

Improved Non-Interactive Zero Knowledge with Applications to Post-Quantum Signatures

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

Vladimir Kolesnikov
Georgia Tech
kolesnikov@gatech.edu

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

May 18, 2018

Abstract

Recent work, including ZKBoo, ZKB++, and Ligerio, has developed efficient non-interactive zero-knowledge proofs of knowledge (NIZKPoKs) for arbitrary Boolean circuits based on symmetric-key primitives alone using the “MPC-in-the-head” paradigm of Ishai et al. We show how to instantiate this paradigm with MPC protocols in the *preprocessing* model; once optimized, this results in an NIZKPoK with shorter proofs (and comparable computation) as in prior work for circuits containing roughly 300–100,000 AND gates. In contrast to prior work, our NIZKPoK also supports witness-independent preprocessing, which allows the prover to move most of its work to an offline phase before the witness is known.

We use our NIZKPoK to construct a signature scheme based only on symmetric-key primitives (and hence with “post-quantum” security). The resulting scheme has shorter signatures than the scheme built using ZKB++ (with comparable signing/verification time), and is even competitive with hash-based signature schemes.

To further highlight the flexibility and power of our NIZKPoK, we also use it to build efficient ring and group signatures based on symmetric-key primitives alone. To our knowledge, the resulting schemes are the most efficient constructions of these primitives that offer post-quantum security.

1 Introduction

Zero-knowledge (ZK) proofs¹ [GMR85, GMW91] and non-interactive zero-knowledge (NIZK) proofs [BFM88, BDSMP91, FLS90] are among the most fundamental and versatile cryptographic primitives. In the present context, (NI)ZK proofs of knowledge (PoKs) [GMR85, BG93, DP92] allow a prover to convince a verifier, who holds a circuit C , that the prover knows an input (or “witness”) w for which $C(w) = 1$. Historically, research on (NI)ZK proofs has followed two largely independent paths: some work gave asymptotic improvements for *generic* proofs that can be used for arbitrary circuits (by reducing a circuit to an NP-complete problem) but had poor concrete efficiency; other work explored *efficient* proofs tailored to certain classes of circuits related to specific, “algebraic” computations (e.g., group exponentiation).

¹In this paper we do not distinguish between computational and information-theoretic soundness, and thus refer to both arguments and proofs simply as “proofs.”

More recently, researchers have merged these two directions and have developed (NI)ZKPoKs for general circuits with good concrete efficiency. We highlight in particular recent constructions of succinct non-interactive arguments of knowledge (SNARKs) [GGPR13, PHGR13, BCG⁺13, CFH⁺15, Gro16], which can be used to generate proofs for arbitrary circuits that are shorter than the input w itself. (Prior work showing ZK proofs with size sublinear in the input [Kil92, Mic94] served as a feasibility result, but was concretely inefficient.) Nevertheless, existing SNARKs suffer from several disadvantages. First, although they have excellent proof size and verification time, the computational efficiency of the prover is poor. Second, most constructions require public parameters generated and published by some semi-trusted party. Security of existing SNARKs is based on non-standard and poorly understood *knowledge assumptions*, and to some extent this is inherent [GW11]. Finally, most existing constructions of SNARKs rely on bilinear maps and are thus insecure against quantum attackers. (An exception is the recent work of Boneh et al. [BISW17]. This scheme still has the other drawbacks, and its concrete efficiency is unclear.)

In another recent line of work, researchers have shown how to use protocols for secure multi-party computation (MPC) to obtain (NI)ZKPoKs with excellent concrete performance (even if the communication is no longer sublinear in the circuit size). Jawurek et al. [JKO13] (see also [FNO15]) showed that two-party protocols based on garbled circuits could be used to construct efficient ZKPoKs; unfortunately, that approach appears to inherently require interaction as well as the use of public-key primitives. Ishai et al. (IKOS) [IKOS07], using the so-called “MPC-in-the-head” approach, showed how to use MPC protocols to obtain ZKPoKs whose security can be based entirely on symmetric-key primitives; their approach yields public-coin protocols, and so can be made non-interactive via the Fiat-Shamir transform [FS87]. Based on the IKOS approach, Giacomelli et al. [GMO16] developed and implemented a protocol called ZKBoo that supports efficient NIZKPoKs for arbitrary circuits. (Concurrently, Ranellucci et al. [RTZ16] proposed an NIZKPoK with similar asymptotic performance.) Chase et al. [CDG⁺17] introduced ZKB++, which improves the performance of ZKBoo; they also showed that ZKB++ could be used to construct an efficient signature scheme based on symmetric-key primitives alone. Ames et al. [AHIV17] subsequently proposed Liger, which has asymptotically shorter proofs than either ZKBoo or ZKB++. Further comparison is given in Section 3.1.

1.1 Our Contributions

We show here a novel way to instantiate the “MPC-in-the-head” approach that leads to NIZKPoKs with shorter proofs than (and comparable computation to) the best prior work [GMO16, RTZ16, CDG⁺17, AHIV17] for “medium-size” circuits with roughly 300–100,000 gates. As we show, this range covers many “cryptographically interesting” circuits that incorporate hashing, block-cipher evaluations, etc.

Our main insight leading to improved performance is that the “MPC-in-the-head” paradigm can be instantiated using MPC protocols *designed in the preprocessing model*. (A detailed overview of our approach is given in Section 1.2.) This greatly increases the space of MPC protocols that can be used, and in particular makes it competitive to use protocols designed for a large number of parties, in contrast to prior work, and hence achieve better efficiency for the same level of security.

As in the work of Chase et al. [CDG⁺17], we may apply our NIZKPoK to construct an efficient signature scheme based on symmetric-key primitives alone, and thus offering “post-quantum” security. Our resulting scheme has signatures roughly $3.2\times$ shorter than those of Chase et al.; in fact, the signature length of our scheme is even better than some leading constructions of (stateless)

hash-based signature schemes [BHH⁺15], with comparable signing/verification times. We can also trade off signature length and computation, giving some flexibility in the design.

As an additional application highlighting the power of our new NIZKPoK, we also show how to use it to build efficient ring and group signatures based on symmetric-key primitives alone.² The main challenge here is to build schemes in which signing corresponds to generating an NIZKPoK for a circuit that is not too large; we defer additional details to Section 4. Our resulting schemes are, to the best of our knowledge, the most efficient constructions of ring/group signatures that offer post-quantum security.

1.2 Overview of Our Techniques

We provide the high-level intuition for our improved NIZKPoK. Here (and in Section 2.2) we describe a 5-round, public-coin proof of knowledge that is honest-verifier zero knowledge (HVZK). In Section 2.4 we then show how it can be “compressed” to three rounds (while remaining public-coin and HVZK). Either of these proofs can, in turn, be compiled to achieve zero knowledge against a malicious verifier using standard techniques. Alternatively—and of most interest to us—the 3-round version can be made non-interactive in the random-oracle model using the Fiat-Shamir transform [FS87].

Our starting point is the “MPC-in-the-head” paradigm of Ishai et al. [IKOS07] for constructing public-coin HVZK proofs of knowledge based on protocols for secure computation. Roughly speaking, the main idea is for the prover to simulate an execution of an n -party secure-computation protocol Π that evaluates a given circuit C on the prover’s input w , and then for the prover to commit to the views of the parties in the protocol. The verifier chooses a random subset of those parties (we call those parties “opened”), and the prover then opens their corresponding views. The verifier then checks that the views of the opened parties are consistent with each other and with an honest execution of Π that yields output 1.

As observed in prior work [GMO16], efficiency considerations for the protocol Π used in the “MPC-in-the-head” paradigm differ from efficiency considerations for standard applications of MPC. For example, in this context Π can freely rely on “oblivious-transfer channels” in addition to standard point-to-point channels; in fact, Π can freely use any (deterministic) two-party functionality as a building block, since such functionalities incur no additional cost for the prover to simulate them, and consistency between any pair of opened parties can still be checked by the verifier. We stress that this does not immediately extend to using t -party functionalities for $t > 2$, since it may no longer be possible to verify consistent execution unless all t parties using the functionality are opened.

Giacomelli et al. [GMO16] and Chase et al. [CDG⁺17] gave concrete instantiations of the “MPC-in-the-head” approach, and determined that using a 3-party protocol Π resulted in optimal performance for the space of protocols they considered. Since the number of parties is small, the soundness of the resulting HVZK proof is poor; a large number of parallel repetitions is thus needed to obtain negligible soundness. (We discuss the work of Ames et al. [AHIV17] in Section 3.1.)

Our main conceptual insight is that the “MPC-in-the-head” approach can be modified so the underlying n -party protocol Π ($n > 2$) can rely on any n -party functionality, as long as that functionality is called during a preprocessing phase. (We stress that this preprocessing is virtual

²An earlier version of this work already containing these ideas was submitted to Eurocrypt 2018 in September 2017. Similar ideas, but using ZKB++, have since been pursued by other groups [DRS17, BEF18].

since everything is being simulated by the prover.) At first glance, using an n -party functionality does not seem compatible with the IKOS paradigm: the only way to verify that the functionality was run honestly is to open *all* parties, but then privacy (which translates to zero-knowledge) cannot possibly hold. We overcome this difficulty by having the verifier issue its challenges in *two* phases, thus giving a 5-round HVZK proof. Roughly, the HVZK proof thus works as follows:

1. The prover commits to multiple (independent) executions of the preprocessing phase. Each such commitment includes a separate commitment to the state of each party at the end of that phase.
2. The verifier challenges the prover to open some of the executions of the preprocessing phase. For each challenged execution, the prover opens the state of *all* parties, thus allowing the verifier to check those executions.
3. Using each of the remaining (unchallenged) executions of the preprocessing phase, the prover simulates an execution of Π that evaluates C on input w . It then commits to each party’s view in the protocol.
4. As in the original IKOS approach, for each execution of Π the verifier chooses a random subset of the parties to open. Now, however, opening a party involves opening the commitment to that party’s state following the preprocessing phase (from step 1) as well as its view in the online execution of the protocol (from step 3). The verifier then checks that all the opened parties are consistent with each other and with an honest execution of Π (using the state from the preprocessing phase) that yields output 1.

The fact that the n -party functionality is executed during a preprocessing phase is essential; the execution of that functionality is independent of the parties’ inputs, and so opening the views of all parties from the preprocessing phase does not reveal any information about w .

As a consequence of being able to rely on preprocessing, the space of possible protocols Π we can use is greatly expanded. We find that we obtain much shorter proofs by using an n -party protocol (secure against semi-honest corruption of all-but-one of the parties) with n as high as 64. The ability to rely on preprocessing is critical here: the communication complexity of traditional MPC protocols (that do not rely on preprocessing) with security against all-but-one corruption is quadratic in the number of parties, but by relying on preprocessing we can obtain communication complexity independent of n .

Further optimizations and specific parameter choices for the above proof are discussed in the remainder of the paper.

1.3 Outline of the Paper

In Section 2, we present the details of the MPC protocol Π we use, and describe more fully how to convert it into a 3-round, public-coin HVZK proof whose security relies on symmetric-key primitives alone. By applying the Fiat-Shamir transform to our 3-round proof, we obtain an NIZKPoK; the resulting construction is analyzed and compared to prior work in Section 3.1, where we show that it yields shorter proofs than prior work for a large range of circuit sizes. As in prior work, we can use our NIZKPoK to build a signature scheme based only on symmetric-key primitives (and hence offering “post-quantum” security); performance of this scheme is discussed in Section 3.2. We explore the applicability of our work to constructing “post-quantum” ring and group signatures in Section 4.

2 HVZK Proofs of Knowledge

We construct our HVZK proof by instantiating the “MPC-in-the-head” paradigm using an MPC protocol Π designed in the preprocessing model. We first describe the MPC protocol we rely on, and then discuss how to build an HVZK proof based on it.

Notation. We denote the n -out-of- n (XOR-based) secret sharing of a bit x by $[x]$. We let $|C|$ denote the number of AND gates in a circuit C . Although in this paper we focus on Boolean circuits, our ideas can be extended to arithmetic circuits in a straightforward manner. Throughout, κ and ρ denote computational and statistical security parameters, respectively.

2.1 A Suitable MPC protocol

In this section we describe an n -party protocol Π run by parties S_1, \dots, S_n in the preprocessing model, secure against semi-honest corruption of all-but-one of the parties. The protocol is inspired by recent work of Wang et al. [WRK17].

Recall that our goal is to use this protocol to instantiate the “MPC-in-the-head” approach efficiently, and so in particular we are interested in minimizing the communication needed to verify consistency and correct computation for any $n - 1$ parties chosen by the verifier. As explained in more detail in the next section, our protocol Π ends up having communication comparable to the 3-party MPC protocols used by ZKBoo and ZKB++, but because we set $n > 3$ we are able to achieve much lower soundness error per protocol execution and thus better communication complexity overall for our resulting interactive proof system.

Our protocol Π maintains the invariant that, for each wire in the circuit, the parties hold an n -out-of- n secret sharing of a random mask along with the (public) masked value of the wire. Specifically, if we let z_α denote the value of wire α in the circuit C when evaluated on input w , then the parties will hold $[\lambda_\alpha]$ (for uniform $\lambda_\alpha \in \{0, 1\}$) along with the value $\hat{z}_\alpha \stackrel{\text{def}}{=} z_\alpha \oplus \lambda_\alpha$.

Preprocessing phase. In the preprocessing phase, shares are set up among the parties as follows. For each wire α that is either an input wire or the output wire of an AND gate, the parties are given $[\lambda_\alpha]$, where $\lambda_\alpha \in \{0, 1\}$ is uniform. For an XOR gate with input wires α, β and output wire γ , define $\lambda_\gamma \stackrel{\text{def}}{=} \lambda_\alpha \oplus \lambda_\beta$; note the parties can compute $[\lambda_\gamma]$ locally. Finally, for each AND gate with input wires α, β , the parties are given $[\lambda_{\alpha, \beta}]$, where $\lambda_{\alpha, \beta} \stackrel{\text{def}}{=} \lambda_\alpha \cdot \lambda_\beta$.

A key observation is that the shares of the $\{\lambda_\alpha\}$ are uniform, and so can be generated by having each party S_i apply a pseudorandom generator to a short, random seed seed_i given to that party, and then (implicitly) defining the $\{\lambda_\alpha\}$ based on the resulting shares. All-but-one of the shares of the $\{\lambda_{\alpha, \beta}\}$ can also be generated in this way, but the final share is constrained by the values of $\lambda_\alpha, \lambda_\beta$. To ensure that the shares of the $\{\lambda_{\alpha, \beta}\}$ are correct, S_n can be given an additional $|C|$ “correction bits” that indicate its share of $\lambda_{\alpha, \beta}$ for each AND gate with input wires α, β .

To summarize: each S_i is given a κ -bit seed $\text{seed}_i \in \{0, 1\}^\kappa$, and S_n is additionally given $|C|$ bits denoted by aux_n . We refer to this information as the *state* of the parties, and denote the state of S_i by state_i . In the online phase of the protocol, each party S_i uses seed_i to generate its shares of the $\{\lambda_\alpha\}$; for $1 \leq i \leq n - 1$, party S_i also uses seed_i to generate its shares of the $\{\lambda_{\alpha, \beta}\}$. Party S_n uses aux_n to determine its shares of the $\{\lambda_{\alpha, \beta}\}$.

Protocol execution. In the description that follows, we rely on a simple subroutine we call *share reconstruction* that allows the parties to publicly reconstruct a shared value. In our context, where

all parties are semi-honest, we perform share reconstruction of $[x]$ by simply having each party broadcast its share of x .

We assume the parties begin the protocol holding a masked value \hat{z}_α for each input wire α . (In our context these masked values will be provided to the parties by the prover who is simulating execution of the protocol.) Note that these masked values, along with the corresponding $\{[\lambda_\alpha]\}$, define an effective input to the protocol. During the online phase of the protocol, the parties inductively compute \hat{z}_α for all wires in the circuit. Specifically, for each gate of the circuit with input wires α, β and output wire γ , where the parties already hold $\hat{z}_\alpha, \hat{z}_\beta$, the parties do:

- If the gate is an XOR gate, then the parties locally compute $\hat{z}_\gamma := \hat{z}_\alpha \oplus \hat{z}_\beta$.
- If the gate is an AND gate, the parties locally compute

$$[s] := \hat{z}_\alpha[\lambda_\beta] \oplus \hat{z}_\beta[\lambda_\alpha] \oplus [\lambda_{\alpha,\beta}] \oplus [\lambda_\gamma],$$

and then publicly reconstruct s . Finally, they compute $\hat{z}_\gamma := s \oplus \hat{z}_\alpha \hat{z}_\beta$. One can verify that $\hat{z}_\gamma = z_\gamma \oplus \lambda_\gamma$.

Once the parties have computed \hat{z}_α for the output wire α , the output value z_α is computed by reconstructing λ_α and then setting $z_\alpha := \hat{z}_\alpha \oplus \lambda_\alpha$.

We remark that the online phase of this protocol is deterministic. Also observe that all communication is due to share reconstruction: for a circuit with $|C|$ AND gates, $|C|$ share reconstructions are needed. We will see in the following sections that these properties are advantageous for our application.

2.2 Our Basic HVZK Proof

In this section we describe a basic, 5-round, HVZK proof based on the MPC protocol from the previous section. We defer discussion of further optimizations to the following section, and a description of a 3-round variant to Section 2.4.

Notation. In this section, H is a collision-resistant hash function. (Later in the paper, we will model H as a random oracle.) We let Com denote a non-interactive³ commitment scheme, where a commitment to a value x is generated by choosing a uniform $r \in \{0, 1\}^\kappa$ and computing $\text{com} := \text{Com}(x; r)$; decommitment is done by simply revealing x and r . Computational binding and hiding suffices.

Overview. Recall from Section 1.2 that, because we rely on an MPC protocol Π designed in the preprocessing model, we modify the “MPC-in-the-head” approach and construct a 5-round proof of the following format: First, the prover commits to m executions of the preprocessing phase, where each such commitment involves a commitment to the state of each party. The verifier then challenges the prover to open all-but-one of those executions, and verifies that all the state information given to the parties in the opened executions is correct. The remaining (unopened) execution of the preprocessing phase is then used by the prover to simulate the online phase of the n -party protocol Π on a shared input w , and the prover commits to the view of each party in this execution. The verifier then chooses all-but-one of the parties to be opened; for any chosen party, the prover opens that party’s state from the preprocessing phase as well as its view from the online

³We could also use a two-round commitment scheme by adding one additional round to the proof system.

An HVZK proof

Inputs: Both parties have a circuit C ; the prover also holds w with $C(w) = 1$. Values m, n are parameters of the protocol.

Round 1 For each $j \in [m]$, the prover emulates the preprocessing phase as follows:

1. Choose uniform $\text{seed}_j^* \in \{0, 1\}^\kappa$ and use it to generate values $\text{seed}_{j,1}, r_{j,1}, \dots, \text{seed}_{j,n}, r_{j,n}$. Also compute $\text{aux}_{j,n} \in \{0, 1\}^{|C|}$ as described in the text. For $i = 1, \dots, n-1$ let $\text{state}_{j,i} := \text{seed}_{j,i}$; let $\text{state}_{j,n} := \text{seed}_{j,n} \parallel \text{aux}_{j,n}$.
2. For $i \in [n]$, compute $\text{com}_{j,i} := \text{Com}(\text{state}_{j,i}; r_{j,i})$.
3. Compute $h_j := H(\text{com}_{j,1}, \dots, \text{com}_{j,n})$.

Compute $h := H(h_1, \dots, h_m)$ and send h to the verifier.

Round 2 The verifier chooses uniform $c \in [m]$ and sends it to the prover.

Round 3 The prover sends $\{\text{seed}_j^*\}_{j \neq c}$ to the verifier.

The prover simulates the online phase of the n -party protocol Π (as described in the text) using the state generated by the c th preprocessing phase, beginning by computing the masked values $\{\hat{z}_\alpha\}$ for the input wires (based on w and the $\{\lambda_\alpha\}$ defined by the c th preprocessing phase). Let msgs_i denote the messages broadcast by S_i in the protocol execution. The prover sends the initial masked values $\{\hat{z}_\alpha\}$ as well as $h' := H(\text{msgs}_1, \dots, \text{msgs}_n)$ to the verifier.

Round 4 The verifier chooses uniform $p \in [n]$ and sends it to the prover.

Round 5 The prover sends $\{\text{state}_{c,i}, r_{c,i}\}_{i \neq p}$, $\text{com}_{c,p}$, and msgs_p .

Verification The verifier accepts iff all the following checks succeed:

1. For $i \neq p$, the verifier uses $\text{state}_{c,i}$ and $r_{c,i}$ to compute $\text{com}_{c,i}$. It then computes $h_c := H(\text{com}_{c,1}, \dots, \text{com}_{c,n})$.
2. For $j \neq c$, the verifier uses seed_j^* to compute h_j as an honest prover would. It then checks that $h \stackrel{?}{=} H(h_1, \dots, h_m)$.
3. The verifier simulates the online phase of Π among the $\{S_i\}_{i \neq p}$ using $\{\text{state}_{c,i}\}_{i \neq p}$, masked input-wire values $\{\hat{z}_\alpha\}$, and msgs_p . This yields $\{\text{msgs}_i\}_{i \neq p}$ and an output bit b . The verifier checks that $b \stackrel{?}{=} 1$.
4. The verifier checks that $h' \stackrel{?}{=} H(\text{msgs}_1, \dots, \text{msgs}_n)$.

Figure 1: Our basic HVZK proof.

phase. The verifier then checks consistency of the parties' views as well as correct behavior. We provide details about each of these steps next.

Checking the preprocessing phase. Recall from the previous section that, following the preprocessing phase, the state of party S_i for $1 \leq i \leq n-1$ is a seed seed_i , while the state of party S_n is a seed seed_n along with a $|C|$ -bit string aux_n . Thus, in a naive implementation of the idea outlined above, for each execution of the preprocessing phase the prover would send $(\text{com}_1, \dots, \text{com}_n)$, where com_i is a commitment to the state of S_i . Then, for each challenged execution of the preprocessing phase, the prover would send $\text{seed}_1, \dots, \text{seed}_n, \text{aux}_n, r_1, \dots, r_n$, where r_i denotes the randomness used for the i th commitment. For the unchallenged execution (used for the subsequent execution of Π), the prover will be able to decommit to the states of any subset of the parties. The communication complexity of this approach is $O(m|C| + mn\kappa)$.

We improve the communication complexity in several ways (cf. Round 1 and steps (1) and (2) of verification in Figure 1):

1. The prover will send $H(\text{com}_1, \dots, \text{com}_n)$ instead of sending all the commitments individually. It can in fact hash each of those results and send just a single hash value to the verifier.
2. When opening a challenged execution, it is unnecessary for the prover to send aux_n since the correct value of aux_n can be computed from $\text{seed}_1, \dots, \text{seed}_n$. This results in an asymptotic improvement in the communication.
3. By generating the $\{\text{seed}_i\}$ and the $\{r_i\}$ from a “master” seed $\text{seed}^* \in \{0, 1\}^\kappa$, the prover can open a challenged execution of the preprocessing phase by simply sending seed^* .

The overall communication complexity of the initial checking phase can thus be reduced to $O(m\kappa)$, which is independent of $|C|$ and n .

Checking the online execution. An execution of our protocol proceeds gate-by-gate, with the processing of each AND gate requiring reconstruction of one shared value. Although the communication complexity of share reconstruction in our protocol is n bits (one bit per party), for our purposes we do not need the prover to send n bits per gate in order to prove consistent execution. This is because the verifier only needs to obtain the protocol messages sent by the (single) unopened party in order to check the execution of the $n-1$ opened parties. Thus, it suffices for the prover to send just a single bit per AND gate.

In addition to the protocol messages sent by the unopened party, the prover also needs to reveal the state (from the preprocessing phase) of every opened party. For each opened party S_i , $i \neq n$, this involves just $O(\kappa)$ bits; if S_n is opened then this requires $|C| + O(\kappa)$ bits due to aux_n . In either case the marginal communication complexity per AND gate is *independent* of the number of parties n .

The above ideas (along with a few small optimizations) are encapsulated in Figure 1. Security of the scheme is given by the following theorem.

Theorem 2.1. *Let H be collision-resistant, and let Com be a secure commitment scheme. Then the protocol in Figure 1 is an honest-verifier zero-knowledge proof of knowledge with soundness/knowledge error $\max\{\frac{1}{m}, \frac{1}{n}\}$.*

Proof. Completeness of the protocol is immediate, and we thus focus on proving honest-verifier zero knowledge and soundness.

Honest-verifier zero knowledge follows easily from semi-honest security of Π (for corruption of any $n-1$ parties) and hiding of the commitment scheme. Specifically, if we let Sim_Π denote a

simulator for Π , then we can construct an honest-verifier zero-knowledge simulator for our protocol as follows:

- Choose uniform $c \in [m]$ and $p \in [n]$.
- Run Sim_Π to simulate the views of parties $\{S_i\}_{i \neq p}$ in an execution of Π when evaluating C with output 1. This results in values $\{\text{state}_i\}_{i \neq p}$, masked input-wire values $\{\hat{z}_\alpha\}$, and msgs_p . From the respective views, $\{\text{msgs}_i\}_{i \neq p}$ can be computed. The simulator computes h' as the honest prover would.
- For $j \neq c$, the simulator computes h_j exactly as the honest prover would. For $i \neq p$, the simulator sets $\text{state}_{c,i} := \text{state}_i$, chooses uniform $r_{c,i}$, and computes $\text{com}_{c,i}$ as the honest prover would. It computes $\text{com}_{c,p}$ as a commitment to a 0-string of the appropriate length. It then computes h_c and h as the honest prover would.
- The simulator outputs the transcript with h , c , $\{\text{seed}_j^*\}_{j \neq c}$, $\{\hat{z}_\alpha\}$, h' , p , $\{\text{state}_{c,i}, r_{c,i}\}_{i \neq p}$, $\text{com}_{c,p}$, and msgs_p .

A straightforward hybrid argument shows that transcripts output by the simulator are computationally indistinguishable from transcripts of real executions of the protocol with an honest verifier.

We next show that given accepting transcripts for challenges (c, p) , (c', \star) , and (c, p') , with $c \neq c'$ and $p \neq p'$, it is possible to extract a witness w for which $C(w) = 1$ unless there is a hash collision or binding of the commitment scheme fails. It follows using standard techniques that the protocol has soundness/knowledge error $\max\{\frac{1}{m}, \frac{1}{n}\}$.

In the discussion that follows we assume for simplicity that no hash collisions occur and that binding of the commitment scheme never fails. From the accepting transcript with challenges (c', \star) , we can compute for any $c \neq c'$ values $\text{state}_{c,1}, \dots, \text{state}_{c,n}$ that are mutually consistent (i.e., are a correct outcome of the preprocessing phase); moreover, those values are equal to the corresponding values sent in round 5 of the other two accepting transcripts (with initial challenge c). These values of the state define a value λ_α for each input wire α ; from these and the masked inputs $\{\hat{z}_\alpha\}$ used in either of the other accepting transcripts we may compute an effective input w .

We now show that $C(w) = 1$:

- From the transcript with challenges (c, p) , we obtain the state of all parties besides S_p ; these will be exactly equal to the values $\{\text{state}_{c,i}\}_{i \neq p}$ computed above. We also derive $\{\text{msgs}_i\}_{i \in [n]}$. Moreover, these values are all consistent with an honest execution of Π by parties $\{S_i\}_{i \neq p}$ resulting in output 1.
- From the transcript with challenges (c, p') , we can obtain $\{\text{state}_{c,i}\}_{i \neq p'}$ along with protocol messages $\{\text{msgs}_i\}_{i \in [n]}$ that must be equal to those obtained above. These values are consistent with an honest execution of Π by parties $\{S_i\}_{i \neq p'}$ resulting in output 1.

Since the masked input-wire values $\{\hat{z}_\alpha\}$ are fixed in the above, the $\{\text{state}_{c,i}\}_{i \in [n]}$ and the $\{\text{msgs}_i\}_{i \in [n]}$ must all be mutually consistent with an honest execution of Π , beginning with correct state and effective input w , resulting in output 1. This implies $C(w) = 1$. \square

n	$\rho = 128$						$\rho = 256$					
	4	8	16	32	64	128	4	8	16	32	64	128
M	218	252	352	462	631	916	456	533	781	1024	1662	2540
τ	65	44	33	27	23	20	129	87	65	53	44	38

Table 1: Sample values of n , M , and τ to achieve statistical security $\rho \in \{128, 256\}$.

2.3 Additional Optimizations

Here we describe some additional optimizations that can be used to further improve the concrete performance of our protocol.

Improved commitments. In the basic protocol we use a commitment scheme to hide the states of the parties following the preprocessing phase. If we are willing to model H as a random oracle, then we can streamline things by observing that all committed values have min-entropy at least κ and can thus be committed by direct evaluation of H (i.e., we may set $\text{com}_{j,i} := H(\text{state}_{j,i})$). This reduces communication since no additional randomness is needed for decommitment.

Reducing the number of random seeds. In the c th emulation of the preprocessing phase, the prover generates n seeds $\text{seed}_{c,1}, \dots, \text{seed}_{c,n}$ from a master seed seed_c^* , commits to the n generated seeds, and then reveals $n - 1$ of those seeds to the verifier. The second step requires $(n - 1) \cdot \kappa$ bits of communication.

Motivated by the NNL scheme for stateless revocation [NNL01], we observe that we can reduce the communication by generating the seeds in a more structured way. Namely, imagine labeling the root of a binary tree of depth $\log n$ with seed_c^* , and then inductively labeling the children of each node with the output of a pseudorandom generator applied to the node's label. The $\{\text{seed}_{c,i}\}_{i \in [n]}$ will be the labels of the n leaves of the tree. To reveal $\{\text{seed}_{c,i}\}_{i \neq p}$, it suffices to reveal the labels on the siblings of the path from the root of the tree to leaf p . Those labels allow the verifier to reconstruct $\{\text{seed}_{c,i}\}_{i \neq p}$ while still hiding $\text{seed}_{c,p}$. Applying this optimization reduces the communication complexity to $\kappa \cdot \log n$ for revealing the seeds used by the $n - 1$ opened parties.

We can, in fact, apply the same idea to the master seeds $\{\text{seed}_j^*\}_{j=1}^m$ used for the different emulations of the preprocessing phase; this reduces the communication required to reveal all-but-one of those seeds in round 3 from $(m - 1) \cdot \kappa$ bits to $\kappa \cdot \log m$ bits. Looking ahead to the next optimization, we remark that we are not limited to revealing all-but-one of the leaf labels; more generally, the scheme just described supports revealing all-but- τ of the leaf labels using communication at most $\kappa \cdot \tau \log \frac{m}{\tau}$ bits (cf. [NNL01]).

Beating parallel repetition. The basic protocol analyzed in the previous section has soundness error $\max\{\frac{1}{m}, \frac{1}{n}\}$; we can achieve soundness error $2^{-\rho}$ by running $\tau = \frac{\rho}{\log \min\{m,n\}}$ independent, parallel repetitions of the basic protocol. This would require τm emulations of the preprocessing phase, followed by τ executions of an n -party protocol.

We can do better by performing a more general cut-and-choose over the emulations of the preprocessing phase. That is, the prover now runs M emulations of the preprocessing phase and the verifier selects $M - \tau$ of those to check; the remaining τ executions of the preprocessing phase are then used to run τ (independent) instances of Π . Each of those instances of Π is then verified by revealing the view of all-but-one party as before. If a cheating prover generates $k \geq M - \tau$

A 3-round HVZK proof

Inputs: Both parties have a circuit C ; the prover also holds w with $C(w) = 1$. Values M, n, τ are parameters of the protocol.

Round 1 For each $j \in [M]$, the prover does:

1. Choose uniform $\text{seed}_j^* \in \{0, 1\}^\kappa$ and use it to generate values $\text{seed}_{j,1}, r_{j,1}, \dots, \text{seed}_{j,n}, r_{j,n}$. Also compute $\text{aux}_j \in \{0, 1\}^{|C|}$ as described in the text. For $i = 1, \dots, n-1$, let $\text{state}_{j,i} := \text{seed}_{j,i}$; let $\text{state}_{j,n} := \text{seed}_{j,n} \parallel \text{aux}_j$.
2. For $i \in [n]$, compute $\text{com}_{j,i} := \text{Com}(\text{state}_{j,i}; r_{j,i})$.
3. The prover simulates the online phase of the n -party protocol Π (as described in the text) using $\{\text{state}_{j,i}\}_i$, beginning by computing the masked inputs $\{\hat{z}_{j,\alpha}\}$ (based on w and the $\{\lambda_{j,\alpha}\}$ defined by the preprocessing). Let $\text{msgs}_{j,i}$ denote the messages broadcast by S_i in this protocol execution.
4. Let $h_j := H(\text{com}_{j,1}, \dots, \text{com}_{j,n})$ and $h'_j := H(\{\hat{z}_{j,\alpha}\}, \text{msgs}_{j,1}, \dots, \text{msgs}_{j,n})$.

Compute $h := H(h_1, \dots, h_M)$ and $h' := H(h'_1, \dots, h'_M)$ and send $h^* := H(h, h')$ to the verifier.

Round 2 The verifier chooses a uniform τ -sized set $\mathcal{C} \subset [M]$ and $\mathcal{P} = \{p_j\}_{j \in [\tau]}$ where each $p_j \in [n]$ is uniform. Send $(\mathcal{C}, \mathcal{P})$ to the prover.

Round 3 For each $j \in [M] \setminus \mathcal{C}$, the prover sends seed_j^*, h'_j .

For each $j \in \mathcal{C}$, the prover sends $\{\text{state}_{j,i}, r_{j,i}\}_{i \neq p_j}$, com_{j,p_j} , $\{\hat{z}_{j,\alpha}\}$, and msgs_{j,p_j} to the verifier.

Verification The verifier accepts iff all the following checks succeed:

1. For every $j \in \mathcal{C}$, $i \neq p_j$, the verifier uses $\text{state}_{j,i}$ and $r_{j,i}$ to compute $\text{com}_{j,i}$. It then computes $h_j := H(\text{com}_{j,1}, \dots, \text{com}_{j,n})$.
2. For $j \in [M] \setminus \mathcal{C}$, the verifier uses seed_j^* to compute h_j as an honest prover would. It then computes $h := H(h_1, \dots, h_M)$.
3. For each $j \in \mathcal{C}$, the verifier simulates an execution of Π among the $\{S_i\}_{i \neq p_j}$ using $\{\text{state}_{j,i}\}_{i \neq p_j}$, masked input-wire values $\{\hat{z}_{j,\alpha}\}$, and msgs_{j,p_j} . This yields $\{\text{msgs}_i\}_{i \neq p_j}$ and an output bit b . The verifier checks that $b \stackrel{?}{=} 1$ and computes $h'_j := H(\{\hat{z}_{j,\alpha}\}, \text{msgs}_{j,1}, \dots, \text{msgs}_{j,n})$ as well as $h' := H(h'_1, \dots, h'_m)$.
4. The verifier checks that $H(h, h') \stackrel{?}{=} h^*$.

Figure 2: Our 3-round HVZK proof.

correct emulations of the preprocessing phase and $M - k$ incorrect emulations, then its probability of successfully passing the first phase of the proof is at most $\binom{k}{M-\tau} \cdot \binom{M}{M-\tau}^{-1}$; conditioned on

passing the first phase, its probability of passing the second phase is at most $1/n^{k-M+\tau}$. The soundness error is therefore

$$\epsilon(M, n, \tau) \stackrel{\text{def}}{=} \max_{M-\tau \leq k \leq M} \left\{ \frac{\binom{k}{M-\tau}}{\binom{M}{M-\tau} \cdot n^{k-M+\tau}} \right\}. \quad (1)$$

In Table 1 we show some values of M , n , and τ that give soundness 2^{-128} or 2^{-256} . (For an interactive protocol, soundness 2^{-40} may be sufficient; however, for our later applications to non-interactive protocols better soundness is needed.)

2.4 A 3-Round Protocol

For applications where an interactive proof suffices, the 5-round protocol described previously can be used. If a non-interactive proof is needed, however, then it is easier to start with a 3-round protocol. (Although it is possible to apply the Fiat-Shamir transform to any *constant-round*, public-coin, HVZK protocol [AABN02], we obtain better efficiency for a given security level by starting with a 3-round protocol.) We present a 3-round version of our proof here.

The high-level idea of our 3-round variant is to have the prover now simulate the online phase of Π for *every* emulation of the preprocessing phase, and then commit to all the resulting executions. The verifier then challenges the prover to open some of the preprocessing phases, as well as all-but-one party from each of the unopened preprocessing phases, as before. The modified protocol is shown in Figure 2. In that figure, we also incorporate the cut-and-choose optimization discussed in the previous section. A proof of the following theorem relies on the same ideas as in the proof of Theorem 2.1, and so is omitted.

Theorem 2.2. *Let H be collision-resistant, and let Com be a secure commitment scheme. Then the protocol in Figure 2 is an honest-verifier zero-knowledge proof of knowledge with soundness/knowledge error $\epsilon(M, n, \tau)$.*

As mentioned in the previous section, instead of computing $h' := H(h'_1, \dots, h'_M)$ the prover can compute h' as the root of a Merkle tree with the $\{h'_j\}$ as leaves. Then, in the third round it suffices for the prover to send at most $\tau \log \frac{M}{\tau}$ values instead of sending the $M - \tau$ values $\{h'_j\}_{j \notin \mathcal{C}}$. Applying this optimization, and assuming the output of H is 2κ bits long, the communication complexity of our protocol is at most

$$2\kappa + \tau \cdot \log \frac{M}{\tau} \cdot 3\kappa + \tau \cdot (\kappa \log n + 2|C| + |w| + 2\kappa)$$

bits.

3 Removing Interaction

In this section we describe how our 3-round protocol can be converted into a general NIZKPoK as well as a signature scheme. We also report on implementations of the resulting schemes. All experiments (here and in Section 4) were carried out using a single core of an Intel Xeon E5-2666v3 CPU running at 2.9 GHz.

3.1 An NIZKPoK

We can apply the Fiat-Shamir transform [FS87] to our 3-round protocol to obtain a *non-interactive* ZKPoK. In more detail, in this modified proof system the prover will compute the first-round message as before, and then derive a challenge by hashing the first message using a hash function G modeled as a random oracle. The proof then consists of the first-round message plus the response to that challenge; it can be verified in the standard way.

Zero-knowledge of the resulting non-interactive proof system (in the random-oracle model) follows immediately from the fact that the 3-round interactive protocol is honest-verifier zero knowledge. Soundness—or, equivalently, the fact that the proof system is a proof of knowledge—can be proved using standard results about the Fiat-Shamir transform [PS00]. Directly applying those results unfortunately results in a relatively loose concrete-security bound due to use of the “forking lemma.” However, we observe that if H is also modeled as a random oracle, and is additionally used to instantiate the commitment scheme as discussed in Section 2.3, then it is possible to extract (with high probability) a valid witness w from a successful forgery without any rewinding at all, and thus obtain a tight security reduction. In a bit more detail: consider some attacker \mathcal{A} making q_G queries to the hash function G . Each time \mathcal{A} makes a query of the form $G(h^*)$, where h^* represents a first-round message from the 3-round protocol, the previous H -queries made by \mathcal{A} define⁴ values corresponding to M executions of the entire protocol Π (i.e., both the preprocessing phase and the online phase). If any of those M executions is “correct,” then that execution defines a value w such that $C(w) = 1$, i.e., a witness is extracted. Otherwise, the probability that the response to that G -query allows \mathcal{A} to generate a valid proof is bounded by $\epsilon(M, n, \tau)$. The overall probability that \mathcal{A} can generate a valid proof, yet a witness cannot be extracted, is thus dominated by $q_G \cdot \epsilon(M, n, \tau)$. (A full analysis would also take into account the negligible-probability events from footnote 4. Details will be given in the final version.)

We remark that Chase et al. do not claim tight security for the non-interactive version of ZKB++, and the observation above has, to the best of our knowledge, not been made previously in the context of efficient NIZKPoKs.

Security against quantum algorithms. The preceding analysis assumes a classical attacker (as in the work of Chase et al. [CDG⁺17]). We are not aware of any quantum polynomial-time attacks against NIZKPoKs derived using the Fiat-Shamir transform (assuming no quantum attacks against the underlying interactive protocol), and it is in this sense that we claim “post-quantum” security for our NIZKPoK. Nevertheless, we note that there are also no known proofs of security for the Fiat-Shamir transform in the quantum setting, except in some special cases [Unr17, CHR⁺18].

It is possible to apply Unruh’s transformation [Unr15] to our 3-round proof system to obtain a non-interactive scheme with provable security against quantum attacks (in the quantum random-oracle model), but this would significantly impact the efficiency. We leave open the question of obtaining an efficient non-interactive version of our protocol (and/or the scheme of Chase et al. [CDG⁺17]) with provable security against quantum attackers.

Implementation optimizations. One potential source of overhead in our scheme is the large number of parties in our underlying MPC protocol. To reduce the impact of this, we utilize bit-level parallelization by packing all parties’ shares into one word. This is efficient because most of

⁴It is possible that \mathcal{A} has not made H -queries defining those values, but in that case it is easy to see that the probability \mathcal{A} is able to use the response from the G -query to generate a valid proof is negligible. Alternately, a hash collision would mean that the values are ill-defined; such collisions occur with negligible probability.

Circuit size:	1,000			10,000			100,000		
	$ \sigma $	t_{total}	t_{online}	$ \sigma $	t_{total}	t_{online}	$ \sigma $	t_{total}	t_{online}
$n = 64$	37 KB	128 ms	10 ms	136 KB	851 ms	220 ms	1126 KB	7865 ms	2207 ms
$n = 32$	39 KB	51 ms	7 ms	159 KB	420 ms	149 ms	1351 KB	4015 ms	1498 ms
$n = 16$	44 KB	28 ms	7 ms	190 KB	250 ms	105 ms	1652 KB	2408 ms	1039 ms
$n = 8$	50 KB	16 ms	4 ms	246KB	190 ms	87 ms	2203 KB	1886 ms	883 ms

Table 2: Performance of our NIZKPoK. Parameters of all schemes were set to achieve 128-bit quantum security, as discussed in the text. $|\sigma|$ is the proof size; t_{total} is the total computation time; t_{online} is the online time.

the computations on parties’ shares are identical, and therefore SIMD-friendly. This optimization works regardless of the topology of the circuit C .

Note, however, that each party’s share is generated from a different seed, and so that step is not compatible with bit-packing. To address this, our implementation uses efficient bit-matrix transposition based on SSE instructions [WMK17, KOS16]: namely, we first generate a matrix of random bits contiguously aligned in memory, where each row is generated from one random seed, and then apply bit-level matrix transposition such that each column is a word containing one share from each party. Now, operating on shares from all parties takes only one instruction.

Our HVZK proofs use an MPC protocol that is designed in the preprocessing model. Interestingly, we can leverage this feature in our resulting proof system so as to allow the prover to perform most of the computation before the witness is even known. Specifically, the prover can compute h (cf. Figure 2) without knowledge of w . We believe this feature can be useful in many applications, where minimizing online time is crucial. Note that for other MPC-in-the-head protocols not based on preprocessing, it is not possible to move a lot of computation before knowing the witness, since the prover cannot emulate the MPC protocol without the witness.

Computation/communication trade off. We implemented our protocol in C++ and will make the code available as open-source. In Table 2 we show the proof size and running time of the prover for circuits of various sizes. (A comparison to prior work is given below.) For computation time, we show both the total time (t_{total}) to compute a proof, as well as the online time (t_{online}), assuming all witness-independent work is done by the prover during an offline phase. The running time of the verifier is always roughly the same as the (total) running time of the prover, and so we do not report it. In all cases we set the parameters of the protocols so as to achieve 128-bit security against known quantum attacks; in particular, we use SHA-256 as the hash function and AES-256 for pseudorandom number generation, and set M, n, τ so that $\epsilon(M, n, \tau) \leq 2^{-256}$.

Interestingly, our protocol offers the ability to trade off communication and computation by varying n . We show the effect of varying n in Table 2. By increasing n from 8 to 64, we are able to reduce the proof size by half; however, this results in a $4\times$ slowdown in the computation. Nevertheless, there may be applications that prefer to optimize proof length vs. computation time, or vice versa.

Comparison to prior work. We compare the performance of our scheme to the most relevant prior work (namely, ZKB++ and Liger). As in those works, we view the size of the resulting proofs as the primary metric of interest and so focus on that.

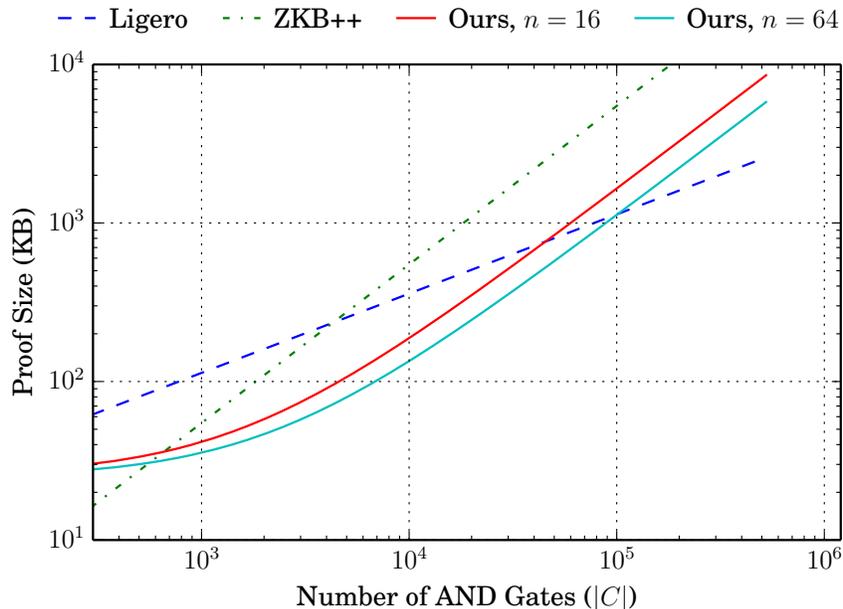


Figure 3: Proof size of ZKB++, Ligerio, and our protocols. Observe that our protocol yields the smallest proof size for circuits containing 300–100,000 AND gates.

Scheme	$ \sigma $	Gen	Sign	Vrfy	$ \text{sk} $	$ \text{pk} $	Assumption
This paper ($n = 64$)	38.5 KB	0.01 ms	129 ms	129 ms	32 B	64 B	ROM
This paper ($n = 16$)	45.9 KB	0.01 ms	28 ms	28 ms	32 B	64 B	ROM
ZKB++ [CDG ⁺ 17]	118.5 KB	0.01 ms	25 ms	17 ms	32 B	64 B	ROM
SPHINCS-256 [BHH ⁺ 15]	41 KB	0.82 ms	13 ms	0.58 ms	1088 B	1056 B	Collision-resistant hashing

Table 3: Performance of signature schemes with 128-bit quantum security. We report times for key generation, signing, and verification, as well as the lengths of the signature, private key, and public key.

We remark that neither ZKB++ nor Ligerio is available as an open-source implementation for experimentation: ZKB++ is only available as a signature scheme (cf. the next section) rather than as a full-fledged NIZKPoK, and code for Ligerio is not available. We calculate communication complexity based on equations provided in their papers, but are unable to report on the computational time required by those other protocols.

For a circuit C having $|C|$ AND gates, the proofs in our NIZKPoK have length $O(\kappa \cdot |C| / \log n)$ whereas those in ZKB++ have length $O(\kappa \cdot |C|)$. Since in our scheme n can vary (in contrast to ZKB++ which fixes $n = 3$), we can obtain asymptotically shorter proofs. More importantly, as discussed below, we obtain concretely shorter proofs than in ZKB++ for medium-size circuits even when we fix $n \in \{16, 64\}$ as we do in our experiments.

The communication complexity of our scheme is asymptotically worse than that of Ligerio, which has proofs of size $O(\kappa \cdot \sqrt{|C| \log |C|})$. However, the computational complexity of the prover

in Ligerio is $O(|C| \log^2 |C|)$, which is significant for large circuits. In contrast, the computational complexity of the prover in our protocol (after applying all the optimizations described in the text) is $O(|C|)$. Concretely, we note that in our protocol all operations can use hardware acceleration, whereas Ligerio requires field operations that are expected to be much slower.

Notwithstanding the fact that the communication complexity in Ligerio is asymptotically better than in ours, there is a range of circuit sizes for which the concrete communication of our protocol is best. This is demonstrated in Figure 3, where we plot the communication complexity of ZKB++, Ligerio⁵, and our protocol (for two different values of n) as the circuit size varies, for a fixed input length $|w| = 256$. (In all cases, parameters were set so as to achieve 128-bit quantum security.) Figure 3 shows that for very small circuits (below ≈ 300 AND gates), ZKB++ offers the lowest communication complexity; for very large circuits (above $\approx 100,000$ AND gates), Ligerio is best. But in the range of 300–100,000 AND gates, our protocol outperforms both of those. This range captures many circuits of cryptographic interest, including those used in our signature scheme and the more advanced schemes introduced in Section 4.

3.2 A Post-Quantum Signature Scheme

As in the work of Chase et al. [CDG⁺17], we can use our NIZKPoK to obtain a signature scheme whose security is based on symmetric-key primitives only. Specifically, the private key is a uniform value $k \in \{0, 1\}^\kappa$; the public key is $y := \text{PRF}_k(0^\kappa)$, where PRF is a block cipher; and a signature is an NIZKPoK of k , where the message being signed is included as input to the hash function G used to compute the challenge. To achieve 128-bit quantum security, we use a block cipher with 256-bit key and block length.

We have implemented the resulting signature scheme, instantiating PRF with the LowMC block cipher [ARS⁺15], which was designed to be efficient for MPC applications. Our implementation was designed to enable an “apples-to-apples” comparison to the signature scheme of Chase et al. [CDG⁺17] (which is the most relevant prior work) at a 128-bit quantum security level. In particular, we use the same hash function, pseudorandom generator, and instantiation of PRF as they do.

In Table 3 we display the performance of our scheme and compare it to other post-quantum signatures based on symmetric-key primitives. Performance numbers for the ZKB++ signature scheme were obtained from the most recent technical report [PPRR17] that is slightly better than what is reported in [CDG⁺17].⁶ For our scheme, we present two sets of parameters with $n = 16$ and $n = 64$. We find that both versions of our scheme achieve better signature length compared to the scheme of Chase et al., with an improvement of 2.6–3 \times . The running time of our scheme is comparable to their scheme when $n = 16$. When $n = 64$, our running time is noticeably worse, but nevertheless still quite practical. Interestingly, the signature length of our scheme is competitive with hash-based signature schemes (not constructed using NIZKPoKs); in particular, when $n = 64$ our signatures are 6% shorter than those of SPHINCS-256 [BHH⁺15].

⁵Numbers for Ligerio are conservative estimates, since their paper only considers 40-bit security in an interactive setting.

⁶We take the numbers for the version of their scheme using the same LowMC circuit we use.

4 Constructing Ring and Group Signatures

To highlight the flexibility and power of our NIZKPoK, we also use it to build efficient ring/group signatures based on symmetric-key primitives alone.⁷ To the best of our knowledge, the resulting schemes are the most efficient constructions of these primitives that can claim to offer post-quantum security.

4.1 Ring Signatures

Ring signatures were introduced by Rivest, Shamir, and Tauman [RST01]. Roughly, they allow a user to adaptively define a *ring* R of users (of which it is a member), and then generate a signature that can be verified as having been generated by some user in that ring (without revealing precisely which one). Importantly, this can be done without the involvement of the other users in the ring.

Formally, a ring signature scheme consists of three probabilistic polynomial-time algorithms ($\text{Gen}, \text{Sign}, \text{Vrfy}$) defined as follows:

- The key-generation algorithm Gen takes as input the security parameter κ and generates public key pk and associated private key sk .
- The signing algorithm Sign takes as input a set (or “ring”) $R = \{\text{pk}_i\}$ of distinct public keys, a secret key sk corresponding to one of the public keys in R , and a message M . It outputs a signature σ .
- The verification algorithm Vrfy takes as input a ring R of distinct public keys, a message M , and a signature σ . It outputs a single bit indicating acceptance or rejection.

Correctness requires that for any collection of keys $\{(\text{pk}_i, \text{sk}_i)\}_{i \in [\ell]}$ output by Gen , any message M , and any $j \in [\ell]$, we have

$$\text{Vrfy}(\{\text{pk}_i\}, M, \text{Sign}(\{\text{pk}_i\}, \text{sk}_j, M)) = 1.$$

Ring signatures have two security requirements: unforgeability and anonymity. We present formal definitions below, following Bender et al. [BKM09].

Unforgeability. Intuitively, unforgeability means that an adversary not in R should not be able to generate a valid signature σ on a message M relative to a ring R unless some honest user in R had previously signed M (relative to the same ring) [BKM09, Definition 8].

Definition 4.1. *Ring signature scheme* $(\text{Gen}, \text{Sign}, \text{Vrfy})$ is unforgeable if, for any PPT adversary \mathcal{A} and any polynomial ℓ , the probability that \mathcal{A} succeeds in the following experiment is negligible in κ :

1. Keys $\{(\text{pk}_i, \text{sk}_i)\}_{i=1}^{\ell}$ are generated by $\text{Gen}(1^\kappa)$. The public keys $S \stackrel{\text{def}}{=} \{\text{pk}_i\}_{i=1}^{\ell}$ are given to \mathcal{A} .
2. \mathcal{A} may query an oracle $\text{Sign}'(\cdot, \cdot, \cdot)$, where $\text{Sign}'(R, i, M)$ (with $\text{pk}_i \in R$) outputs $\text{Sign}(R, \text{sk}_i, M)$. (Note that we do not require $R \subseteq S$.)
3. \mathcal{A} may also query a corruption oracle Corrupt that on input i returns sk_i . If \mathcal{A} queries $\text{Corrupt}(i)$ then we say that pk_i is corrupted. We let C be the set of corrupted public keys at the end of the experiment.

⁷Note that group signatures with full anonymity imply public-key encryption [AW04], and thus the group signatures we construct will meet a slightly weaker definition.

4. \mathcal{A} outputs M^*, R^*, σ^* . It succeeds if (1) $\text{Vrfy}(R^*, M^*, \sigma^*) = 1$; (2) $R^* \subseteq S \setminus C$; and (3) \mathcal{A} never queried $\text{Sign}'(R^*, \star, M^*)$.

Anonymity. Intuitively, anonymity ensures that a valid signature with respect to a ring R does not reveal which secret key (corresponding to some public key in R) was used to generate the signature. Our definition is based on [BKM09, Definition 5], slightly simplified for our context.

Definition 4.2. *Ring signature scheme* $(\text{Gen}, \text{Sign}, \text{Vrfy})$ is anonymous if, for any PPT adversary \mathcal{A} and any polynomial ℓ , the probability that \mathcal{A} succeeds in the following experiment is at most $1/2 + \text{negl}(\kappa)$.

1. Keys $\{(\text{pk}_i, \text{sk}_i)\}_{i=1}^{\ell}$ are generated by $\text{Gen}(1^\kappa)$ and all keys (both public and private) are given to \mathcal{A} .
2. \mathcal{A} outputs a message M , a ring R , and $i_0, i_1 \in [\ell]$. A uniform $b \in \{0, 1\}$ is chosen, and \mathcal{A} is given $\text{Sign}(R', \text{sk}_{i_b}, M)$, where $R' = R \cup \{\text{pk}_{i_0}, \text{pk}_{i_1}\}$.
3. \mathcal{A} outputs a bit b' , and succeeds if $b' = b$.

Constructing a ring signature scheme. The main idea is to generate a ring signature with respect to a ring R by generating an NIZKPoK of the secret key corresponding to one of the public keys in R . The main question is how to do so while reducing the size of the proof as much as possible.

A naive solution to building a ring signature scheme, based on the signature scheme from the previous section, is as follows. As before, key generation chooses a uniform key $k \in \{0, 1\}^\kappa$; the corresponding public key is $y := \text{PRF}_k(0^\kappa)$. Given a ring $R = \{y_i\}_{i=1}^{\ell}$ of public keys, let C be the circuit that takes as input a key k and outputs 1 iff $\text{PRF}_k(0^\kappa) = y_i$ for some i . (Alternately, it can take as input k, i and output 1 iff $\text{PRF}_k(0^\kappa) = y_i$. This does not materially affect the parameters.) A signature will then be an NIZKPoK of an input k_i such that $C(k_i) = 1$. The size of C (and hence the size of a signature) grows linearly in $\ell = |R|$.

Efficiency can be improved using a Merkle tree. Given a ring R as before, we define a (public) mechanism for computing a Merkle root h^* starting from values y_1, \dots, y_ℓ at the leaves; assume for simplicity that $\ell = 2^q$ for some integer q . Let C be a circuit that takes as input k, i , and an auxiliary value `path`; computes $y := \text{PRF}_k(0^\kappa)$; and then outputs 1 if `path` is a valid Merkle proof (with respect to root h^*) for value y at leaf i . As before, a signature is an NIZKPoK of an input for which C evaluates to 1. Importantly, the size of C (and hence the signature length) is now only logarithmic in ℓ .

A proof of security for the resulting ring signature scheme is immediate. Anonymity is ensured by the zero-knowledge property of the NIZKPoK. Given an adversary who forges a signature for some ring $R^* \subset R \setminus C$ (and assuming the hash function used for the Merkle tree is collision-resistant), we can use the knowledge extractor of the NIZKPoK to extract some (i, k) such that $\text{PRF}_{k'}(0^\kappa) = y_i$ for some party i who is not corrupted by the adversary. That is, the adversary is able to find a private key for some uncorrupted party, which happens with negligible probability. As in the previous section, since the scheme relies on symmetric-key primitives alone it can be said to offer post-quantum security. We refer to Section 4.3 for a performance evaluation of the scheme.

Ring/group size:	2^7		2^{10}		2^{13}	
	$ \sigma $	t	$ \sigma $	t	$ \sigma $	t
Derler et al. [DRS17]	982 KB	—	1.35 MB	—	1.72 MB	—
Here	285 KB	2.0 s	388 KB	2.8 s	492 KB	3.6 s
Boneh et al. [BEF18]	1.37 MB	—	1.85 MB	—	—	—
Here	315 KB	2.3 s	418 KB	3.0 s	532 KB	3.8 s

Table 4: Performance of our ring/group signature schemes, and comparison to prior work. $|\sigma|$ is the signature size, and t is the signing/verification time. The top two schemes are ring signatures; the bottom two are group signatures.

4.2 Group Signatures

A group signature scheme [Cv91] also allows users in a group to sign messages without revealing their identity (beyond the fact that they belong to the group). The main differences between group signatures and ring signatures are that (1) in the context of group signatures, a *group manager* determines membership in the group; and (2) the group manager has the ability to violate anonymity by *tracing* the user who generated a given signature.

We focus on the case of static groups for simplicity. Formally, a group signature scheme consists of four probabilistic polynomial-time algorithms (Gen , Sign , Vrfy , Open) with the following functionality:

- The key-generation algorithm Gen takes as input the security parameter κ and the number of group members ℓ , and generates a master public key gpk , master private key gmsk , and signing keys $\text{gsk}_1, \dots, \text{gsk}_\ell$.
- The signing algorithm Sign takes as input a signing key gsk_i and a message M , and outputs a signature σ .
- The verification algorithm Vrfy takes as input the master public key gpk , a message M , and a signature σ . It outputs a bit indicating acceptance or rejection.
- The tracing algorithm Open takes as input the master private key gmsk , a message M , and a signature σ . It outputs an index $i \in [\ell]$ or a failure symbol \perp .

Correctness requires that for any $\text{gpk}, \text{gmsk}, \text{gsk}_1, \dots, \text{gsk}_\ell$ output by Gen , any message M , and any $i \in [\ell]$, we have

$$\text{Vrfy}(\text{gpk}, M, \text{Sign}(\text{gsk}_i, M)) = 1$$

and

$$\text{Open}(\text{gmsk}, M, \text{Sign}(\text{gsk}_i, M)) = i.$$

Bellare et al. [BMW03] define two security requirements for group signatures: traceability and anonymity. We follow their definitions here with some differences described below.

Traceability. The traceability property requires that an attacker who colludes with some set C of group members cannot output a valid signature that fails to be traced by the group manager to some member of C .

Definition 4.3. *Group signature scheme* $(\text{Gen}, \text{Sign}, \text{Vrfy}, \text{Open})$ is traceable if, for any PPT adversary \mathcal{A} and any polynomial ℓ , the probability that \mathcal{A} succeeds in the following experiment is negligible in κ :

1. Keys $\text{gpk}, \text{gmsk}, \text{gsk}_1, \dots, \text{gsk}_\ell$ are output by $\text{Gen}(1^\kappa, 1^\ell)$, and gpk, gmsk are given to \mathcal{A} .
2. \mathcal{A} may query an oracle $\text{Sign}'(\cdot, \cdot)$, where $\text{Sign}'(M, i)$ outputs $\text{Sign}(M, \text{gsk}_i)$.
3. \mathcal{A} may also query a corruption oracle **Corrupt** that on input i returns gsk_i ; in this case we say that user i is corrupted. We let C denote the set of corrupted users at the end of the experiment.
4. \mathcal{A} outputs M, σ , and succeeds if: (1) $\text{Vrfy}(\text{gpk}, M, \sigma) = 1$; (2) $\text{Open}(\text{gmsk}, M, \sigma) \notin C$; and (3) the adversary never queried $\text{Sign}'(M, \star)$.

As shown by Bellare et al. [BMW03], traceability implies unforgeability (simply consider $C = \emptyset$ in the above definition). One could also consider a weaker definition of traceability (that still implies unforgeability) in which gmsk is not given to the adversary. Our scheme can be made more efficient if that weaker definition suffices.

Anonymity. As in the case of ring signatures, anonymity implies that a valid signature does not reveal the user who generated that signature (except to the group manager who holds gmsk); this should hold even if the adversary is given access to a tracing oracle to which it can submit any other signature. The notion of full anonymity considered by Bellare et al. [BMW03] requires signatures produced by a user i to remain anonymous even if the secret key gsk_i of that user is known to the adversary. A weaker notion [CG05] requires anonymity only for uncorrupted users. Previous work [AW04] shows that fully anonymous group signatures imply public-key encryption. Since our goal is to construct a scheme based on symmetric-key primitives alone, we consider the weaker notion of anonymity here.

Definition 4.4. *Group signature scheme* $(\text{Gen}, \text{Sign}, \text{Vrfy}, \text{Open})$ satisfies weak anonymity if, for any PPT adversary \mathcal{A} and any polynomial ℓ , the probability that \mathcal{A} succeeds in the following experiment is at most $1/2 + \text{negl}(\kappa)$:

1. Keys $\text{gpk}, \text{gmsk}, \text{gsk}_1, \dots, \text{gsk}_\ell$ are output by $\text{Gen}(1^\kappa, 1^\ell)$, and gpk is given to \mathcal{A} .
2. \mathcal{A} is given access to the following oracles:
 - A signing oracle that on input M, i outputs $\text{Sign}(M, \text{gsk}_i)$.
 - A corruption oracle that on input i outputs gsk_i . We let C denote the set of queries made by \mathcal{A} to this oracle at the end of the experiment.
 - A tracing oracle that outputs $\text{Open}(\text{gmsk}, M, \sigma)$.
3. At some point, \mathcal{A} outputs a message M^* and $i_0, i_1 \in [\ell]$. Then a uniform $b \in \{0, 1\}$ is chosen, $\sigma^* \leftarrow \text{Sign}(M^*, \text{gsk}_{i_b})$ is computed, and \mathcal{A} is given σ^* .
4. \mathcal{A} may continue to query all the oracles above, except that it may not query its tracing oracle on M^*, σ^* .
5. \mathcal{A} outputs a bit b' , and succeeds if $i_0, i_1 \notin C$ and $b' = b$.

Constructing a group signature scheme. Analogous to our construction of a ring signature scheme, the basic idea here is for a user to generate a signature by generating an NIZKPoK of a secret key corresponding to the group public key. Here, however, we must also ensure traceability, which adds additional complications.

Assume for simplicity that $\ell = 2^q$ for some integer q . In our scheme, the private key gsk_i of a user now includes two uniform and independent keys k_i^0, k_i^1 ; we set $y_i^b := \text{PRF}_{k_i^b}(0^\kappa)$. The master private key gmsk consists of the $\{k_i^0\}$ but, importantly, does not include the $\{k_i^1\}$. A Merkle root h^* is then computed for the Merkle tree with values $(y_1^0, y_1^1), \dots, (y_\ell^0, y_\ell^1)$ at the ℓ leaves of the tree. The master public key is h^* , and the private key gsk_i of a user includes the Merkle proof (with respect to h^*) for (y_i^0, y_i^1) .

Let $C_{x,y}$ be a circuit that has values x, y hardcoded; takes as input k^0, k^1, i , and an auxiliary value path ; computes $y^0 := \text{PRF}_{k^0}(0^\kappa)$ and $y^1 := \text{PRF}_{k^1}(0^\kappa)$; and then outputs 1 iff $y = \text{PRF}_{k^0}(x)$ and path is a valid Merkle proof (with respect to h^*) for the value (y^0, y^1) at leaf i . To sign a message M , a user i holding k_i^0, k_i^1 computes $y = \text{PRF}_{k_i^0}(H(M))$ and then generates an NIZKPoK of an input for which the circuit $C_{H(M),y}$ outputs 1; the signature includes both y and the resulting proof. Tracing is done by finding the key k_i^0 for which $\text{PRF}_{k_i^0}(H(M)) = y$. We refer to Section 4.3 for a performance evaluation of the scheme.

Traceability and weak anonymity of the above construction follows using standard techniques; we omit the details.

4.3 Implementation and Performance

We have implemented both the ring and group signature schemes described in the previous sections, using our NIZKPoK from Section 3.1 with $n = 64$ so as to optimize the signature length. As in the previous section, we target 128-bit quantum security. Both of our schemes involving constructing a circuit that can verify a Merkle path, which requires several hash-function computations. To minimize the circuit size, we derived a fixed-length hash function mapping 512-bit inputs to 256-bit outputs by applying the Davies-Meyer construction to the LowMC cipher with 256-bit key/block size. We summarize the performance of our schemes in Table 4.

The signature length in our ring signature scheme asymptotically matches what is achieved by the state-of-the-art, lattice-based, ring signature scheme by Libert et al. [LLNW16], but the concrete efficiency of our scheme (both in terms of communication and computation) appears to be significantly better. (Libert et al. do not provide an implementation of their scheme.) As compared to their scheme, ours has the advantage of not having to rely on trusted public parameters.

Recent work by Ling et al. [LNWX18] gives a group signature scheme based on lattices for which the signature size is independent of the number of users in the group. However, they do not investigate the concrete efficiency of their scheme, and it appears that our scheme will out-perform theirs for practical settings of the parameters.

Subsequent to our work (cf. footnote 2), two other groups [DRS17, BEF18] have proposed constructions of ring and group signature schemes using ideas similar to ours, but based on ZKB++ as the underlying NIZKPoK. As shown in Table 4, by using our NIZKPoK in place of ZKB++ we are able to generate signatures roughly 3.5–4.4× shorter than in their work.

Acknowledgments

This material is based on work supported by NSF awards #1111599 and #1563722. Portions of this work were also supported by DARPA and SPAWAR under contract N66001-15-C-4065. Second and third authors were supported by the Office of Naval Research (ONR) contract number N00014-14-C0113. Work partially done while second and third authors were at Bell Labs. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes not withstanding any copyright notation thereon. The views, opinions, and/or findings expressed are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

References

- [AABN02] Michel Abdalla, Jee Hea An, Mihir Bellare, and Chanathip Namprempre. From Identification to Signatures via the Fiat-Shamir Transform: Minimizing Assumptions for Security and Forward-Security. In Lars R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 418–433, Amsterdam, The Netherlands, April 28 – May 2, 2002. Springer, Heidelberg, Germany.
- [AHIV17] Scott Ames, Carmit Hazay, Yuval Ishai, and Muthuramakrishnan Venkitasubramanian. Liger: Lightweight Sublinear Arguments Without a Trusted Setup. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 17*, pages 2087–2104, Dallas, TX, USA, October 31 – November 2, 2017. ACM Press.
- [ARS⁺15] Martin R. Albrecht, Christian Rechberger, Thomas Schneider, Tyge Tiessen, and Michael Zohner. Ciphers for MPC and FHE. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 430–454, Sofia, Bulgaria, April 26–30, 2015. Springer, Heidelberg, Germany.
- [AW04] Michel Abdalla and Bogdan Warinschi. On the Minimal Assumptions of Group Signature Schemes. In Javier López, Sihan Qing, and Eiji Okamoto, editors, *ICICS 04*, volume 3269 of *LNCS*, pages 1–13, Malaga, Spain, October 27–29, 2004. Springer, Heidelberg, Germany.
- [BCG⁺13] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, Eran Tromer, and Madars Virza. SNARKs for C: Verifying Program Executions Succinctly and in Zero Knowledge. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 90–108, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Heidelberg, Germany.
- [BDSMP91] Manuel Blum, Alfredo De Santis, Silvio Micali, and Giuseppe Persiano. Noninteractive zero-knowledge. *SIAM Journal on Computing*, 20(6):1084–1118, 1991.
- [BEF18] Dan Boneh, Saba Eskandarian, and Ben Fisch. Post-Quantum EPID Group Signatures from Symmetric Primitives. Cryptology ePrint Archive, Report 2018/261, 2018. <https://eprint.iacr.org/2018/261>.

- [BFM88] Manuel Blum, Paul Feldman, and Silvio Micali. Non-Interactive Zero-Knowledge and Its Applications (Extended Abstract). In *20th ACM STOC*, pages 103–112, Chicago, IL, USA, May 2–4, 1988. ACM Press.
- [BG93] Mihir Bellare and Oded Goldreich. On Defining Proofs of Knowledge. In Ernest F. Brickell, editor, *CRYPTO'92*, volume 740 of *LNCS*, pages 390–420, Santa Barbara, CA, USA, August 16–20, 1993. Springer, Heidelberg, Germany.
- [BHH⁺15] Daniel J. Bernstein, Daira Hopwood, Andreas Hülsing, Tanja Lange, Ruben Niederhagen, Louiza Papachristodoulou, Michael Schneider, Peter Schwabe, and Zooko Wilcox-O'Hearn. SPHINCS: Practical Stateless Hash-Based Signatures. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 368–397, Sofia, Bulgaria, April 26–30, 2015. Springer, Heidelberg, Germany.
- [BISW17] Dan Boneh, Yuval Ishai, Amit Sahai, and David J. Wu. Lattice-Based SNARGs and Their Application to More Efficient Obfuscation. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part III*, volume 10212 of *LNCS*, pages 247–277, Paris, France, April 30 – May 4, 2017. Springer, Heidelberg, Germany.
- [BKM09] Adam Bender, Jonathan Katz, and Ruggero Morselli. Ring Signatures: Stronger Definitions, and Constructions without Random Oracles. *Journal of Cryptology*, 22(1):114–138, January 2009.
- [BMW03] Mihir Bellare, Daniele Micciancio, and Bogdan Warinschi. Foundations of Group Signatures: Formal Definitions, Simplified Requirements, and a Construction Based on General Assumptions. In Eli Biham, editor, *EUROCRYPT 2003*, volume 2656 of *LNCS*, pages 614–629, Warsaw, Poland, May 4–8, 2003. Springer, Heidelberg, Germany.
- [CDG⁺17] Melissa Chase, David Derler, Steven Goldfeder, Claudio Orlandi, Sebastian Ramacher, Christian Rechberger, Daniel Slamanig, and Greg Zaverucha. Post-Quantum Zero-Knowledge and Signatures from Symmetric-Key Primitives. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 17*, pages 1825–1842, Dallas, TX, USA, October 31 – November 2, 2017. ACM Press.
- [CFH⁺15] Craig Costello, Cédric Fournet, Jon Howell, Markulf Kohlweiss, Benjamin Kreuter, Michael Naehrig, Bryan Parno, and Samee Zahur. Geppetto: Versatile Verifiable Computation. In *2015 IEEE Symposium on Security and Privacy*, pages 253–270, San Jose, CA, USA, May 17–21, 2015. IEEE Computer Society Press.
- [CG05] Jan Camenisch and Jens Groth. Group Signatures: Better Efficiency and New Theoretical Aspects. In Carlo Blundo and Stelvio Cimato, editors, *SCN 04*, volume 3352 of *LNCS*, pages 120–133, Amalfi, Italy, September 8–10, 2005. Springer, Heidelberg, Germany.
- [CHR⁺18] Ming-Shing Chen, Andreas Hülsing, Joost Rijneveld, Simona Samardjiska, and Peter Schwabe. SOFIA: MQ-Based Signatures in the QROM. In Michel Abdalla and Ricardo Dahab, editors, *PKC 2018, Part II*, volume 10770 of *LNCS*, pages 3–33, Rio de Janeiro, Brazil, March 25–29, 2018. Springer, Heidelberg, Germany.

- [Cv91] David Chaum and Eugène van Heyst. Group Signatures. In Donald W. Davies, editor, *EUROCRYPT'91*, volume 547 of *LNCS*, pages 257–265, Brighton, UK, April 8–11, 1991. Springer, Heidelberg, Germany.
- [DP92] Alfredo De Santis and Giuseppe Persiano. Zero-Knowledge Proofs of Knowledge Without Interaction (Extended Abstract). In *33rd FOCS*, pages 427–436, Pittsburgh, Pennsylvania, October 24–27, 1992. IEEE Computer Society Press.
- [DRS17] David Derler, Sebastian Ramacher, and Daniel Slamanig. Post-Quantum Zero-Knowledge Proofs for Accumulators with Applications to Ring Signatures from Symmetric-Key Primitives. Cryptology ePrint Archive, Report 2017/1154, 2017.
- [FLS90] Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple Non-Interactive Zero Knowledge Proofs Based on a Single Random String (Extended Abstract). In *31st FOCS*, pages 308–317, St. Louis, Missouri, October 22–24, 1990. IEEE Computer Society Press.
- [FNO15] Tore Kasper Frederiksen, Jesper Buus Nielsen, and Claudio Orlandi. Privacy-Free Garbled Circuits with Applications to Efficient Zero-Knowledge. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 191–219, Sofia, Bulgaria, April 26–30, 2015. Springer, Heidelberg, Germany.
- [FS87] Amos Fiat and Adi Shamir. How to Prove Yourself: Practical Solutions to Identification and Signature Problems. In Andrew M. Odlyzko, editor, *CRYPTO'86*, volume 263 of *LNCS*, pages 186–194, Santa Barbara, CA, USA, August 1987. Springer, Heidelberg, Germany.
- [GGPR13] Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic Span Programs and Succinct NIZKs without PCPs. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 626–645, Athens, Greece, May 26–30, 2013. Springer, Heidelberg, Germany.
- [GMO16] Irene Giacomelli, Jesper Madsen, and Claudio Orlandi. ZKBoo: Faster Zero-Knowledge for Boolean Circuits. In *USENIX Security*, 2016.
- [GMR85] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The Knowledge Complexity of Interactive Proof-Systems (Extended Abstract). In *17th ACM STOC*, pages 291–304, Providence, RI, USA, May 6–8, 1985. ACM Press.
- [GMW91] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs That Yield Nothing But Their Validity Or All Languages in NP Have Zero-Knowledge Proof Systems. *Journal of the ACM*, 38(3):691–729, 1991.
- [Gro16] Jens Groth. On the Size of Pairing-Based Non-interactive Arguments. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 305–326, Vienna, Austria, May 8–12, 2016. Springer, Heidelberg, Germany.
- [GW11] Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In Lance Fortnow and Salil P. Vadhan, editors, *43rd ACM STOC*, pages 99–108, San Jose, CA, USA, June 6–8, 2011. ACM Press.

- [IKOS07] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge from secure multiparty computation. In David S. Johnson and Uriel Feige, editors, *39th ACM STOC*, pages 21–30, San Diego, CA, USA, June 11–13, 2007. ACM Press.
- [JKO13] Marek Jawurek, Florian Kerschbaum, and Claudio Orlandi. Zero-knowledge using garbled circuits: how to prove non-algebraic statements efficiently. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 13*, pages 955–966, Berlin, Germany, November 4–8, 2013. ACM Press.
- [Kil92] Joe Kilian. A Note on Efficient Zero-Knowledge Proofs and Arguments (Extended Abstract). In *24th ACM STOC*, pages 723–732, Victoria, British Columbia, Canada, May 4–6, 1992. ACM Press.
- [KOS16] Marcel Keller, Emmanuela Orsini, and Peter Scholl. MASCOT: Faster Malicious Arithmetic Secure Computation with Oblivious Transfer. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 16*, pages 830–842, Vienna, Austria, October 24–28, 2016. ACM Press.
- [LLNW16] Benoît Libert, San Ling, Khoa Nguyen, and Huaxiong Wang. Zero-Knowledge Arguments for Lattice-Based Accumulators: Logarithmic-Size Ring Signatures and Group Signatures Without Trapdoors. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 1–31, Vienna, Austria, May 8–12, 2016. Springer, Heidelberg, Germany.
- [LNWX18] San Ling, Khoa Nguyen, Huaxiong Wang, and Yanhong Xu. Constant-Size Group Signatures from Lattices. In Michel Abdalla and Ricardo Dahab, editors, *PKC 2018, Part II*, volume 10770 of *LNCS*, pages 58–88, Rio de Janeiro, Brazil, March 25–29, 2018. Springer, Heidelberg, Germany.
- [Mic94] Silvio Micali. CS Proofs (Extended Abstracts). In *35th FOCS*, pages 436–453, Santa Fe, New Mexico, November 20–22, 1994. IEEE Computer Society Press.
- [NNL01] Dalit Naor, Moni Naor, and Jeffery Lotspiech. Revocation and Tracing Schemes for Stateless Receivers. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 41–62, Santa Barbara, CA, USA, August 19–23, 2001. Springer, Heidelberg, Germany.
- [PHGR13] Bryan Parno, Jon Howell, Craig Gentry, and Mariana Raykova. Pinocchio: Nearly Practical Verifiable Computation. In *2013 IEEE Symposium on Security and Privacy*, pages 238–252, Berkeley, CA, USA, May 19–22, 2013. IEEE Computer Society Press.
- [PPRR17] L Perrin, Angela Promitzer, Sebastian Ramacher, and Christian Rechberger. Improvements to the Linear Layer of LowMC: A Faster Picnic. Cryptology ePrint Archive, Report 2017/1148, 2017. <https://eprint.iacr.org/2017/1148>.
- [PS00] David Pointcheval and Jacques Stern. Security Arguments for Digital Signatures and Blind Signatures. *Journal of Cryptology*, 13(3):361–396, 2000.

- [RST01] Ronald L. Rivest, Adi Shamir, and Yael Tauman. How to Leak a Secret. In Colin Boyd, editor, *ASIACRYPT 2001*, volume 2248 of *LNCS*, pages 552–565, Gold Coast, Australia, December 9–13, 2001. Springer, Heidelberg, Germany.
- [RTZ16] Samuel Ranellucci, Alain Tapp, and Rasmus Winther Zakarias. Efficient Generic Zero-Knowledge Proofs from Commitments (Extended Abstract). In Anderson C. A. Nascimento and Paulo Barreto, editors, *ICITS 16*, volume 10015 of *LNCS*, pages 190–212, Tacoma, WA, USA, August 9–12, 2016. Springer, Heidelberg, Germany.
- [Unr15] Dominique Unruh. Non-Interactive Zero-Knowledge Proofs in the Quantum Random Oracle Model. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 755–784, Sofia, Bulgaria, April 26–30, 2015. Springer, Heidelberg, Germany.
- [Unr17] Dominique Unruh. Post-quantum Security of Fiat-Shamir. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part I*, volume 10624 of *LNCS*, pages 65–95, Hong Kong, China, December 3–7, 2017. Springer, Heidelberg, Germany.
- [WMK17] Xiao Wang, Alex J. Malozemoff, and Jonathan Katz. Faster Secure Two-Party Computation in the Single-Execution Setting. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part III*, volume 10212 of *LNCS*, pages 399–424, Paris, France, April 30 – May 4, 2017. Springer, Heidelberg, Germany.
- [WRK17] Xiao Wang, Samuel Ranellucci, and Jonathan Katz. Authenticated Garbling and Efficient Maliciously Secure Two-Party Computation. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 17*, pages 21–37, Dallas, TX, USA, October 31 – November 2, 2017. ACM Press.