1255.

[27] E. D. Cristofaro and G. Tsudik, "Practical private set intersection protocols with linear complexity," in *Financial Cryptography and Data Security, 14th International Conference, FC 2010*. Springer, 2010.

[28] K. Czajkowski, I. Foster, C. Kesselman, V. Sander, and S. Tuecke, "Snap: A protocol for negotiating service level agreements and coordinating resource management in distributed systems," in *Workshop on Job Scheduling Strategies for Parallel Processing*. Springer, 2002, pp. 153–183.

[29] D. Dachman-Soled, T. Malkin, M. Raykova, and M. Yung, "Efficient robust private set intersection," in *International Conference on Applied Cryptography and Network Security*. Springer, 2009, pp. 125–142.

[30] A. Dalskov, C. Orlandi, M. Keller, K. Shrishak, and H. Shulman, "Securing dnssec keys via threshold ecdsa from generic mpc," in *European Symposium on Research in Computer Security*. Springer, 2020, pp. 654–673.

[31] I. Damgård, M. Keller, E. Larraia, V. Pastro, P. Scholl, and N. P. Smart, "Practical covertly secure mpc for dishonest majority–or: breaking the spdz limits," in *European Symposium on Research in Computer Security*. Springer, 2013, pp. 1–18.

[32] I. Damgård, V. Pastro, N. Smart, and S. Zakarias, "Multiparty computation from somewhat homomorphic encryption," in *Annual Cryptology Conference*. Springer, 2012, pp. 643–662.

[33] A. de la Piedra, M. Venema, and G. Alpár, "ABE Squared: Accurately benchmarking efficiency of attribute-based encryption," *Cryptology ePrint Archive*, 2022.

[34] Z. Durumeric, F. Li, J. Kasten, J. Amann, J. Beekman, M. Payer, N. Weaver, D. Adrian, V. Paxson, M. Bailey *et al.*, "The matter of heartbleed," in *Proceedings of the 2014 conference on internet measurement conference*, 2014, pp. 475–488.

[35] M. J. Freedman, K. Nissim, and B. Pinkas, "Efficient private matching and set intersection," in *International conference on the theory and applications of cryptographic techniques*. Springer, 2004, pp. 1–19.

[36] B. C. Ghosh, S. Patranabis, D. Vinayagamurthy, V. Ramakrishna, K. Narayanam, and S. Chakraborty, "Private certifier intersection (full version)," Cryptology ePrint Archive, Paper 2022/1302, 2022, https://eprint.iacr.org/2022/1302. [Online]. Available: https://eprint.iacr.org/2022/1302

[37] B. C. Ghosh, D. Vinayagamurthy, V. Ramakrishna, K. Narayanam, and S. Chakraborty, "Privacy-preserving negotiation of common trust anchors across blockchain networks (short paper)," in *2022 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, 2022.

[38] L. Grassi, C. Rechberger, D. Rotaru, P. Scholl, and N. P. Smart, "Mpc-friendly symmetric key primitives," in *ACM CCS 2016*. ACM, 2016, pp. 430–443.

[39] P. E. Hoffman and W. Wijngaards, "Elliptic Curve Digital Signature Algorithm (DSA) for DNSSEC," RFC 6605, Apr. 2012. [Online]. Available: https://www.rfc-editor.org/info/rfc6605

[40] Y. Huang, D. Evans, and J. Katz, "Private set intersection: Are garbled circuits better than custom protocols?" in *NDSS*, 2012.

[41] D. Johnson, A. Menezes, and S. Vanstone, "The elliptic curve digital signature algorithm (ecdsa)," *International journal of information security*, vol. 1, no. 1, pp. 36–63, 2001.

[42] M. Keller, "MP-SPDZ: A versatile framework for multi-party computation," in *ACM CCS*, 2020.

[43] M. Keller, E. Orsini, and P. Scholl, "Mascot: faster malicious arithmetic secure computation with oblivious transfer," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 2016, pp. 830–842.

[44] L. Kissner and D. Song, "Privacy-preserving set operations," in *Annual International Cryptology Conference*. Springer, 2005, pp. 241–257.

[45] B. Lynn, "Pbc library-pairing-based cryptography," *http://crypto. stanford. edu/pbc/*.

[46] B. Moeller, N. Bolyard, V. Gupta, S. Blake-Wilson, and C. Hawk, "Elliptic Curve Cryptography (ECC) Cipher Suites for Transport Layer Security (TLS)," RFC 4492, May 2006. [Online]. Available: https://www.rfc-editor.org/info/rfc4492

[47] M. Nemec, D. Klinec, P. Svenda, P. Sekan, and V. Matyas, "Measuring popularity of cryptographic libraries in internet-wide scans," in *Proceedings of the 33rd Annual Computer Security Applications Conference*, 2017, pp. 162–175.

[48] B. Pinkas, T. Schneider, G. Segev, and M. Zohner, "Phasing: Private set intersection using permutation-based hashing," in {*USENIX*} *Security*, 2015, pp. 515–530.

[49] B. Pinkas, T. Schneider, and M. Zohner, "Faster private set intersection based on {OT} extension," in *23rd {USENIX} Security Symposium ({USENIX} Security 14)*, 2014, pp. 797–812.

[50] D. Reed, M. Sporny, D. Longley, C. Allen, R. Grant, M. Sabadello, and J. Holt, "Decentralized identifiers (dids) v1.0," 2020, (Last accessed: May 13, 2022). [Online]. Available: https://w3c.github.io/did-core/

[51] P. Rindal and M. Rosulek, "Malicious-secure private set intersection via dual execution," in *ACM CCS*, 2017, pp. 1229–1242.

[52] N. P. Smart and Y. Talibi Alaoui, "Distributing any elliptic curve based protocol," in *IMA International Conference on Cryptography and Coding*. Springer, 2019, pp. 342–366.

[53] M. Sporny, D. Longley, and D. Chadwick, "Verifiable credentials data model v1.1," 2021, (Last accessed: May 13, 2022). [Online]. Available: https://w3c.github.io/vc-data-model/

[54] A. Tobin and D. Reed, "The inevitable rise of self-sovereign identity," *The Sovrin Foundation*, vol. 29, no. 2016, 2016.

[55] R. S. Wahby and D. Boneh, "Fast and simple constant-time hashing to the BLS12-381 elliptic curve," *Cryptology ePrint Archive*, 2019.

[56] M. Winslett, T. Yu, K. E. Seamons, A. Hess, J. Jacobson, R. Jarvis, B. Smith, and L. Yu, "Negotiating trust in the web," *IEEE Internet Computing*, vol. 6, no. 6, pp. 30–37, 2002.

[57] S. Yonezawa, T. Kobayashi, and T. Saito, "Pairing-friendly curves," *Network Working Group. Internet-Draft. January*, 2019.