



Post-quantum EPID Signatures from Symmetric Primitives

Dan Boneh, Saba Eskandarian^(✉), and Ben Fisch

Stanford University, Stanford, USA
{dabo,saba,bfisch}@cs.stanford.edu

Abstract. EPID signatures are used extensively in real-world systems for hardware enclave attestation. As such, there is a strong interest in making these schemes post-quantum secure. In this paper we initiate the study of EPID signature schemes built only from symmetric primitives, such as hash functions and PRFs. We present two constructions in the random oracle model. The first is a scheme satisfying the EPID signature syntax and security definitions needed for private hardware attestation used in Intel’s SGX. The second achieves significantly shorter signatures for many applications, including the use case of remote hardware attestation. While our EPID signatures for attestation are longer than standard post-quantum signatures, they are short enough for applications where the data being signed is large, such as analytics on large private data sets, or streaming media to a trusted display. We evaluate several instantiations of our schemes so that the costs and benefits of these constructions are clear. Along the way we also give improvements to the zero-knowledge Merkle inclusion proofs of Derler et al. (2017).

1 Introduction

Enhanced Privacy ID, or EPID, signatures allow members of a group to anonymously sign messages on behalf of the group, with the added property that a group manager can revoke the credentials of a misbehaving or compromised group member [15, 36].

In recent years, EPID signatures have become an important privacy mechanism in real-world systems, most prominently in trusted hardware attestation such as Intel’s SGX. Attestation is a process by which a hardware enclave running on a client device proves the authenticity of its execution environment to a remote party. EPID lets the client device attest without revealing its identity to the remote party. However, EPID signatures used today are not post-quantum secure [15]. An adversary with a quantum computer could subvert the attestation process and break a hardware enclave’s security in the worst possible way.

In light of the above, there is strong interest in developing post-quantum secure EPID signatures. One way to do so is to construct an EPID signature scheme using only symmetric primitives, which are believed to be post-quantum secure. This is analogous to constructing a standard signature scheme from hash

functions [9, 16, 20, 43, 44] to obtain a signature scheme whose post-quantum security is virtually assured.

Can we build efficient and secure EPID signatures from symmetric primitives? Bellare et al. [6] give a generic construction of a group signature [21], a related primitive, from a standard signature scheme, public-key encryption, and a non-interactive zero-knowledge (NIZK) proof. In this generic construction, the group manager adds a member to the group by signing that member’s public key. The member can then sign messages anonymously by first using the private key to sign the message, and then computing a NIZK proof of knowledge of both this signature and the group manager’s signature on the corresponding public key. This NIZK proof is the member’s group signature. With some work, their framework can be adapted to support the EPID group signature definition of Brickell and Li [15] and to only use symmetric primitives. The NIZK can be built from the “MPC in the Head” technique of Ishai et al. [3, 30, 35] using random oracles, and the standard signature scheme can also be built from one-way functions and collision-resistant hashing [9, 20, 31, 43]. Camenisch and Groth [17] give such a scheme from one-way functions and NIZKs. However, without careful optimization, this generic approach leads to very inefficient signatures due to the need for NIZK proofs on complex circuits (the proof size and prover time of these NIZKs is proportional to the number of multiplication gates in the arithmetic circuit representing the statement).

1.1 Our Contributions

We construct an EPID signature scheme from symmetric primitives, and take a significant step towards reducing the signature size.

Towards this goal, we build two signature schemes. Our first construction greatly reduces the size of the NIZK statement in the signature by using PRFs instead of signatures wherever possible. In particular, we are able to replace the inner group member’s signature in the generic approach with a PRF evaluation. Our construction does not treat the given primitives as a black-box and performs best when instantiated with NIZK-friendly PRFs and CRHFs. In particular, we evaluate the scheme using the LowMC cipher [2], also including a comparison to AES to show the benefit of choosing the right instantiations for our primitives.

Next, we show how to significantly improve our EPID signature by adapting it to the specific real-world use case where signature verification requires an interaction with the group manager to ensure that the signer has not been revoked. We take advantage of this structure to dramatically reduce the signature size by moving many heavy verification steps outside of the NIZK, without compromising anonymity or affecting security. This significantly shrinks the signature size over the first construction.

Along the way, we develop a technique for proving membership in a Merkle tree, without revealing the leaf location, using a *third* preimage resistant hash function (Sect. 5.4). This also provides an improvement to the recent post-quantum accumulators of Derler et al. [24].

Performance and Use Cases. In Sect. 5 we discuss options for instantiating our schemes, and measure the sizes of the resulting signatures under different security assumptions. For the circuit sizes needed inside NIZKs in our construction, ZKB++ [20] provides the most efficient proofs. We report sizes for both the Random Oracle and Quantum Random Oracle models [11] (using the Fiat-Shamir [27] and Unruh [48] transforms, respectively), and find that our second signature scheme, designed for attestation, can support groups of over a million members with 3.45 MB signatures at 128-bit post-quantum security. While these signatures are not short, it is important to keep in mind that several megabytes of traffic for attestation is quite acceptable for many applications of trusted hardware, especially where the data transfer needs of the higher-level application dwarf the size of the attestation.

One example is the case of analytics over large private data sets, an area of heavy investment, both in terms of research and financial resources [29, 51]. In this setting, nodes in a distributed network (or the server in a client-server setting) provide a single remote attestation and then exchange a great deal of data. As the quantity of data transferred exceeds millions of database records, the size of the initial attestation ceases to present a major bottleneck.

The case of digital rights management (DRM), for which hardware enclaves such as Intel SGX seem particularly well-suited [22], is another setting where the size of our signatures are acceptable. Consider the common situation where a content provider wishes to stream a movie (easily a few gigabytes in size) to a subscriber while preventing redistribution or unauthorized viewing of copyrighted content [50, 52]. The few additional megabytes of an attestation do not matter next to a film or television series several hundred times its size.

1.2 Additional Related Work

Trusted Hardware and Attestation. Hardware enclaves, particularly Intel’s SGX [22], have recently been used for a variety of security applications [28, 45]. One of the primary cryptographic components of SGX is its use of direct anonymous attestation, a primitive introduced by Brickell et al. [14]. The EPID attestation mechanism currently in use by SGX, is due to Brickell et al. [15, 36].

Group Signatures. Anonymous attestation and EPID signatures bear a great deal of similarity to group signatures. Group signatures [21] allow members of a group to anonymously produce signatures on behalf of the group, with the added restriction that a group manager has the power to police the behavior of members, e.g. by revoking their group credentials or stripping their anonymity. The most frequently used definitions of group signatures are described by Bellare et al. [6, 7]. Subsequent work on group signatures has led to various schemes, for example, those of Lysyanskaya and Camenisch [18, 19], Boneh et al. [10, 12], and a scheme of Groth [34]. These constructions are not post-quantum secure.

Post-quantum Signatures and Proofs. Lattice-based cryptography is a popular candidate for post-quantum security. Lattice group signatures were introduced by Gordon et al. [33] and extended in several subsequent works [39–42].

The resulting group signatures are shorter than the ones developed here, but rely on qualitatively stronger post-quantum assumptions.

Another set of post-quantum tools come from the “MPC in the Head” technique [35] for zero-knowledge proofs. This idea has been extended by ZKBoo [30], ZKB++ [20], and Ligerio [3]. In particular, Chase et al. use ZKB++ to construct two post-quantum signature schemes Fish and Picnic [20]. The recent development of zk-STARKS [8] opens another avenue to post-quantum zero-knowledge proofs. In concurrent work, El Bansarkhani and Misoczki [4] describe a stateful group signature scheme based on hash functions. Their work features small signature sizes but large keys. Derler et al. [24] and Katz et al. [37] also concurrently study post-quantum group and ring signatures from symmetric primitives.

2 Preliminaries

Notation. Let $x \leftarrow F(y)$ denote the assignment of the output of $F(y)$ to x , and let $x \stackrel{\mathbb{R}}{\leftarrow} S$ denote assignment to x of a uniformly random element sampled from set S . We use λ to refer to a security parameter and sometimes omit it if its presence is implicit. The notation $[k]$ represents the set of integers $1, 2, \dots, k$, and \emptyset denotes the empty set. We use \mathcal{A}^H to denote that \mathcal{A} has oracle access to some function H . A function $\text{negl}(x)$ is *negligible* if for all $c > 0$, there is an x_0 such that for any $x > x_0$, $\text{negl}(x) < \frac{1}{x^c}$. We omit x if the parameter is implicit. We use $f(x) \approx g(x)$ to mean that for two functions f, g , $|f(x) - g(x)| < \text{negl}(x)$. PPT stands for probabilistic polynomial time. We use the notation $\text{Func}_{\mathcal{A}, \mathcal{B}}\langle a, b \rangle$ to refer to a protocol Func between parties \mathcal{A} and \mathcal{B} with inputs a and b , respectively. Finally, we allow algorithms to output \perp to indicate failure.

Proof Systems. We briefly review the definitions of proof systems that we will need in later sections. The main notion we will use is that of a non-interactive zero knowledge proof of knowledge in the random oracle model. We use the definitions of [26], which modify prior common reference string-based definitions of non-interactive zero-knowledge for use in the Random Oracle Model.

Definition 1 (Non-interactive Proof System). *A non-interactive proof system Π for a relation R consists of prover algorithm that on input x, w outputs a proof π and a verifier algorithm that on input x, π outputs a bit b . We say that (P, V) is correct and sound if it satisfies the following properties:*

- $(x, w) \in R \rightarrow V(x, P(x, w)) = 1$
- $(x, w) \notin R \rightarrow \Pr[V(x, P^*(x, w)) = 1] < \text{negl}$ for any (potentially cheating) prover P^* .

For convenience and clarity of notation, we use $P(\text{public}(\cdot), \text{private}(\cdot), R)$ to indicate that the public parts of the input to a prover P for relation R correspond to the statement x and that the private parts correspond to the witness w .

The zero-knowledge property [32] informally requires that a proof reveals nothing about (x, w) except that $(x, w) \in R$. Formally, we model this property

by describing a *simulator* that can provide a legitimate proof given only x and not w [5].

Extractability, informally, is a strengthening of the soundness property that requires any acceptable proof to have an *extractor* algorithm that can efficiently recover w with high probability given the ability to interact with the prover. We refer to Bellare and Goldreich [5] for a full definition. Simulation-sound extractability [34, 46, 47] further strengthens the extractability requirement of proofs of knowledge to enable extracting a witness even after seeing many simulated proofs.

EPID Signatures. We construct our EPID signature to match the syntax and security requirements as defined by Brickel and Li [15]. In this section we state the EPID syntax and sketch security requirements. Full definitions and security games appear in the full version of this paper. First, anonymity must ensure that the group manager colluding with any number of group members cannot uncover the identity of the signer. In particular, we do not want the group manager to have a tracing key that lets it compromise a group member’s identity from a signature. Nevertheless, we will later briefly explain how to extend our scheme to achieve traceability, if desired.

Second, we want a revocation property where a group manager can revoke a user’s ability to sign by either:

- adding a revoked user’s leaked signing key to a revocation list KEY-RL, or
- adding a revoked user’s EPID signature to a revocation list SIG-RL.

A user is revoked if its key is included in the list KEY-RL, or if any of its signatures are included in the list SIG-RL.

With this setup, we define the syntax and security properties for an EPID signature scheme as follows.

Definition 2 (EPID Signature). *An EPID signature scheme \mathcal{G} involving a group manager \mathcal{M} and n group members, parties \mathcal{P}_1 to \mathcal{P}_n , consists of algorithms *Init*, *Join*, *GPSign*, *GPVerify*, *RevokeKey* and *RevokeSig*:*

- $(gsk, gpk) \leftarrow \text{Init}(1^\lambda)$: *This algorithm takes as input a security parameter 1^λ and outputs a key pair (gsk, gpk) .*
- $\langle \text{cert}_i, (sk_i, \text{cert}_i) \rangle \leftarrow \text{Join}_{\mathcal{M}, \mathcal{P}_i}((gsk, gpk), gpk)$: *This is a protocol between the group manager and a group member \mathcal{P}_i where each party has its keys as input, and both parties get party \mathcal{P}_i ’s certificate as output. \mathcal{P}_i also gets its secret key sk_i as an output.*
- $\perp / \text{sig} \leftarrow \text{GPSign}(gpk, sk_i, \text{cert}_i, m, \text{SIG-RL})$: *This algorithm takes as input the public key, a signature revocation list SIG-RL, and party \mathcal{P}_i ’s secret key and certificate. The output is an EPID signature sig .*
- $1/0 \leftarrow \text{GPVerify}(gpk, m, \text{KEY-RL}, \text{SIG-RL}, \text{sig})$: *This algorithm verifies an EPID signature sig on a message m given the group public key and key/signature revocation lists KEY-RL, SIG-RL. It outputs 1 to accept the signature and 0 to reject it.*

- $KEY\text{-}RL \leftarrow \text{RevokeKey}(gpk, KEY\text{-}RL, sk_i)$: This algorithm adds a secret key sk_i to a key revocation list, so signatures created with this key will no longer be accepted.
- $SIG\text{-}RL \leftarrow \text{RevokeSig}(gpk, KEY\text{-}RL, SIG\text{-}RL, m, sig)$: This algorithm adds a signature sig to a signature revocation list, so signatures created with the same key as sig will no longer be accepted.

The algorithms must satisfy Correctness, Anonymity, and Unforgeability.

For correctness, we require that if a group member has successfully completed the Join procedure and neither its key nor any of its signatures have been revoked, then that group member's signatures should successfully verify.

We define anonymity via the Anonymity game. Informally, the property of being Anonymous requires that signatures in \mathcal{G} hide the identity of the signer against any coalition of group members (including the group manager) except the signer herself. The definition of anonymity also implies notions of unlinkability between a signer and her signatures. The game allows the adversary to create users, sign messages, and corrupt users of its choosing before attempting to distinguish which of two uncorrupted users produced a signature on a challenge message of the adversary's choice.

Finally, we define unforgeability. Our unforgeability game consists of an adversary who can add arbitrary parties to a group and corrupt arbitrarily many members of a group. Security holds if this adversary cannot forge the signature of an uncorrupted party on a message of its own choosing.

3 Post-quantum EPID Signatures

In this section we describe and prove the security of our first post-quantum EPID signature scheme. Our construction uses a standard signature scheme where each group member has its own key pair and a *certificate* from the group manager. Instead of signature keys, however, we construct our scheme so that each group member has a unique PRF secret key that will be used to issue EPID signatures. As we will see, this leads to significant savings over the general framework of Bellare et al. [6]. We still need a signature scheme for the group manager to produce certificates, but the NIZK proof is done over a circuit that verifies a single signature (the group manager's) along with a few evaluations of the PRF. An overview of the construction is as follows. Each member generates its own secret key sk for a PRF f . During the join procedure it obtains a challenge c from the group manager, sends $t = f(sk, c)$ to the manager, and obtains back a signature σ on t . To sign a message, the member first reveals $t' = (f(sk, r), r)$ for random r and then a signature of knowledge, where the proof witness is sk consistent with t' as well as σ , i.e. a signature on $f(sk, c^*)$ for some c^* . Including t' in the clear is used for signature revocation. Note that signatures need to be verified relative to the same signature revocation lists under which they were signed.

Collision Resistant PRF. We state and prove security of our scheme using a function $f : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$ that is both a secure PRF and a collision resistant

function. In fact, it suffices that f be collision-resistant *on the keyspace*, meaning that for a target input $x \in \mathcal{X}$ chosen by the adversary, it should be hard to find $k_0 \neq k_1 \in \mathcal{K}$ such that $f(k_0, x) = f(k_1, x)$. We explain how to construct an MPC-friendly function with this property in Sect. 5.

Construction 1 (EPID Signature). Our EPID signature scheme $\mathcal{G} = (\text{Init}, \text{Join}, \text{GPSign}, \text{GPVerify}, \text{RevokeKey}, \text{RevokeSig})$ with security parameter λ uses a signature scheme $\mathcal{S} = (\text{Keygen}, \text{Sign}, \text{Verify})$, a proof system $\Pi = (P, V)$, and a PRF f that also serves as a collision-resistant hash function.

- $\text{Init}(1^\lambda)$: Group manager \mathcal{M} runs $\text{Keygen}(1^\lambda)$ to get (gpk, gsk) and outputs this tuple (gpk is published and gsk kept secret).
- $\text{Join}_{\mathcal{M}, \mathcal{P}_i}((\text{gsk}, \text{gpk}), \text{gpk})$:
 - Group manager \mathcal{M} sends challenge c_i to member \mathcal{P}_i .
 - \mathcal{P}_i chooses $\text{sk}_i \xleftarrow{\mathbb{R}} \{0, 1\}^\lambda$ and sends $t_i^{\text{join}} = f(\text{sk}_i, c_i)$ back to \mathcal{M} .
 - \mathcal{M} produces signature $\sigma_i = \text{Sign}(\text{gsk}, (t_i^{\text{join}}, c_i))$, and constructs $\text{cert}_i = (t_i^{\text{join}}, c_i, \sigma_i)$, sending a copy to \mathcal{P}_i . If the signature scheme is stateful, then algorithm Join must maintain a counter that is incremented for every user who joins the group.
 - The group member’s private key is sk_i and both parties get copies of cert_i .
- $\text{GPSign}(\text{gpk}, \text{sk}_i, \text{cert}_i, m, \text{SIG-RL})$: Compute the following and output sig :
 - $r \xleftarrow{\mathbb{R}} \{0, 1\}^\lambda \setminus c_i$
 - $t \leftarrow (f(\text{sk}_i, r), r)$
 - $\pi \leftarrow P(\text{public}(\lambda, m, \text{gpk}, t, \text{SIG-RL}, \text{KEY-RL}), \text{private}(\text{sk}_i, \text{cert}_i), R_1)$
 - $\text{sig} \leftarrow (t, \pi)$.

We define the relation R_1 in the proof of knowledge π for $(\text{sk}_i^*, \text{cert}_i^*)$ to be true when the following statements hold:

- $t = (f(\text{sk}_i^*, r), r)$
- $r \neq c_i^*$
- $\text{Verify}(\text{gpk}, (t_i^{\text{join}*}, c_i^*), \sigma_i^*) = 1$
- $t_i^{\text{join}*} = f(\text{sk}_i^*, c_i^*)$
- for each $\text{sig}_j \in \text{SIG-RL}$, $t_{\text{sig}_j} \neq (f(\text{sk}_i^*, r_{\text{sig}_j}), r_{\text{sig}_j})$
- $\text{GPVerify}(\text{gpk}, m, \text{KEY-RL}, \text{SIG-RL}, \text{sig})$:
 - Verify proof π : check $V((\lambda, m, \text{gpk}, t, \text{SIG-RL}, \text{KEY-RL}), \pi) = 1$.
 - For each $\text{sk}_j \in \text{KEY-RL}$, check that $t \neq (f(\text{sk}_j, r), r)$.
 - Check that $\text{sig} \notin \text{SIG-RL}$.
 - Output 1 if all of the above checks return 1; otherwise, output 0.
- $\text{RevokeKey}(\text{gpk}, \text{KEY-RL}, \text{sk}_i)$: Return $\text{KEY-RL} \cup \{\text{sk}_i\}$.
- $\text{RevokeSig}(\text{gpk}, \text{KEY-RL}, \text{SIG-RL}, m, \text{sig})$: return $\text{SIG-RL} \cup \{\text{sig}\}$ if $\text{GPVerify}(\text{gpk}, m, \text{KEY-RL}, \text{SIG-RL}, \text{sig}) = 1$. Otherwise, return SIG-RL .

Revocation. Although the difference between the two forms of revocation does not affect our scheme’s security, the effect of revocation differs in practice depending on whether a group member is revoked by key or by signature. A revocation by key renders all signatures, past or future, invalid for that user, whereas a

revocation by signature only applies to future signatures because past signatures need to be verified with respect to the SIG-RL in place at the time of signing. This does not matter for the purposes of the security game because the attempted forgery is always the last signature produced in the game. For the same reason, the decision to include the check that $\text{sig} \notin \text{SIG-RL}$ during GPVerify does not affect security for the purpose of the proof and can be omitted. We include it only to better capture behavior that may be expected of revocation in practice.

Traceable Signatures. Our approach can also be used to achieve *traceability*. Traceability requires that the group manager have the power to learn the identity of a signer. We presented our scheme without a tracing property in order to guarantee a stronger anonymity property against the group manager, but a similar approach could be used to achieve traceability. The group manager could give each group member a signed secret token sk_i'' , and every signature would include the token $t' = (f(\text{sk}_i'', r'), r')$, for a newly picked random r' , along with a proof of knowledge of a signature on sk_i'' . Now the group manager can trace a signature by trying to reconstruct t' with the value of sk_i'' for each signer, but anonymity will still hold against any other group member.

Camenisch and Groth [17] give a traceable group signature scheme from one-way functions and NIZKs. Although their scheme can be instantiated under the same assumptions as ours, they (loosely speaking) include a commitment to a credential for each group member in their public key and give a proof of knowledge that a signature corresponds to one of those credentials. By avoiding this cost, our scheme shrinks both the public key size and signature size by a factor $O(N)$. Our public key can also be published at group initialization time before any members have joined the group.

Security Theorems. We now state our various theorems regarding the security of our scheme and give a brief intuition to justify them. Proofs are deferred to the full version of this paper. Correctness follows almost immediately from the construction with the caveat that we must ensure that the revocation checks do not accidentally cause a signature from a legitimate key to be rejected.

Theorem 3. *Assuming the correctness of signature scheme \mathcal{S} and proof system Π and the pseudorandomness of f , \mathcal{G} is a correct EPID signature scheme.*

Anonymity follows from the zero-knowledge and pseudorandomness properties of the primitives used in our construction. Intuitively, the scheme achieves anonymity against the group manager by having each group member generate its own PRF secret key sk , and from all other parties because the signatures are zero-knowledge signatures of knowledge.

Theorem 4. *Assuming that Π is a zero-knowledge proof system and that f is a PRF, \mathcal{G} is an anonymous EPID signature scheme.*

The high level intuition for unforgeability is as follows. If the adversary A has not obtained a signature from the group manager on $t = f(\text{sk}, c)$ then it

cannot produce a PoK of valid signature on t by unforgeability of the group manager’s signature scheme and soundness of the PoK. Second, if A does not know sk for some $t = f(sk, c)$ that has been signed, then even though it sees many $f(sk, r), r$ inside signatures it cannot produce $f(sk, c^*)$ on a fresh c^* by the security of the PRF. (Note that even if it were able to do this it has to actually learn sk to forge a signature as otherwise it breaks PoK extractability). Finally, collision-resistance of the PRF ensures that A who has a signature on $f(sk, c)$ for revoked sk cannot find $sk' \neq sk$ and r such that $f(sk', c) = t$ and $f(sk', r) \neq f(sk, r)$.

Theorem 5. *Assuming that Π is a zero knowledge proof of knowledge proof system with simulation-sound extractability, \mathcal{S} is an unforgeable signature scheme, that f is a PRF, and that f is additionally a collision-resistant hash function, \mathcal{G} is an unforgeable EPID signature scheme.*

4 Practical Post-quantum Signatures for Attestation

Attestation schemes (such as that used in Intel SGX [22,36]) involve interaction with an attestation service on every attestation, among other reasons to obtain an updated revocation list. In the case of SGX, this attestation service is also the group manager. In this section, we present a significantly smaller post-quantum EPID-like signature scheme appropriate for this setting where frequent interaction with the group manager is allowed.

The main bottleneck for signature size in our first construction was including verification of the group manager’s signature on a group member’s certificate inside the PoK (i.e. this contributed the most to arithmetic complexity). We remove this signature in our new scheme by making each group member’s certificate a leaf in a Merkle tree. The group manager signs only the root, providing each group member an inclusion proof during Join. The signature on the root can be public as it leaks nothing about the identity of a member. Signers now only need to include the Merkle inclusion proof inside the proof of knowledge instead of a hash-based signature. The verification of an inclusion proof requires a much smaller circuit.

This modification has several implications for security. As a new Merkle tree root will need to be published each time a group member joins, this reduces the size of anonymity sets. In an extreme case the group manager could issue a sequence of Merkle roots where each tree only included a valid credential for one group member, uniquely identifying the member’s signatures.

Fortunately, the continuing contact between group members and the group manager enforced by attestation in practice enable effective mitigations for these concerns. Group members can periodically “re-join” the group to update the Merkle root relative to which they provide membership proofs, thereby increasing the size of their anonymity sets. In practice, we can ensure that subsequent Merkle roots issued by the group manager only ever add new credentials to the group and never omit previous ones by using a *Merkle consistency proof* such

as the one proposed by the Certificate Transparency standard [38] and proven secure by Dowling et al. [25]. We model the Merkle trees used in our proofs as accumulators with zero-knowledge membership proofs and discuss how we instantiate this primitive with an improved construction in Sect. 5.

4.1 Definitions

In this section we define accumulators and EPID-like signatures for attestation. We begin with a special case of the formalization of accumulators by [23].

Definition 6 (Accumulator). *A static accumulator is a tuple of efficient algorithms (A_{Gen} , A_{Eval} , $A_{WitCreate}$, A_{Verify} , $A_{ProveCon}$, $A_{CheckCon}$) which are defined as follows:*

- $A_{Gen}(1^\lambda)$: *This algorithm takes a security parameter λ and returns a public key pk_Λ .*
- $A_{Eval}(pk_\Lambda, \mathcal{X})$: *This deterministic algorithm takes a key pk_Λ and a set \mathcal{X} to be accumulated and returns an accumulator $\Lambda_\mathcal{X}$.*
- $A_{WitCreate}(pk_\Lambda, \Lambda_\mathcal{X}, \mathcal{X}, x_i)$: *This algorithm takes a key pk_Λ , an accumulator $\Lambda_\mathcal{X}$, the set \mathcal{X} , and a value x_i . It returns \perp if $x_i \notin \mathcal{X}$ and a witness wit_{x_i} for x_i otherwise.*
- $A_{Verify}(pk_\Lambda, \Lambda_\mathcal{X}, wit_{x_i}, x_i)$: *This algorithm takes a public key pk_Λ , an accumulator $\Lambda_\mathcal{X}$, a witness wit_{x_i} , and a value x_i . It returns 1 if wit_{x_i} is a witness for $x_i \in \mathcal{X}$ and 0 otherwise.*

We require accumulators to be correct, meaning that A_{Verify} will accept an honestly generated witness for $x_i \in \mathcal{X}$. We also require a soundness property dubbed *collision-freeness*, formally defined below.

Definition 7 (Collision Freeness). *An accumulator is collision free if for all PPT adversaries \mathcal{A} , we have that*

$$Pr[A_{Verify}(pk_\Lambda, \Lambda^*, wit_{x_i}^*, x_i^*) = 1 \wedge x_i^* \notin \mathcal{X}^* | pk_\Lambda \leftarrow A_{Gen}(1^\lambda, \Lambda^*), \Lambda^* \leftarrow Eval_{r^*}(pk_\Lambda, \mathcal{X}^*), (wit_{x_i}^*, x_i^*, \mathcal{X}^*) \leftarrow \mathcal{A}(pk_\Lambda, \Lambda^*)] \leq negl(\lambda)$$

The setting of EPID signatures for attestation largely leaves the security definitions of Sect. 3 unaffected up to changes in syntax, so we present the updated syntax in the full version and omit statements of the security definitions. The only notable changes are that (1) in both security games the adversary can now choose to have a group member run the new $GAREjoin$ at any time it chooses, and (2) signatures are only indistinguishable for two signatures produced relative to the same accumulator.

4.2 EPID Signature Construction II

The full construction of the modified EPID signature scheme appears below. Structurally similar to the construction in Sect. 3, the main changes involve the introduction of a post-quantum accumulator and the resulting restructuring of what needs to be proven inside/outside the proof of knowledge π , as described informally at the beginning of this section.

Construction 2 (EPID Signature for Attestation). Our EPID signature scheme for attestation $\mathcal{GA} = (\text{GALnit}, \text{GAJoin}, \text{GAREjoin}, \text{GASign}, \text{GAVerify}, \text{GAREvokeKey}, \text{GAREvokeSig})$ with security parameter λ uses a signature scheme $\mathcal{S} = (\text{Keygen}, \text{Sign}, \text{Verify})$, a proof system $\Pi = (P, V)$, a PRF f that also serves as a collision-resistant hash function, and an accumulator $\mathcal{A} = (\text{AGen}, \text{AEval}, \text{AWitCreate}, \text{AVerify})$.

- $\text{GALnit}(1^\lambda)$: Group manager \mathcal{M} runs $\text{Keygen}(1^\lambda)$ to get $(\text{pk}_{\text{gp}}, \text{sk}_{\text{gp}})$ and runs $\text{AGen}(1^\lambda)$, to get pk_\wedge . It outputs public key $\text{gpk} = (\text{pk}_{\text{gp}}, \text{pk}_\wedge)$ and secret key $\text{gsk} = \text{sk}_{\text{gp}}$.
- $\text{GAJoin}_{\mathcal{M}, \mathcal{P}_i} \langle (\text{gsk}, \text{gpk}, \mathcal{X}), \text{gpk} \rangle$:
 - Group manager \mathcal{M} sends challenge c_i to member \mathcal{P}_i .
 - \mathcal{P}_i picks $\text{sk}_i \xleftarrow{\mathbb{R}} \{0, 1\}^\lambda$ and sends $t_i^{\text{join}} = f(\text{sk}_i, c_i)$ back to \mathcal{M} .
 - \mathcal{M} defines $x_i = (t_i^{\text{join}}, c_i)$, sets $\mathcal{X} = \mathcal{X} \cup x_i$, sets $\Lambda = \text{AEval}(\text{pk}_\wedge, \mathcal{X})$, and produces signature $\sigma_\wedge = \text{Sign}(\text{gsk}, \Lambda)$. Next, \mathcal{M} creates $\text{wit}_{x_i} = \text{AWitCreate}(\text{pk}_\wedge, \Lambda, \mathcal{X}, x_i)$ and constructs $\text{cert}_i = (x_i, \text{wit}_{x_i})$, sending a copy to \mathcal{P}_i along with Λ and σ_\wedge .
 - The group member's private key is sk_i and both parties get copies of cert_i , Λ , and σ_\wedge .
- $\text{GAREjoin}_{\mathcal{M}, \mathcal{P}_i} \langle (\text{gsk}, \text{gpk}, \mathcal{X}, \Lambda, \sigma_\wedge), (\text{gpk}, \text{cert}_i) \rangle$:
 - \mathcal{P}_i sends cert_i to \mathcal{M} .
 - First, \mathcal{M} verifies the signature in cert_i , aborting in case of failure. Then it creates a new $\text{wit}_{x_i} = \text{AWitCreate}(\text{pk}_\wedge, \Lambda, \mathcal{X}, x_i)$ and constructs the updated $\text{cert}_i = (x_i, \text{wit}_{x_i})$, sending a copy to \mathcal{P}_i along with Λ and σ_\wedge .
 - \mathcal{P}_i updates its values of cert_i , Λ , and σ_\wedge .
- $\text{GASign}(\text{gpk}, \text{sk}_i, \text{cert}_i, m, \text{SIG-RL}, \Lambda, \sigma_\wedge)$: Compute the following and output sig :
 - $\text{Verify}(\text{pk}_{\text{gp}}, \sigma_\wedge, \Lambda)$ (abort if it outputs 0)
 - $r \xleftarrow{\mathbb{R}} \{0, 1\}^\lambda \setminus c_i$
 - $t = (f(\text{sk}_i, r), r)$
 - $\pi = P(\text{public}(\lambda, m, \text{gpk}, t, \text{SIG-RL}, \text{KEY-RL}, \Lambda), \text{private}(\text{sk}_i, \text{cert}_i), R_2)$
 - $\text{sig} = (t, \pi, \Lambda, \sigma_\wedge)$.

We define R_2 as a relation in the proof of knowledge of $(\text{sk}_i^*, \text{cert}_i^*)$ such that the following statements hold:

- $t = (f(\text{sk}_i^*, r), r)$
- $r \neq c_i^*$
- $\text{AVerify}(\text{pk}_\wedge, \Lambda, \text{wit}_{x_i}^*, (t_i^{\text{join}*}, c_i^*))$
- $t_i^{\text{join}*} = f(\text{sk}_i^*, c_i^*)$
- for each $\text{sig}_j \in \text{SIG-RL}$, $t_{\text{sig}_j} \neq (f(\text{sk}_i^*, r_{\text{sig}_j}), r_{\text{sig}_j})$
- $\text{GAVerify}(\text{gpk}, m, \text{KEY-RL}, \text{SIG-RL}, \text{sig})$:
 - Verify signature σ_\wedge : check $\text{Verify}(\text{pk}_{\text{gp}}, \sigma_\wedge, \Lambda) = 1$
 - Verify proof π : check $V((\lambda, m, \text{gpk}, t, \text{SIG-RL}, \text{KEY-RL}, \Lambda), \pi) = 1$.
 - For each $\text{sk}_j \in \text{KEY-RL}$, check that $t \neq (f(\text{sk}_j, r), r)$.
 - Check that $\text{sig} \notin \text{SIG-RL}$.
 - If all of the above checks return 1, output 1. Else, output 0.
- $\text{GAREvokeKey}(\text{gpk}, \text{KEY-RL}, \text{sk}_i)$: Return $\text{KEY-RL} \cup \text{sk}_i$.

- $\text{GARevokeSig}(\text{gpk}, \text{KEY-RL}, \text{SIG-RL}, m, \text{sig})$: If $\text{GAVerify}(\text{gpk}, m, \text{KEY-RL}, \text{SIG-RL}, \text{sig}) = 1$, return $\text{SIG-RL} \cup \text{sig}$. Otherwise, return SIG-RL .

Security Theorems. Correctness and anonymity proofs for \mathcal{GA} are almost completely unchanged from our standard EPID signature scheme, so we only state the corresponding theorems. The only proof that needs some tweaking is that of unforgeability, which we sketch in the full version of this paper.

Theorem 8. *Assuming the correctness of signature scheme \mathcal{S} , proof system Π , and accumulator \mathcal{A} , as well as the pseudorandomness of f , \mathcal{GA} is a correct EPID signature scheme.*

Theorem 9. *Assuming that Π is a zero-knowledge proof system and that f is a PRF, \mathcal{GA} is an anonymous EPID signature scheme.*

Theorem 10. *Assuming that Π is a proof system for zero-knowledge proofs of knowledge with simulation-sound extractability, \mathcal{S} is an unforgeable signature scheme, that f is a PRF, that f is additionally a collision-resistant hash function, and that \mathcal{A} is a collision-free (sound) accumulator, \mathcal{GA} is an unforgeable EPID signature scheme.*

5 Instantiation of Protocols

We have now described and proven the security of our constructions, but the post-quantum security of each construction relies on the existence of post-quantum secure instantiations of the various primitives required. In particular we require a PRF that is also a collision-resistant hash function, a signature scheme, zero knowledge proofs of knowledge (ZKPoKs), and an accumulator. In this section we describe options for instantiating each primitive under different security assumptions about the underlying ciphers used and report the signature sizes of our instantiated schemes in both the Random Oracle (RO) and Quantum Random Oracle (QRO) models [11].

5.1 Zero Knowledge Proofs of Knowledge

In principle, standard symmetric primitives (AES, SHA) suffice for post-quantum security so long as we double our security parameters. However, our schemes uses these primitives in a non-black box manner by running them inside of a ZKPoK. In particular, the following ZKPoKs contribute significantly to signature sizes:

1. ZKPoK of a PRF key k such that $f(k, r) = t$, for a PRF that is collision-resistant on its key space.
2. ZKPoK of a signature σ on a message m such that $\text{Verify}(m, \sigma) = 1$ for a post-quantum signature scheme $\mathcal{S} = (\text{Keygen}, \text{Sign}, \text{Verify})$.
3. ZKPoK of membership of element x_i in accumulator Λ for set \mathcal{X} .

We restrict our choice of ZKPoK proof system to those systems which rely only on symmetric primitives. This includes works following the “MPC in the Head” approach of Ishai et al. [35] – ZKBoo [30], ZKB++ [20], and Ligerio [3] – as well as zk-STARKs [8]. Although Ligerio and zk-STARKs offer proofs asymptotically sublinear in the size of the circuit to be proven, a preliminary analysis suggested that, for our relatively small proof circuits, ZKB++ provides the smallest signature sizes in practice without requiring heavy computing costs for the signer. Moreover, ZKB++ has proofs of security in both the Random Oracle and Quantum Random Oracle models, whereas Ligerio and zk-STARKs only have proofs in the classical RO model. As such, we choose to instantiate our signatures and measure signature size using ZKB++ as our underlying ZKPoK.

In ZKB++ [20], the underlying statement to be proven is represented as an arithmetic circuit over $\text{GF}(2)$, and the proof size is proportional to the multiplicative complexity (i.e., number of AND gates) in the circuit. The most important practical consideration in signature schemes is signature size; therefore our main criterion in instantiating the PRF and outer signature scheme is to minimize their multiplicative complexity over $\text{GF}(2)$.

5.2 PRF and Collision-Resistant Hash Function

Recently the ciphers LowMC [2] and MiMC [1] have been proposed as alternatives to AES that have significantly lower multiplicative complexity as arithmetic circuits over finite fields.¹ Although relatively new and less extensively studied, these ciphers were shown to resist statistical cryptanalytic attacks, similar to other state-of-the-art designs. A number of works have already proposed using LowMC as the best candidate to-date for instantiating ciphers inside ZKB++-style proofs [20, 24]. The most recent public version of the LowMC cipher with parameters set for 128-bit post-quantum security (256-bit key, 256-bit block size) involves only 1374 AND gates, a significant improvement over the 7616 AND gates in AES-256 [2].

Derler et al. [24] also suggest using the LowMC round function in the sponge framework (as described in [1]) to construct a collision-resistant hash function with low multiplicative complexity. However, since only a collision-resistant compression function on a fixed message length is needed (rather than full-blown indistinguishability from a random oracle), we propose applying the much simpler Davies-Meyer transformation to the LowMC cipher. Collision resistance of Davies-Meyer is proved in the ideal cipher model [13], which is only marginally stronger than the security assumption underlying the sponge transformation. Given an ideal cipher $E(k, x)$ on equal sized key and message space, the Davies-Meyer compression function is $H(m_1 || m_2) = E(m_1, m_2) \oplus m_2$. For a collision-resistant PRF we would use $F(k, x) = E(k, x) \oplus x$; as long as E is a PRF then F is also a PRF. Note that the multiplicative complexity of F is the same as E .

¹ LowMC optimizes multiplicative complexity over $\text{GF}(2)$ while MiMC optimizes complexity over larger finite fields. In ZKB++ the underlying circuit is represented in $\text{GF}(2)$, which is why we prefer LowMC.

To obtain a PRF that is collision-resistant only on its keyspace we can rely on a slightly weaker assumption than the ideal cipher model. The ideal cipher model assumes that E with *any* key is indistinguishable from a random permutation, whereas we only need to assume there is an explicit fixed key k_{fix} on which $E(k_{\text{fix}}, \cdot)$ is indistinguishable from a random permutation. Then we can define $\Pi(y) = E(k_{\text{fix}}, y)$, and define $F'(k, x) = \Pi(E(k, x)) \oplus E(k, x)$. (The inner evaluation of $E(k, x)$ ensures the PRF property while $\Pi(y) \oplus y$ is collision resistant as a special case of Davies-Meyer).

5.3 Post-quantum Signature Scheme

Choices for post-quantum signatures that do not rely on stronger lattice assumptions include Merkle signatures [43], Goldreich's stateless signatures [31], SPHINCS signatures [9], or the Fish signatures of Chase et al. [20]. The recent literature on post-quantum signatures has focused on optimizing signature size. When using signatures outside of proofs (in our construction of EPID signatures for attestation) we propose using SPHINCS, which has the smallest signature size. However, since our main EPID signature construction involves verifying the group manager's post-quantum signature inside a ZKPoK, there we care about optimizing the arithmetic multiplicative complexity of signature verification rather than the signature size.

We examine two options for instantiating the group manager's signature scheme for signatures used inside a ZKPoK: one using stateful Merkle signatures, and other using Goldreich's stateless signatures.

Stateful Merkle Signatures. The signer runs a signature setup that generates a large number of one-time signature (OTS) keypairs. We would use Lamport signatures from one-way functions (instantiated with LowMC) for the OTS. The Lamport signature private key consists of 256 pairs of pseudorandom 256-bit strings the public key consists of the 256 pairs of outputs generated by applying the one-way function to each private key string. The signer finalizes the setup by computing a Merkle tree (using a 2-to-1 collision resistant compression function) over the OTS public keys at the leaves of the tree and publishing the root as the public verification key. Signing a message involves signing the message with one of the leaf OTS keys and proving membership of this OTS key in the Merkle tree. The signer needs to maintain state to ensure that no OTS key is used more than once. The stateful requirement is not prohibitive in the setting of managing a group of trusted hardware platforms. The preprocessing of a tree of up to 2^{30} members would take under a day on modern commodity hardware and would require the server to use only several GB of storage.

Stateless Goldreich Signatures. Instead of maintaining state in the Merkle signature scheme above, the signer could choose an OTS key at random. This requires squaring the size of the tree to make collisions unlikely. For a group of 2^{30} members storing a tree of size 2^{60} keys would be prohibitively expensive. However, Goldreich's scheme provides a way to generate this tree pseudorandomly from a

small seed. In this scheme, the signer pseudorandomly generates an OTS keypair for each node of the tree, which can be done by evaluating a PRF on the index of the tree node. The OTS public key at the root of the tree is the overall public key. The OTS key pair on each node of the tree is used to sign the hash of the public keys on each of its two child nodes. To sign a message a random leaf is selected and the signature includes the OTS signatures along the path from this leaf to the root, where each signature signs either a child public key or the actual message at the leaf.

5.4 Reducing Circuit Size for Membership Proofs

As mentioned in Sect. 4, we will use Merkle trees to instantiate our accumulators. A recent work of Derler et al. [24] points out, however, that the circuit used to verify standard Merkle inclusion proofs differs based on the path from the Merkle root to the leaf x_i . The dependence arises based on whether the hash at depth j of the tree becomes the left or right input of the hash at depth $j - 1$. This dependence of the `AVerify` circuit on i must be removed in order to generically create a zero-knowledge inclusion proof with some zero-knowledge proof system. They suggest a modification to the standard inclusion proof that allows the same circuit to verify inclusion regardless of the index i whose inclusion is proven. The idea is as follows: suppose x_i resides in a subtree rooted at internal node a and that a has sibling and parent nodes b and c , respectively. At each level of the Merkle tree, instead of simply calculating $h(a, b)$ and only comparing the result to the root, they evaluate the expression $c = h(a, b) \vee c = h(b, a)$ and reject the inclusion proof if it is not satisfied. This allows the construction of a circuit `AVerify'` with a fixed ordering of inputs to each hash function, since as long as one ordering of inputs matches the node at the next level of the tree, correctness will hold. The cost of this transformation is an extra hash evaluation, an equality check, and a logical OR for each level of the tree.

We propose a solution that eliminates the need for equality checks at each level of the tree and replaces the OR with an XOR, allowing smaller and more efficient zero-knowledge membership proofs. Our idea is to replace the hash function h already used in computing the merkle root with a modified function $h'(x, y) = h(x, y) \oplus h(y, x)$. Using h' in place of h proves that the input x_i is a d^{th} preimage of the merkle root for a tree of depth d without any dependence on the position i of x_i among the tree's leaves. Of course, h' is trivially neither collision-resistant nor second preimage resistant, as a swapping of the inputs x and y results in the same output. Below we prove that h' provides a *third* preimage resistance property and helps build the inclusion proofs we desire.

Definition 11 (Third Preimage Resistance). *We say a hash function H defined over $(\mathcal{M}, \mathcal{T})$ is third preimage resistant if given a random $m = a||b \in \mathcal{M}$ (with $|a| = |b|$) and a different $m' = b||a \in \mathcal{M}$ such that $H(m) = H(m')$, it is difficult to find an $m'' \in \mathcal{M}$ such that $H(m'') = H(m) = H(m')$.*

Lemma 12. *Assuming the hash function $h : \mathcal{M} \times \mathcal{M} \rightarrow \mathcal{M}$ is a random function, the hash function $h'(x, y) = h(x, y) \oplus h(y, x)$ for $x, y \in \mathcal{M}$ is third preimage resistant, provided $x \neq y$.*

Proof. $h'(x, y)$ admits a trivial collision $h'(y, x)$. We argue it is hard to find any other collision unless $x = y$ (since $h'(x, x) = 0$ for all x). To find a third preimage of $h'(x, y)$ an adversary must produce w, z such that either $h'(w, z) = h'(x, y)$ and either $w \neq x$ or $z \neq y$. Since h is a random function and $(x, y), (y, x), (w, z), (z, w)$ are all distinct tuples, $h(x, y), h(y, x), h(w, z)$, and $h(z, w)$ will all be independently random strings. The probability that $h(x, y) \oplus h(y, x) = h(w, z) \oplus h(z, w)$ is therefore negligible in the length $|x| + |y|$. Therefore no efficient adversary can find a third preimage for h' . \square

In order to replace h with h' in our merkle tree construction and retain security for the circuit `AVerify'`, we only need to show that we will have no leaves $x||y$ in the accumulator such that $x = y$. Fortunately, since the elements in the accumulator for our particular case are challenge/response pairs $(f(sk_i, c_i), c_i)$ that serve as group member credentials (where f is collision-resistant and a PRF), the probability that $x = y$ is negligible in our setting.

Practically, our new circuit `AVerify'` reduces the number of equality checks inside a ZKPoK from $2 \log_2(N)$ (where N is the group size) to 1. Additionally, $\log_2(N)$ OR gates are replaced with XORs which do not increase proof size.

5.5 Signature Sizes

As discussed above, we instantiate our signatures using LowMC, Merkle signatures (inside the ZKPoK), SPHINCS signatures (outside the ZKPoK), ZKB++, and Merkle tree accumulators with our modified membership proof circuit.

Figure 1 shows the sizes for our modified EPID signatures for various group sizes under (1) the assumption that LowMC is an ideal cipher and (2) the assumption that LowMC with a public fixed key is a random permutation. Figure 2 presents the same information, but uses the Unruh transform [48] instead of the Fiat-Shamir transform [27] to make the ZKB++ proof noninteractive. The Fiat-Shamir transform is proven secure in the Random Oracle model but only sometimes retains security in the Quantum Random Oracle model [11, 49]. As visible from the figures, groups of size up to 2^{20} could use post-quantum signatures of size 6.74 MB (3.45 MB in RO model) under our scheme, a sufficiently small size for attestation in applications with heavy data transfer requirements. For comparison, the same signatures instantiated with AES-256 would require 33.8 MB (16.9 MB in RO model), meaning the choice of LowMC enables a $5\times$ improvement in signature size.

For comparison, our signature sizes are smaller than the recent ring signatures of Derler et al. [24], which require at least 10.4 MB (5.26 MB in RO Model) for signatures in a ring of 2^{20} members², despite providing a more elaborate

² This size represents an optimized version of the ring signatures instantiated assuming LowMC is an ideal cipher. The original Derler et al. paper claimed slightly larger signatures of size 11.88 MB (8 MB in RO Model) for this ring size.

Signature Sizes in RO Model

Group Size	Ideal Cipher	Random Permutation
2^7	1.37MB	2.28MB
2^{10}	1.85MB	3.21MB
2^{20}	3.45MB	6.31MB
2^{30}	5.05MB	9.41MB
2^{40}	6.65MB	12.5MB

Fig. 1. Signature sizes for construction II under various security assumptions on LowMC, using Fiat-Shamir [27] to make proofs of knowledge noninteractive.

Signature Sizes in QRO Model

Group Size	Ideal Cipher	Random Permutation
2^7	2.64MB	4.45MB
2^{10}	3.59MB	6.30MB
2^{20}	6.74MB	12.5MB
2^{30}	9.89MB	18.6MB
2^{40}	13.0MB	24.8MB

Fig. 2. Signature sizes for construction II under various security assumptions on LowMC, using the Unruh transform [48] to make proofs of knowledge noninteractive.

functionality. The improvement comes from our new accumulator membership proofs, as the accumulator constitutes the most costly component of both constructions. Note that subsequent to our paper, the Derler et al. paper has been updated with new results that shrink their signatures by a factor of 2. Their techniques can reduce signature sizes in our construction II as well.

Our general-purpose EPID signatures require 216.82 MB for signatures in a group of size 2^{30} assuming LowMC is an ideal cipher (110.81 MB in QRO Model), a much larger value than the variation designed for attestation. This motivates the question of how to generalize the specialized version of our construction to apply to a wider range of use-cases, which we leave as an open problem.

6 Conclusion

We presented a general-purpose post-quantum EPID signature scheme as well as a construction of a specialized variant designed for trusted hardware enclave attestation. We also gave an analysis of the concrete sizes of our signatures based on the best possible instantiations with current tools and showed that our signatures for attestation can achieve sizes acceptable for use in some applications.

EPID signatures play an important role in modern trusted hardware. Making them post-quantum secure is an important goal, and we hope this work will spur further research on this question that will further reduce the signature size.

Acknowledgments. We would like to thank David Wu for several helpful conversations. This work is supported by NSF, the DARPA/ARL SAFEWARE project, the Simons foundation, and a grant from ONR. The views expressed are those of the author and do not reflect the official policy or position of the Department of Defense, the National Science Foundation, or the U.S. Government.

References

1. Albrecht, M., Grassi, L., Rechberger, C., Roy, A., Tiessen, T.: MiMC: efficient encryption and cryptographic hashing with minimal multiplicative complexity. In: Cheon, J.H., Takagi, T. (eds.) ASIACRYPT 2016. LNCS, vol. 10031, pp. 191–219. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53887-6_7
2. Albrecht, M.R., Rechberger, C., Schneider, T., Tiessen, T., Zohner, M.: Ciphers for MPC and FHE. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9056, pp. 430–454. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46800-5_17
3. Ames, S., Hazay, C., Ishai, Y., Venkatasubramanian, M.: Liger: lightweight sub-linear arguments without a trusted setup. In: CCS, pp. 2087–2104 (2017)
4. El Bansarkhani, R., Misoczki, R.: G-merkle: a hash-based group signature scheme from standard assumptions. IACR Cryptology ePrint Archive (2018)
5. Bellare, M., Goldreich, O.: On defining proofs of knowledge. In: Brickell, E.F. (ed.) CRYPTO 1992. LNCS, vol. 740, pp. 390–420. Springer, Heidelberg (1993). https://doi.org/10.1007/3-540-48071-4_28
6. Bellare, M., Micciancio, D., Warinschi, B.: Foundations of group signatures: formal definitions, simplified requirements, and a construction based on general assumptions. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 614–629. Springer, Heidelberg (2003). https://doi.org/10.1007/3-540-39200-9_38
7. Bellare, M., Shi, H., Zhang, C.: Foundations of group signatures: the case of dynamic groups. In: Menezes, A. (ed.) CT-RSA 2005. LNCS, vol. 3376, pp. 136–153. Springer, Heidelberg (2005). https://doi.org/10.1007/978-3-540-30574-3_11
8. Ben-Sasson, E., Bentov, I., Horesh, Y., Riabzev, M.: Scalable, transparent, and post-quantum secure computational integrity. IACR Cryptology ePrint Archive (2018)
9. Bernstein, D.J., et al.: SPHINCS: practical stateless hash-based signatures. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9056, pp. 368–397. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46800-5_15
10. Boneh, D., Boyen, X., Shacham, H.: Short group signatures. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 41–55. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-28628-8_3
11. Boneh, D., Dagdelen, Ö., Fischlin, M., Lehmann, A., Schaffner, C., Zhandry, M.: Random oracles in a quantum world. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 41–69. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-25385-0_3
12. Boneh, D., Shacham, H.: Group signatures with verifier-local revocation. In: Proceedings of the 11th ACM Conference on Computer and Communications Security (CCS), pp. 168–177. ACM (2004)
13. Boneh, D., Shoup, V.: A Graduate Course in Applied Cryptography (2017)
14. Brickell, E.F., Camenisch, J., Chen, L.: Direct anonymous attestation. In: CCS, pp. 132–145 (2004)

15. Brickell, E., Li, J.: Enhanced privacy ID from bilinear pairing. IACR Cryptology ePrint Archive, 2009:95 (2009)
16. Buchmann, J., Dahmen, E., Hülsing, A.: XMSS - a practical forward secure signature scheme based on minimal security assumptions. In: Yang, B.-Y. (ed.) PQCrypto 2011. LNCS, vol. 7071, pp. 117–129. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-25405-5_8
17. Camenisch, J., Groth, J.: Group signatures: better efficiency and new theoretical aspects. In: Blundo, C., Cimato, S. (eds.) SCN 2004. LNCS, vol. 3352, pp. 120–133. Springer, Heidelberg (2005). https://doi.org/10.1007/978-3-540-30598-9_9
18. Camenisch, J., Lysyanskaya, A.: A signature scheme with efficient protocols. In: Cimato, S., Persiano, G., Galdi, C. (eds.) SCN 2002. LNCS, vol. 2576, pp. 268–289. Springer, Heidelberg (2003). https://doi.org/10.1007/3-540-36413-7_20
19. Camenisch, J., Lysyanskaya, A.: Signature schemes and anonymous credentials from bilinear maps. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 56–72. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-28628-8_4
20. Chase, M., et al.: Post-quantum zero-knowledge and signatures from symmetric-key primitives. In: CCS, pp. 1825–1842 (2017)
21. Chaum, D., van Heyst, E.: Group signatures. In: Davies, D.W. (ed.) EUROCRYPT 1991. LNCS, vol. 547, pp. 257–265. Springer, Heidelberg (1991). https://doi.org/10.1007/3-540-46416-6_22
22. Costan, V., Devadas, S.: Intel SGX explained. IACR Cryptology ePrint Archive 2016:86 (2016)
23. Derler, D., Hanser, C., Slamanig, D.: Revisiting cryptographic accumulators, additional properties and relations to other primitives. In: Nyberg, K. (ed.) CT-RSA 2015. LNCS, vol. 9048, pp. 127–144. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-16715-2_7
24. Derler, D., Ramacher, S., Slamanig, D.: Post-quantum zero-knowledge proofs for accumulators with applications to ring signatures from symmetric-key primitives. IACR Cryptology ePrint Archive (2017)
25. Dowling, B., Günther, F., Herath, U., Stebila, D.: Secure logging schemes and certificate transparency. In: Askoxylakis, I., Ioannidis, S., Katsikas, S., Meadows, C. (eds.) ESORICS 2016. LNCS, vol. 9879, pp. 140–158. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-45741-3_8
26. Faust, S., Kohlweiss, M., Marson, G.A., Venturi, D.: On the non-malleability of the fiat-shamir transform. In: Galbraith, S., Nandi, M. (eds.) INDOCRYPT 2012. LNCS, vol. 7668, pp. 60–79. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-34931-7_5
27. Fiat, A., Shamir, A.: How to prove yourself: practical solutions to identification and signature problems. In: Odlyzko, A.M. (ed.) CRYPTO 1986. LNCS, vol. 263, pp. 186–194. Springer, Heidelberg (1987). https://doi.org/10.1007/3-540-47721-7_12
28. Fisch, B., Vinayagamurthy, D., Boneh, D., Gorbunov, S.: IRON: functional encryption using intel SGX. In: CCS, pp. 765–782 (2017)
29. Fuller, B., et al.: Sok: cryptographically protected database search. In: IEEE Symposium on Security and Privacy (Oakland), pp. 172–191 (2017)
30. Giacomelli, I., Madsen, J., Orlandi, C.: ZKBoo: faster zero-knowledge for boolean circuits. In: USENIX Security, pp. 1069–1083 (2016)
31. Goldreich, O.: Foundations of Cryptography - Volume 2, Basic Applications. Cambridge University Press, Cambridge (2004)
32. Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof systems. SIAM J. Comput. **18**(1), 186–208 (1989)

33. Gordon, S.D., Katz, J., Vaikuntanathan, V.: A group signature scheme from lattice assumptions. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 395–412. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-17373-8_23
34. Groth, J.: Simulation-sound NIZK proofs for a practical language and constant size group signatures. In: Lai, X., Chen, K. (eds.) ASIACRYPT 2006. LNCS, vol. 4284, pp. 444–459. Springer, Heidelberg (2006). https://doi.org/10.1007/11935230_29
35. Ishai, Y., Kushilevitz, E., Ostrovsky, R., Sahai, A.: Zero-knowledge proofs from secure multiparty computation. *SIAM J. Comput.* **39**(3), 1121–1152 (2009)
36. Johnson, S., Scarlata, V., Rozas, C., Brickell, E., Mckeen, F.: Intel® software guard extensions: EPID provisioning and attestation services (2016)
37. Katz, J., Kolesnikov, V., Wang, X.: Improved non-interactive zero knowledge with applications to post-quantum signatures. *IACR Cryptology ePrint Archive* 2018:475 (2018)
38. Laurie, B., Langley, A., Kasper, E.: Certificate transparency. RFC 6962 (2013)
39. Libert, B., Ling, S., Mouhartem, F., Nguyen, K., Wang, H.: Signature schemes with efficient protocols and dynamic group signatures from lattice assumptions. In: Cheon, J.H., Takagi, T. (eds.) ASIACRYPT 2016. LNCS, vol. 10032, pp. 373–403. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53890-6_13
40. Libert, B., Ling, S., Nguyen, K., Wang, H.: Zero-knowledge arguments for lattice-based accumulators: logarithmic-size ring signatures and group signatures without trapdoors. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016. LNCS, vol. 9666, pp. 1–31. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49896-5_1
41. Ling, S., Nguyen, K., Wang, H.: Group signatures from lattices: simpler, tighter, shorter, ring-based. In: Katz, J. (ed.) PKC 2015. LNCS, vol. 9020, pp. 427–449. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46447-2_19
42. Ling, S., Nguyen, K., Wang, H., Xu, Y.: Lattice-based group signatures: achieving full dynamicity with ease. In: Gollmann, D., Miyaji, A., Kikuchi, H. (eds.) ACNS 2017. LNCS, vol. 10355, pp. 293–312. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-61204-1_15
43. Merkle, R.C.: A digital signature based on a conventional encryption function. In: Pomerance, C. (ed.) CRYPTO 1987. LNCS, vol. 293, pp. 369–378. Springer, Heidelberg (1988). https://doi.org/10.1007/3-540-48184-2_32
44. Naor, M., Yung, M.: Universal one-way hash functions and their cryptographic applications. In: STOC, pp. 33–43. ACM (1989)
45. Nayak, K., et al.: HOP: hardware makes obfuscation practical. In: NDSS (2017)
46. Sahai, A.: Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In: FOCS, pp. 543–553 (1999)
47. De Santis, A., Di Crescenzo, G., Ostrovsky, R., Persiano, G., Sahai, A.: Robust non-interactive zero knowledge. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 566–598. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-44647-8_33
48. Unruh, D.: Non-interactive zero-knowledge proofs in the quantum random oracle model. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9057, pp. 755–784. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46803-6_25
49. Unruh, D.: Post-quantum security of Fiat-Shamir. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017. LNCS, vol. 10624, pp. 65–95. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-70694-8_3

50. Yu, M., Gligor, V.D., Zhou, Z.: Trusted display on untrusted commodity platforms. In: CCS, pp. 989–1003 (2015)
51. Zheng, W., Dave, A., Beekman, J.G., Popa, R.A., Gonzalez, J.E., Stoica, I.: Opaque: an oblivious and encrypted distributed analytics platform. In: NSDI, pp. 283–298 (2017)
52. Zhou, Z., Gligor, V.D., Newsome, J., McCune, J.M.: Building verifiable trusted path on commodity x86 computers. In: IEEE Symposium on Security and Privacy (Oakland), pp. 616–630 (2012)