



Public-Coin Statistical Zero-Knowledge Batch Verification Against Malicious Verifiers

Inbar Kaslasi^{1(✉)}, Ron D. Rothblum¹, and Prashant Nalini Vasudevan²

¹ Technion, Haifa, Israel

{inbark,rothblum}@cs.technion.ac.il

² UC Berkeley, Berkeley, USA

prashvas@berkeley.edu

Abstract. Suppose that a problem Π has a statistical zero-knowledge (SZK) proof with communication complexity m . The question of batch verification for SZK asks whether one can prove that k instances x_1, \dots, x_k all belong to Π with a statistical zero-knowledge proof whose communication complexity is better than $k \cdot m$ (which is the complexity of the trivial solution of executing the original protocol independently on each input).

In a recent work, Kaslasi *et al.* (TCC, 2020) constructed such a batch verification protocol for any problem having a *non-interactive* SZK (NISZK) proof-system. Two drawbacks of their result are that their protocol is private-coin and is only zero-knowledge with respect to the honest verifier.

In this work, we eliminate these two drawbacks by constructing a public-coin malicious-verifier SZK protocol for batch verification of NISZK. Similarly to the aforementioned prior work, the communication complexity of our protocol is $(k + \text{poly}(m)) \cdot \text{polylog}(k, m)$.

Keywords: Statistical zero-knowledge · Batch verification

1 Introduction

The concept of zero knowledge proofs, introduced by Goldwasser, Micali and Rackoff [GMR89], is an incredibly deep and fascinating notion that has proven to be a fundamental component in the construction and design of cryptographic protocols (see, e.g., [GMW87]). A zero-knowledge proof allows a prover to convince an efficient verifier that a given statement is true without revealing anything else to the verifier. This is formalized by requiring that for any (possibly malicious) verifier that participates in such a proof, there is an efficient simulation algorithm that simulates its interaction with the prover.

In this work we focus on *statistical zero-knowledge proofs*. In this variant, both the verifier and the prover are guaranteed information-theoretic (rather

The full version is available on the Cryptology ePrint Archive [KRV21].

© International Association for Cryptologic Research 2021

A. Canteaut and F.-X. Standaert (Eds.): EUROCRYPT 2021, LNCS 12698, pp. 219–246, 2021.

https://doi.org/10.1007/978-3-030-77883-5_8

than computational) security. On the one hand, the verifier knows that even an all-powerful prover could not convince it to accept a false statement (other than with negligible probability). On the other hand, the prover knows that any polynomial-time cheating strategy of the verifier can only reveal a negligible amount of information beyond the validity of the statement being proven.

The class of languages having a statistical zero-knowledge protocol is denoted by SZK. This class contains several natural problems like Graph Nonisomorphism, and many of the problems that are central to cryptography such as quadratic residuosity [GMR89], discrete logarithm [GK93, CP92], and various lattice problems [GG00, MV03, PV08, APS18]. It has been found to possess extremely rich structure [For89, AH91, Oka00, SV03, GSV98, GV99, NV06, OV08] and to have fundamental connections to different aspects of cryptography [BL13, KMN+14, LV16, Ost91, OW93, BDRV18, KY18, BBD+20] and complexity theory [For87, AH91, Aar12, GR14, Dru15, AV19, BCH+20].

In a recent work, Kaslasi *et al.* [KRR+20] raised the question of *batch verification for statistical zero-knowledge proofs*: Suppose Π has a statistical zero-knowledge proof (SZK). Can we prove that $x_1, \dots, x_k \in \Pi$ with communication complexity that beats the naive approach of separately proving that each $x_i \in \Pi$, while still maintaining zero-knowledge? Beyond being of intrinsic interest and teaching us about the structure of SZK, such protocols have potential cryptographic applications such as the batch verification of cryptographic signatures [NMVR94, BGR98, CHP12] or well-formedness of public-keys [GMR98].

The main result of [KRR+20] was such a generic batch verification result for a subclass of languages in SZK – specifically for problems having a *non-interactive* statistical zero-knowledge proof system (NISZK). Kaslasi *et al.* construct an (interactive) SZK protocol for batching k instances of $\Pi \in \text{NISZK}$, with communication complexity $(k + \text{poly}(n)) \cdot \text{polylog}(k, n)$, where n is the length of each of the k inputs, and poly refers to a fixed polynomial that depends only on the specific problem (and not on k). Their result should be contrasted with the naive approach of simply executing the NISZK protocol separately on each input (which has communication complexity $k \cdot \text{poly}(n)$).

Two major drawbacks of the protocol of [KRR+20] are the fact that it is *private-coin* and only *honest-verifier* statistical zero-knowledge (HVSZK). These drawbacks are significant. Recall that private-coin protocols can only be verified by a designated verifier, in contrast to *public-coin* protocols that can be verified by anyone (as long as they can ensure that the coins were truly unpredictable to the prover, e.g., they were generated by some physical phenomenon or a public randomness beacon). Further, public-coin protocols have the added benefit that they can be transformed into *non-interactive* arguments via the Fiat-Shamir transform (either heuristically [FS86], in the random-oracle model [PS96], or, more recently, under concrete cryptographic assumptions (see, e.g., [CCH+19])).

The second drawback is arguably even more significant. Recall that honest-verifier zero-knowledge is a relatively weak privacy guarantee which, in a nutshell, only guarantees that verifiers that follow the protocol to the letter learn nothing in the interaction. Usually this weak privacy guarantee is only used as a stepping

stone towards getting full-fledged zero-knowledge (i.e., zero-knowledge that holds even against arbitrary polynomial-time cheating verifiers).

At first glance it may seem straightforward to overcome both drawbacks of the protocol of [KRR+20] by employing the known generic transformations from *private-coin honest-verifier statistical zero-knowledge* to *public-coin malicious-verifier statistical zero-knowledge* [Oka00, GSV98, HRV18]. Unfortunately, these transformations incur a large polynomial overhead in communication that we cannot afford in our context (see also Remark 1 below).

1.1 Our Results

In this paper we eliminate the two major drawbacks mentioned above by constructing a *public-coin malicious-verifier SZK* batch verification protocol for every problem in NISZK. The communication complexity in our protocol is similar to that of [KRR+20].

Theorem 1 (Informally Stated, see Theorem 7). *Let $\Pi \in \text{NISZK}$. There exists a public-coin SZK protocol for verifying that $x_1, \dots, x_k \in \Pi$, with communication complexity $(k + \text{poly}(n)) \cdot \text{polylog}(n, k)$. The verifier's running time is $k \cdot \text{poly}(n, \log k)$, and the number of rounds is $k \cdot \text{polylog}(n, k)$.*

Our high-level approach for proving Theorem 1 follows a classical approach for constructing malicious-verifier zero-knowledge proofs: first construct a public-coin honest-verifier zero-knowledge batching protocol, and then show how to transform it to be *malicious-verifier* zero-knowledge. The main challenge that we must overcome is in actually implementing these two steps without incurring the exorbitant price of the generic transformations for SZK [Oka00, SV03, GV99, GSV98, HRV18]. Thus, our two main steps are:

1. Construct an efficient *public-coin HVSZK* batch verification protocol.
2. Transform it to be zero-knowledge against malicious verifiers, while preserving its efficiency.

Our first main technical contribution is in implementing Step 1.

Theorem 2 (Informally Stated, see Theorem 6). *Let $\Pi \in \text{NISZK}$. There exists a public-coin HVSZK protocol for verifying that $x_1, \dots, x_k \in \Pi$, with communication complexity $(k + \text{poly}(n)) \cdot \text{polylog}(n, k)$. The verifier's running time is $k \cdot \text{poly}(n, \log k)$, and the number of rounds is $O(k)$.*

Theorem 2 already improves on the main result of [KRR+20], since it gives a *public-coin* batch verification protocol. However, we would like to go further and obtain security even against malicious verifiers. It is tempting at this point to apply the generic transformations of [GSV98, HRV18] from public-coin *honest-verifier* zero-knowledge, to *malicious-verifier*. Unfortunately, the overhead introduced in these transformations is too large and applying them to the protocol of Theorem 1 would yield a trivial result (see Remark 1 for details).

Rather, as our second technical contribution (which may be of independent interest), we show that the communication complexity of the [GSV98] transformation can be significantly improved for protocols satisfying a strong notion of soundness. Specifically, we refer to the notion of *round-by-round soundness*, introduced in a recent work of Canetti *et al.* [CCH+19].

Theorem 3 (Informally Stated, see Theorem 5). *Any public-coin HVSZK protocol with negligible round-by-round soundness error can be efficiently transformed into a public-coin SZK protocol. In particular, a message of length ℓ in the original protocol grows to length $\text{poly}(\ell)$ in the transformed protocol.*

Note that the growth of each message in the transformation above depends only on its own length and not on n or k – this allows us to take advantage of the fact that all but one of the messages in the protocol of Theorem 2 have polylogarithmic length. We show that the protocol of Theorem 2 indeed has round-by-round soundness which, in combination with Theorem 3, yields Theorem 1.

Remark 1 (On Generic Transformations from the Literature). We discuss here a few known generic transformations for the class SZK from the literature, and why they are not applicable in our context.

Okamoto [Oka00] showed how to transform any *private-coin* HVSZK protocol into a public-coin one. Unfortunately, we cannot use his transformation to derive Theorem 1 from the private-coin batching protocol of [KRR+20], due to the overhead involved. In more detail, Okamoto’s protocol starts by taking a t -fold parallel repetition of the private-coin protocol, where $t = \ell^9 \cdot r^9$, where ℓ is the round complexity and r is the randomness complexity of the simulator. In our context $\ell = k$ and $r = \text{poly}(n)$ and so the overhead from Okamoto’s transformation would yield a trivial result (as a matter of fact, we could not even afford an overhead of $t = \ell$, which seems inherent to Okamoto’s approach).

Similarly, we cannot derive Theorem 1 from Theorem 2 by applying the generic transformation of [GSV98, HRV18] for transforming honest-verifier public-coin SZK proofs to malicious-verifier ones. In more detail, the transformation of [GSV98] starts by applying an ℓ -fold parallel repetition of the honest-verifier protocol (where again ℓ is the number of rounds). In the context of Theorem 2, $\ell = \Theta(k)$, and so, applying the [GSV98] transformation yields a protocol with communication complexity $k \cdot \text{poly}(n)$, which we cannot afford.

The more recent work of Hubáček *et al.* [HRV18] gives an efficiency-preserving transformation from honest-verifier to malicious-verifier. This transformation also incurs a polynomial overhead in the communication complexity. In particular, [HRV18] rely on the instance-dependent commitments of Ong and Vadhan [OV08], which in turn use the SZK completeness of the Entropy Difference (or ED) problem [GV99]. The known reduction from SZK to ED (see [Vad99, Theorem 3.3.13]) generates circuits whose input size is roughly $\ell \cdot r$. This size would correspond to the size of the decommitments in the [HRV18] protocol and again would lead to an overhead that we cannot incur.

Indeed, our work motivates the study of *communication-preserving* transformations for SZK protocols. In particular, obtaining a *generic* communication-preserving transformation from honest-verifier to malicious-verifier SZK is an interesting open question.

1.2 Technical Overview

First, in Sect. 1.2, we describe the public-coin honest-verifier statistical zero-knowledge (HVSZK) batching protocol. Then, in Sect. 1.2, we show how to compile honest-verifier protocols *efficiently* to be secure against malicious verifiers.

Public-Coin HVSZK Batching. Our starting point is the aforementioned recent work of Kaslasi *et al.* [KRR+20], which gave a *private-coin* HVSZK batching protocol. As their first step, they introduced a new (promise) problem called *approximate injectivity* (AI) and showed that it is NISZK-complete. They then designed a private-coin HVSZK batch verification protocol for AI. We follow a similar route, except that we construct a *public-coin* HVSZK batch verification protocol for AI.

The inputs to the AI problem, which is parameterized by a real number $\delta \in [0, 1]$, are circuits $C : \{0, 1\}^n \rightarrow \{0, 1\}^m$. YES instances are circuits C for which all but a δ fraction of the inputs are mapped injectively to their corresponding outputs (i.e., $\Pr_x [|C^{-1}(C(x))| > 1] < \delta$). NO instances are circuits for which at most a δ fraction of the inputs are mapped injectively (i.e., $\Pr_x [|C^{-1}(C(x))| > 1] > 1 - \delta$).

Since AI_δ was shown to be NISZK complete, to prove Theorem 2 it suffices to show a *public-coin* HVSZK batching protocol for AI_δ . Our main technical result is precisely such a protocol achieving communication roughly $k + \text{poly}(n)$, where k is the number of instances being batched. For simplicity, we first focus on the case that $\delta = 0$ – namely, distinguishing circuits that are injective from those in which no input is mapped injectively to its output. A discussion of the case of $\delta > 0$ is deferred to the end of this overview.

In order to present our approach, following [KRR+20], we will first consider a drastically easier case in which the goal is to distinguish between circuits that are *permutations* (rather than merely being injective) or are far from such.

Warmup: Batch Verification for Permutations. We first consider a variant of AI, called PERM. The inputs to PERM are length-preserving circuits $C : \{0, 1\}^n \rightarrow \{0, 1\}^n$. YES instances are circuits that compute permutations over $\{0, 1\}^n$, whereas NO instances are circuits that are *far* from being permutations in the sense that no input is mapped injectively (i.e., every image has at least two preimages).

Consider the following batch verification protocol for PERM. Given common input $C_1, \dots, C_k : \{0, 1\}^n \rightarrow \{0, 1\}^n$ the parties proceed as follows. The verifier V samples $y_k \in \{0, 1\}^n$ and sends it to the prover P . Then, P responds with an x_1 s.t. $C_k(C_{k-1}(\dots C_1(x_1))) = y_k$. The verifier checks that indeed $C_k(C_{k-1}(\dots C_1(x_1))) = y_k$ and accepts or rejects accordingly.

To argue that completeness holds, observe that when all of the C_i 's are permutations, there exists a *unique* sequence $x_1, y_1, x_2, y_2, \dots, x_{k-1}, y_{k-1}, x_k$, where $x_i = y_{i-1}$ for every $i \in [2, k]$, that is consistent with y_k . That is, $x_i = C_i^{-1}(y_i)$, for every $i \in [k]$. The prover can thus make the verifier accept (with probability 1) by sending x_1 as its message.

For soundness, let $i^* \in [k]$ denote the maximal index i s.t. C_i is a NO instance. Since C_{i^*} is a NO instance, each of its images has at least two preimages and so the size of its image is at most 2^{n-1} . Since y_k is sampled uniformly and C_{i^*+1}, \dots, C_k are permutations, we have that y_{i^*} is also random in $\{0, 1\}^n$. In particular this means that with probability at least half it has no preimage under C_{i^*} , and the verifier will eventually reject no matter what value of x_1 the prover sends. Note that the soundness error, which is $1/2$ here, can be amplified by repetition.

For zero-knowledge, consider the simulator that first samples $x_1 \in \{0, 1\}^n$ and then computes $y_k = C_k(C_{k-1}(\dots C_1(x_1)))$. Since all circuits are permutations, y_k is distributed uniformly over $\{0, 1\}^n$ as in the real interaction between the honest-verifier and the prover.

The above batch verification protocol for PERM, while simple, will be the basic underlying idea also for our batch verification protocol for Al_0 .

Public-Coin Batch Verification for Al_0 . Let $C_1, \dots, C_k : \{0, 1\}^n \rightarrow \{0, 1\}^m$ be instances of Al_0 . Since the output size of each circuit C_i is not compatible with the input size for the following circuit C_{i+1} , we cannot directly compose the circuits as we did for PERM.

A natural idea that comes to mind is to use hashing. Namely, choose a hash function¹ g to map C_i 's output $y_i \in \{0, 1\}^m$ to the next circuit input $x_{i+1} \in \{0, 1\}^n$, for every $i \in [k]$. Based on this idea it is natural to consider a minor modification of the batch verification protocol for PERM, where the only difference is that we interleave applications of g as we compose the circuits. Note that we have to hash the last circuit output y_k as well so that we can specify $x_{k+1} = g(y_k)$ to the prover as a genuine unstructured random string.

If we could guarantee that the hash function g maps the image of each circuit *injectively* into the domain of the subsequent circuit, then a similar analysis as in the protocol for PERM could be applied and we would be done. However, finding such a hash function in general seems incredibly difficult. Thus, instead, we choose g to be a random function.²

In what follows, for every $i \in [k]$, denote the image of C_i by $S_i \subseteq \{0, 1\}^m$. Note that if C_i is a YES instance (i.e., injective), the size of S_i is exactly 2^n , and if C_i is a NO instance (entirely non-injective) the size of S_i is at most 2^{n-1} .

When using a random function g as our hash function, we run into two key challenges that did not exist in the protocol for PERM. The first of these two challenges arises from the fact that for any YES instance circuit C_i , with high

¹ The specific type of hash function that we use is left vague for now and will be discussed in detail shortly.

² Jumping ahead, we note that we cannot afford for g to be entirely random, and will have to settle for some de-randomization. For the moment we ignore this issue.

probability over the choice of $g : \{0, 1\}^m \rightarrow \{0, 1\}^n$, a constant fraction of the elements in $\{0, 1\}^n$ have no preimages (under g) in the set S_i .³ If such a situation occurs for *any* of the circuits, a situation which is exceedingly likely, then even the honest prover will not be able to find a suitable preimage x_1 and we lose completeness.

The way we solve this difficulty is relatively simple: we add to the hash function g a short random auxiliary input that will be chosen independently for each of the k applications of g . We denote the auxiliary input for the i^{th} application by z_i and its length by d (which we set to $\text{polylog}(n, k)$). Observe that if $g : \{0, 1\}^m \times \{0, 1\}^d \rightarrow \{0, 1\}^n$ is chosen at random, then we expect all x 's in the domain of C_{i+1} to have roughly the same number of preimages (under g) that lie in the set $S_i \times \{0, 1\}^d$.⁴

This brings us to the second challenge in using a random hash function g , which is slightly more subtle. When considering a YES instance circuit C_i , even ignoring the additional auxiliary input, a constant fraction of the domain $\{0, 1\}^n$ of C_{i+1} will have *more than one* preimage (under g) which falls in S_i . Needless to say, this issue is further exacerbated by the addition of the auxiliary input. At first glance this may not seem like much of an issue when we consider a YES circuit. However, on further inspection, we observe that we may very well be in the case that all of the circuits except for the first circuit C_1 are YES instances (i.e., injective) and only C_1 is a NO instance (i.e., non-injective).

If such is the case, due to the collisions that occur in g , it is likely that y_k will have an exponential (in k) number of preimages x_2 that are consistent with it. If the prover has so much flexibility in its choice of x_2 then it is likely that, even though C_1 is non-injective, the prover will be able to find a consistent preimage x_1 and we lose soundness.

Borrowing an idea from [KRR+20], we solve this challenge using interaction. The high-level approach is for the prover to *commit* to x_i *before* we reveal the auxiliary information for the next circuit. Thus, the protocol proceeds iteratively, where in each iteration first the prover commits to x_i , and then the verifier reveals the auxiliary input, which the prover uses to recover y_{i-1} , and so on. The commitment has the property that as long as we are processing YES input circuits, with high probability, there is a *unique* x_i that is consistent with the interaction. In particular, when we reach the first NO instance circuit C_{i^*} (recall that i^* denotes the maximal i s.t. C_i is a NO instance) there is a unique x_{i^*+1} that is consistent with the transcript.

Distinguishing Injective Circuits from Non-Injective Circuits. Recall that our goal is to distinguish an injective circuit from a highly non-injective circuit. Following our approach thus far, assume that we have processed the circuit

³ This is similar to the fact that when throwing N balls into N bins, in expectation, a constant fraction of the bins remain empty. Here the images of C_i play the role of the balls and the elements in the domain $\{0, 1\}^n$ of C_{i+1} play the role of the bins.

⁴ Here we rely on the fact that when throwing $N \cdot \text{polylog}(N)$ balls into N bins, with high probability, all bins will contain very close to the expected $\text{polylog}(N)$ balls.

up to the circuit C_i and moreover, that there is a unique x_{i+1} that is consistent with the interaction up to this point.

Recall also that if C_i is injective then $|S_i| = 2^n$, whereas if it is non-injective then $|S_i| \leq 2^{n-1}$. Thus, our approach is to employ a set lower bound protocol (a la [GS89]) as follows. The verifier chooses a “filter” function $h_i \in H_n$, where H_n is a family of pairwise independent hash functions from domain $\{0, 1\}^m \times \{0, 1\}^d$ to an appropriately chosen range size, as well as a target element α_i . If C_i is injective then we expect each x_{i+1} to have very close to 2^d preimages $(y_i, z_i) \in S_i \times \{0, 1\}^d$ under g . On the other hand, if C_i is non-injective, then we expect each x_{i+1} to have roughly 2^{d-1} such preimages or less. Thus, by setting the range size to be roughly 2^d , the probability that one of these preimages hashes correctly via h_i to α_i is larger (by a constant) in the YES case than in the NO case.

Balancing Completeness, Soundness and Zero-Knowledge. At this point we observe that even if g and h_i were random functions, the set lower bound approach only yields a small constant gap between completeness and soundness. This is insufficient since the completeness error is accumulating across the k different circuits. Note that we cannot afford a k -fold repetition of the set lower bound protocol (since it would be prohibitively expensive in our parameter setting). Moreover, since we cannot generically amplify zero-knowledge, we also need the zero-knowledge error accumulated by the YES instance circuits to be negligible.

We resolve this issue by considering a new variant of the approximate injectivity problem which we denote by $\text{Al}_{L,\delta}$. The YES instances in this variant are identical to the YES instance in Al_δ – namely circuits that are injective on all but a δ fraction of their domain. However, a circuit C is a NO instance if at least $1 - \delta$ fraction of its inputs have at least L “siblings” (i.e., inputs that are mapped to the same output). Thus, the standard Al_δ problem corresponds to the case $L = 2$. Using a large enough L increases the gap between the number of preimages in YES and NO instances, letting us set the range size of the h_i ’s to obtain a larger gap between completeness and soundness.

We show that $\text{Al}_{2,\delta}$ reduces to $\text{Al}_{L,\delta'}$, where $L = 2^{\text{polylog}(n,k)}$ and δ' is related to δ . The idea for the reduction is to simply concatenate sufficiently many copies of the input circuit.⁵ Thus, in the sequel it suffices to consider batch verification for k instances of $\text{Al}_{L,0}$ (and the case that $\delta > 0$ will be discussed later on).

Over-Simplified Batch Verification Protocol for $\text{Al}_{L,0}$. With the foregoing insights in mind, consider the following over-simplified batching protocol for $\text{Al}_{L,0}$ (see also the diagram in Fig. 1 which gives a bird’s eye view of the flow of the protocol).

1. V samples g and $x_{k+1} \leftarrow \{0, 1\}^n$ and sends both to P .
2. For $i = k, \dots, 1$:

⁵ In more detail, we transform an instance C of $\text{Al}_{2,\delta}$ to an instance C' of $\text{Al}_{2^\ell,\delta \cdot \ell}$ by concatenating ℓ copies of C . It is not hard to see that if C were (almost) injective then C' is (almost) injective. But, if (almost) every image of C has at least two preimages then (almost) every image of C' has at least 2^ℓ preimages.

- (a) V samples a filter function $h_i \leftarrow H_n$ and target value α_i (of appropriate length) and sends both to P .
- (b) P selects at random a pair (x_i, z_i) s.t.
 - i. $g(C_i(x_i), z_i) = x_{i+1}$
 - ii. $h_i(C_i(x_i), z_i) = \alpha_i$
 and sends (x_i, z_i) to V .
- 3. V sets $x'_1 = x_1$ and for $i = 1, \dots, k$:
 - (a) Computes $y'_i = C_i(x'_i)$.
 - (b) Verifies that $h_i(y'_i, z_i) = \alpha_i$.
 - (c) Computes $x'_{i+1} = g(y'_i, z_i)$.
 - (d) Verifies that $x'_{i+1} = x_{i+1}$.
- 4. If all of V 's checks pass then she accepts. Otherwise she rejects.




Fig. 1. Sampling Process of the Protocol

Note that the communication complexity in this protocol is $\Omega(k \cdot n)$ (since the prover sends x_1, \dots, x_k), which is more than we can afford. This issue will be resolved shortly and so we ignore it for now.

For completeness, if all C_i 's are injective, in every iteration i , given that there exists a consistent x_i (i.e., x_i for which there exists z where $g(C_i(x_i), z) = x_{i+1}$ and $h_i(C_i(x_i), z) = \alpha_i$), w.h.p. there exists also a consistent x_{i-1} . Hence, it is not hard to show, by induction, that with high probability, after the last iteration there exists an x_1 that passes the verifier's checks, and so V accepts.

For soundness, consider the iteration i^* which, as defined in the protocol for PERM, denotes the maximal index of a NO instance circuit. Since the number of preimages of C_{i^*} is less than $2^n/L$, by the foregoing discussion it is very likely that there does not exist a pair (x_{i^*}, z_{i^*}) that is consistent with the rest of the messages, i.e., s.t. $g(C_{i^*}(x_{i^*}), z_{i^*}) = x_{i^*+1}$ and $h_{i^*}(C_{i^*}(x_{i^*}), z_{i^*}) = \alpha_{i^*}$, and therefore V will eventually reject.

Before arguing why the protocol is zero-knowledge, we first discuss the hash function family that g is sampled from.

The Hash Function g . One important consideration that arises when derandomizing g is that a cheating prover P^* has some flexibility in the choice of the x_i

that she sends for rounds $i > i^*$. Indeed, by design there will be many such x_i 's. While the honest prover should choose x_i at random (from this set), a cheating prover may try to cleverly choose x_i 's that help her cheat. Since all these choices are made after the function g was revealed, we cannot assume that x_{i^*+1}, \dots, x_k are uniformly random relative to g . Thus, for our analysis it does not suffice that a random x_{i+1} has a suitable number of preimages under g .

Instead, we will seek a much stronger, worst-case guarantee, from g . Specifically, that for every $x_{i+1} \in \{0, 1\}^n$, the number of preimages $(y, z) \in g^{-1}(x_{i+1})$, where $y \in S_i$, is close to its expectation, i.e., around 2^d for $C_i \in \text{YES}$ and around $2^d/L$ for $C_i \in \text{NO}$.

As shown by Alon *et al.* [ADM+99], this type of guarantee is not offered by pairwise independent hash function or even more generally by randomness extractors. On the other hand, what gives us hope is that a totally random function does have such a worst-case guarantee. Thus, we wish to construct a small hash function family G where with high probability over the choice of $g \leftarrow G$, every image has roughly the same number of preimages from a predetermined set.

Following a work of Celis *et al.* [CRSW11], we refer to such functions as *load-balancing hash functions*. A function in this family maps the set $\{0, 1\}^m$ together with some auxiliary input $\{0, 1\}^d$ to the set $\{0, 1\}^n$. We require that for any set $S \subseteq \{0, 1\}^n$ s.t. $|S| \geq 2^n/L$ and for every $x \in \{0, 1\}^n$, w.h.p over $g \leftarrow G_n$, it holds that $|\{(y, z) : y \in S, g(y, z) = x\}|$ is roughly $\frac{|S| \cdot 2^d}{2^n}$. We construct a suitable load-balancing functions based on $\text{poly}(n)$ -wise independent hash functions (combined with almost pairwise independent hash functions).

Honest Verifier Statistical Zero-Knowledge. To argue zero-knowledge, consider a simulator that first samples an initial x_1 . Then, in each iteration i the simulator samples z_i , computes $y_i = C_i(x_i)$ and $x_{i+1} = g(y_i, z_i)$. Then it samples h_i and computes $\alpha_i = h_i(y_i, z_i)$.




Fig. 2. Direction of Protocol Progress vs. Simulator Progress – solid lines represent the protocol, and dashed lines represent the simulator.

Statistical zero-knowledge is not obvious since the protocol and the simulator progress in opposite directions. The simulation progress direction is from circuit

C_1 up to circuit C_k , while the protocol progress direction is from circuit C_k down to circuit C_1 , as shown in Fig. 2. However, due to the special property of the the load-balancing function g , we manage to achieve statistical zero-knowledge.

We define the hybrid distribution H_i that is sampled as follows:

1. Sample $g \in G_n$ and $x_i \in \{0, 1\}^n$.
2. Generate $(x_j, z_j, h_j, \alpha_j)_{j \in \{1, \dots, i-1\}}$ according to the protocol.
3. Generate $(z_j, h_j, \alpha_j, x_{j+1})_{j \in \{i, \dots, k\}}$ according to the simulator.
4. Output g, x_{k+1} and $(x_j, z_j, h_j, \alpha_j)_{j \in \{1, \dots, k\}}$.

Note that the simulator distribution is identical to H_1 while the protocol distribution is identical to H_{k+1} . We bound the statistical difference between H_i and H_{i+1} and use the hybrid argument.

Note that conditioned on $(g, x_i, z_i, h_i, \alpha_i, x_{i+1})$ the rest of the variables in H_i, H_{i+1} are distributed identically. Therefore it is enough to bound the statistical differences between those variables as sampled in H_i and H_{i+1} . Since g is sampled identically in both hybrids, we can fix it and bound the statistical difference of those variables conditioned on some g .

For the hybrid H_{i+1} , the variables $(x_i, z_i, h_i, \alpha_i, x_{i+1})$ are sampled according to the protocol whereas for the hybrid H_i , the variables $(x_i, z_i, h_i, \alpha_i, x_{i+1})$ are sampled according to the simulator. Let X_S, X_P and X_U denotes the distributions of any random variable X according to the simulator, protocol and the uniform distribution respectively.

Consider the distribution of $(x_i, z_i, h_i, \alpha_i, x_{i+1})_P$ which is sampled according to the protocol, i.e., (h_i, α_i, x_{i+1}) are sampled uniformly at first, and then (x_i, z_i) are chosen uniformly from the set of (x, z) s.t. $g(C_i(x), z) = x_{i+1}$ and $h_i(C_i(x), z) = \alpha_{i+1}$. Consider also the distribution $(x_i, z_i, h_i, \alpha_i, x_{i+1})_S$ which is sampled according to the simulator, i.e., (x_i, z_i, h_i) are sampled uniformly at first and then it sets $x_{i+1} = g(C_i(x_i), z_i)$ and $\alpha_i = h_i(C_i(x_i), z_i)$. Note that the distributions $(x_i, z_i)_P$ and $(x_i, z_i)_S$ are identical conditioned on specific values of (h_i, α_i, x_{i+1}) , and therefore, it is enough to bound $\Delta((h_i, \alpha_i, x_{i+1})_P, (h_i, \alpha_i, x_{i+1})_S)$.

We define the function $\varphi_i(x_i, z_i, h_i) = (h_i, \alpha_i, x_{i+1})$, where $x_{i+1} = g(C_i(x_i), z_i)$ and $\alpha_i = h_i(C_i(x_i), z_i)$. Note that

$$\Delta((h_i, \alpha_i, x_{i+1})_P, (h_i, \alpha_i, x_{i+1})_S) = \Delta\left((h_i, \alpha_i, x_{i+1})_U, \varphi_i((x_i, z_i, h_i)_U)\right).$$

Therefore, what we have left is to show that φ_i 's output on uniform input is close to uniform.

Consider uniform (x_i, z_i, h_i) and set $x_{i+1} = g(C_i(x_i), z_i)$ and $\alpha_i = h_i(C_i(x_i), z_i)$. x_{i+1} is close to uniform due to the special property of g that every x_{i+1} has roughly the same number of preimages (y_i, z_i) , and due to the fact that each image of C_i has exactly one preimage. Now fix x_{i+1} . the number of pairs (x_i, z_i) that are mapped to x_{i+1} is roughly 2^d , i.e., there are d bits of entropy on which h_i is applied. Therefore, since h_i is pairwise independent hash function and as such, also a strong extractor, we get that (h_i, α_i) is also close to uniform.

Reducing Communication via Hashing. The foregoing soundness analysis relied on the fact that the prover sent the values (x_1, \dots, x_k) during the interaction. However, as noted above, this requires $n \cdot k$ communication which we cannot afford. In a nutshell, we resolve this inefficiency by having the prover only send short hashes of the x_i 's, details follows.

At the beginning of each round, the verifier V sends a description of a pairwise independent hash function f_i in addition to the filter function h_i and target value α_i . Then, in function to sending z_i , the prover P in her turn also sends a hash value $\beta_i = f_i(x_i)$ (rather than sending x_i explicitly). In order for the hash value to commit the prover to x_i , we would like for the hash function f_i to be injective on the set of consistent (x, z) 's (i.e., those for which $g(C_i(x), z) = x_{i+1}$ and $h_i(C_i(x), z) = \alpha_i$). This set is of size roughly 2^d where $d = \text{polylog}(n, k)$. Therefore, setting the output size of f_i to be $\text{poly}(d)$ ($= \text{polylog}(n, k)$) bits is sufficient. At the very end of the protocol, the prover P still needs to explicitly send x'_1 , so that V can compute $y'_i = C_i(x'_i)$ and $x'_{i+1} = g(y'_i, z_i)$, for every $i \in [k]$, and verify that $h_i(y'_i, z_i) = \alpha_i$, and lastly that $x'_{k+1} = x_{k+1}$. Note that the verifier can only perform these checks at the end of the interaction (as she did in the simplified protocol) since she must obtain the value of x'_1 in order to generate x'_2, \dots, x'_k .

Overall, the communication is dominated by the values x_{k+1} and g , which are sent by the verifier, and the value x'_1 sent by the prover. Each of these messages has length $\text{poly}(n, \log(k))$. All the rest of the messages including the specification of the hash functions h_i, f_i as well as the values α_i, β_i and z_i have length $\text{polylog}(n, k)$. Overall we obtain the desired communication complexity $(k + \text{poly}(n)) \cdot \text{polylog}(n, k)$.

When $\delta > 0$. For the case where $\delta > 0$, the arguments we made for completeness and zero-knowledge still hold in a straightforward manner, but we need to be more careful about soundness. More specifically, for some YES instance circuit C_i (where $i > i^*$) and $x_{i+1} \in \{0, 1\}^n$, there potentially can exist a pair $(y_i, z_i) \in g^{-1}(x_{i+1})$ s.t. $h_i(y_i, z_i) = \alpha_i$ and y_i has an exponential number of preimages. Therefore, the set of consistent (x, z) is of exponential size and therefore, in order to fix the chosen x_i by the prover, β_i must consist polynomial number of bits, which is of course too expensive for us.

However, since there is only a small number of images y_i that have more than one preimage (δ fraction of 2^n), there is also only a small number of pairs $(y_i, z_i) \in g^{-1}(x_{i+1})$ where y_i is such an image. Therefore, w.h.p over (h_i, α_i) , none of those problematic pairs satisfy the condition that $h_i(y_i, z_i) = \alpha_i$, and therefore, their preimages are inconsistent which allows the earlier setting of the output size of f_i to work.

Comparison to the [KRR+20] Protocol. Our public-coin batching protocol bears some resemblance to the *private-coin* batching protocol of [KRR+20]. We highlight here the similarity and differences between our approaches. We note that readers who are unfamiliar with [KRR+20] can safely skip this discussion and proceed directly to Sect. 1.2.

In the protocol of [KRR+20] the verifier V first samples x_1 and auxiliary randomnesses (z_1, \dots, z_k) as part of her setup. She computes $y_i = C_i(x_i)$ and determines the next circuit input as $x_{i+1} = g(y_i, z_i)$, for every $i \in [k]$, where in contrast to our protocol, the function g is simply a (strong) seeded randomness extractor (where y_i is the min-entropy source and z_i is the extractor’s seed).

The actual interaction starts by having the verifier V send y_k to P . The parties then proceed in iterations where in each iteration i , given x_{i+1} , the prover needs to guess x_i using some additional hints that the verifier supplies. The prover’s guesses are communicated by sending a short hash.

While their protocol bears some similarity to ours, we emphasize several fundamental ways in which our approach differs from that of [KRR+20] (beyond the fact that our protocol is public-coin).

- In [KRR+20] the verifier herself samples (z_1, \dots, z_k) and using it computes (x_1, \dots, x_k) as part of her setup, and then she reveals these gradually. This means that in the [KRR+20] there is a ground truth that the verifier can compare. In contrast, in our protocol, each x_i is chosen via an interactive process that involves both parties and happens “online”. In particular, and as discussed above, this means that a cheating prover can bias the distribution of the x_i ’s as we process the circuit.
- On a related note, if the prover commits to a wrong x_i in some iteration i , then, in [KRR+20], the verifier V can immediately detect this and reject. In contrast, in our protocol we are unable to do so and must wait to detect this at the very end of the interaction.
- In our protocol the x_i ’s are computed in reverse order starting from x_k down to x_1 , whereas in the protocol of [KRR+20] the x_i ’s are computed in order, starting from x_1 up to x_k . This may seem like a minor difference but turns out to complicate matters significantly when considering zero-knowledge. Indeed, both the [KRR+20] simulator as well as our simulator compute the x_i ’s in order. This means that in the current work the protocol and the simulator, operate in *different* order. This makes the analysis of the simulation significantly more challenging.
- Lastly, [KRR+20] use extractors in order to map each circuit output y_i to the next circuit input x_{i+1} . As discussed in detail in the above technical overview, the average-case guarantee provided by extractors are not good enough for us and we rely on the stronger notion of load-balancing hash functions.

Efficient Transformation to Public-Coin Malicious Verifier SZK. Goldreich, Sahai, and Vadhan [GSV98] showed that any public-coin HVSZK protocol can be transformed into a public-coin SZK protocol. Applying their transformation to the public-coin honest-verifier batch verification protocol described above would indeed result in a malicious-verifier SZK batch verification protocol for AI, and thus NISZK. This transformation, however, starts by repeating the HVSZK protocol several times in parallel in order to make the soundness error exponentially small in the number of rounds. This would incur a blowup in communication by a factor of $\Omega(k)$, which we cannot afford.

In order to get around this, we show that the transformation of [GSV98], when used on protocols with a stronger soundness guarantee called *round-by-round soundness* [CCH+19], can be performed without the initial repetition step, and thus achieve a much smaller blowup in communication. Then we show that our honest-verifier batching protocol does indeed provide this guarantee, and thus the transformation can be applied to it with this better blowup, giving us the desired result. We now briefly describe the transformation, the round-by-round soundness property, and how they fit together.

The GSV Transformation. Recall that in a public-coin HVSZK protocol, the honest verifier’s messages consist of uniformly random strings. What breaks when the verifier is malicious is that it might choose these strings arbitrarily rather than uniformly at random. In the GSV transformation, the prover and verifier essentially run the given HVSZK protocol, but instead of the verifier sending these random strings, they are sampled by the prover and the verifier together using a Random Selection (RS) protocol. This protocol, which is constructed by [GSV98], uses four messages, is public-coin, and produces as output a string r of some desired length ℓ . It provides, very roughly, the following guarantees when run with security parameter λ :

1. If the prover is honest, the distribution of r is $2^{-\lambda}$ -close to uniform over $\{0, 1\}^\ell$, and the transcript of the protocol is simulatable given output r .
2. If the verifier is honest, for any set $T \subseteq \{0, 1\}^\ell$, we have $\Pr[r \in T] \leq 2^\lambda \cdot |T|/2^\ell$.

The first property above ensures that the resulting protocol is complete and zero-knowledge. The second property ensures that the prover cannot skew the distribution of r , and thus soundness is maintained. Following the analysis in [GSV98], however, it turns out that if the original protocol has soundness error s and r rounds, the bound obtained on the soundness error of the new protocol is roughly $2^{r\lambda} \cdot s$. Thus, we would need to start by decreasing the soundness error of the HVSZK protocol to less than $2^{-r\lambda}$. The only way we know to do this generically is by repetition, which results in a multiplicative blowup of at least $\Omega(r)$ in communication – in our case this is $\Omega(k)$, which is too large for us. We get around this by showing that our batch verification protocol has a stronger soundness property that results in a much better bound on the soundness error when this transformation is applied.

Round-by-Round Soundness. Typically, the soundness property in an interactive proof places requirements on how likely it is that a verifier accepts on a NO input. Round-by-round soundness, introduced by Canetti et al. [CCH+19], instead places requirements on intermediate stages of the protocol. It involves a mapping State from partial transcripts of the protocol to the set $\{\text{alive}, \text{doomed}\}$. A protocol is ϵ -*round-by-round sound* if there exists a mapping State such that for any input x and partial transcript τ with $\text{State}(x, \tau) = \text{doomed}$, and any subsequent prover message α , the probability over the next verifier message β

that $\text{State}(x, \tau, \alpha, \beta) = \text{alive}$ is at most ϵ . Further, any complete transcript that is **doomed** always results in a rejection by the verifier, and for any NO instance x , it is the case that $\text{State}(x, \perp) = \text{doomed}$.

In other words, at a point where a protocol is **doomed**, irrespective of what the prover does, the set of “bad” verifier messages that make the protocol **alive** in the next round has relative size at most ϵ . To see its implications for standard soundness, consider a protocol that has r rounds and is ϵ -round-by-round sound. The probability that the verifier accepts on a NO instance is at most the probability that the complete transcript is **alive**. Thus, since the protocol has to go from **doomed** to **alive** in at least one of the r rounds, the probability that the verifier accepts is at most $r\epsilon$.

Putting it Together. Now, consider passing an ϵ -round-by-round sound protocol through the GSV transformation described above. Here, each verifier message is replaced by the output of the RS protocol. On a NO instance, in order to successfully cheat, the prover has to make at least one of the r verifier messages fall into the “bad” set for that round. Each of these bad sets, however, has relative size at most ϵ , and thus the RS protocol’s output falls in each with probability at most $\epsilon 2^\lambda$. Thus, that total probability that the prover successfully cheats is at most $\epsilon r 2^\lambda$.

So, if the original protocol had round-by-round soundness error somewhat smaller than $(r 2^\lambda)^{-1}$, the resulting protocol would still be sound. This is a much more modest requirement than before, and can be achieved with a multiplicative blowup of at most $O(\lambda \log r)$ in communication. Completeness and zero-knowledge follow from the other properties of the RS protocol, and setting λ to be $\text{polylog}(n, k)$ gives us the desired SZK protocol.

Round-by-Round Soundness of our HVSZK Batching Protocol. To show that our protocol described earlier has round-by-round soundness, we define the **State** function as follows. Consider the partial transcript τ_j that corresponds to its j ’th iteration – this consists of $g, x_{k+1}, (h_i, \alpha_i, f_i, z_i, \beta_i)_{i \in \{k, \dots, j+1\}}$, and (h_j, α_j, f_j) .

- For $j > i^*$: Output **doomed** on the transcript τ_j if, for any prover message (z_j, β_j) that follows, there exists at most one preimage x_j that is consistent with τ_j and (z_j, β_j) . Further, if there is no such x_j , output **doomed** on all future transcripts that extend τ_j .
 - Consider a **doomed** transcript τ_{j+1} , a prover message (z_{j+1}, β_{j+1}) , and the unique x_{j+1} that is consistent with them as promised (if it doesn’t exist, future transcripts are already **doomed**). By our earlier discussion, w.h.p., there are very few x_j ’s which have a z_i such that $g(C_j(x_j), z_j) = x_{j+1}$ and $h_j(C_j(x_j), z_j) = \alpha_j$. Therefore, w.h.p., f_j maps these x_j ’s injectively. Hence for any prover message (z_{j+1}, β_{j+1}) , we have that w.h.p. $\tau_{j+1}, (z_{j+1}, \beta_{j+1}), (h_j, \alpha_j, f_j)$ is also **doomed**.
- For $j = i^*$: We define the **State** function to output **doomed** on the transcript τ_{i^*} if for any prover message (z_{i^*}, β_{i^*}) , there does not exist x_{i^*} that is consistent with τ_{i^*} and (z_{i^*}, β_{i^*}) .

- Consider any **doomed** transcripts τ_{i^*+1} , a prover message $(z_{i^*+1}, \beta_{i^*+1})$, and the unique x_{i^*+1} that is consistent with them. As discussed earlier, w.h.p., there does not exist (x_{i^*}, z_{i^*}) s.t. $g(C_{i^*}(x_{i^*}), z_{i^*}) = x_{i^*+1}$ and $h_{i^*}(C_{i^*}(x_{i^*}), z_{i^*}) = \alpha_{i^*}$ and therefore for every prover message $(z_{i^*+1}, \beta_{i^*+1})$, it holds w.h.p. that $\tau_{i^*+1}, (z_{i^*+1}, \beta_{i^*+1}), (h_{i^*}, \alpha_{i^*}, f_{i^*})$ is **doomed** too.
- For $j < i^*$: We define the **State** function to answer according to the partial transcript τ_{i^*} , and therefore round-by-round soundness in this case is immediate.

We set $\text{State}(x, \perp)$ to be **doomed** if x is a NO instance, and anything that is not **doomed** is set to be **alive**. Lastly, consider a complete transcript that is **doomed**; there does not exist x_{i^*} that is consistent with the beginning of the transcript, and therefore \mathbb{V} must reject. This function now witnesses the round-by-round soundness of our protocol.

1.3 Organization

We start with preliminaries in Sect. 2. In Sect. 3 we formalize our notion of a *load-balancing* hash function and provide a construction based on k -wise independent hash functions. In Sect. 4 we introduce our variant of the approximate injectivity (AI) problem and show that it is NISZK-complete. In Sect. 5 we construct our *public-coin honest-verifier* batch verification for AI. In Sect. 6 we show a generic, and efficient, transformation from public-coin HVSZK (having round-by-round soundness) to public-coin full-fledged SZK. Lastly, in Sect. 7 we use the results obtained in the prior sections to obtain our public-coin SZK batch verification protocol for NISZK.

Due to space restrictions, most of the proofs are deferred to the full version [KRV21].

2 Preliminaries

For a finite non-empty set S , we denote by $x \leftarrow S$ a uniformly distributed element in S . We also use U_ℓ to denote a random variable uniformly distributed over $\{0, 1\}^\ell$.

For a distribution X over a finite set U and a (non-empty) event $E \subseteq U$, we denote by $X|_E$ the distribution obtained by conditioning X on E : namely, $\Pr[X|_E = u] = \Pr[X = u | E]$, for every $u \in U$. The *support* of X is defined as $\text{supp}(X) = \{u \in U : \Pr[X = u] > 0\}$.

Definition 1. Let $\Pi = (\text{YES}, \text{NO})$ be a promise problem, where $\text{YES} = (\text{YES}_n)_{n \in \mathbb{N}}$ and $\text{NO} = (\text{NO}_n)_{n \in \mathbb{N}}$, and let $k = k(n) \in \mathbb{N}$. We define the promise problem $\Pi^{\otimes k} = (\text{YES}^{\otimes k}, \text{NO}^{\otimes k})$, where $\text{YES}^{\otimes k} = (\text{YES}_n^{\otimes k})_{n \in \mathbb{N}}$, $\text{NO}^{\otimes k} = (\text{NO}_n^{\otimes k})_{n \in \mathbb{N}}$ and

$$\text{YES}_n^{\otimes k} = (\text{YES}_n)^k$$

and

$$\text{NO}_n^{\otimes k} = (\text{YES}_n \cup \text{NO}_n)^k \setminus (\text{NO}_n)^k.$$

The statistical distance between two distributions P and Q over a finite set U is defined as $\Delta(P, Q) = \max_{S \subseteq U} (P(S) - Q(S)) = \frac{1}{2} \sum_{u \in U} |P(u) - Q(u)|$.

2.1 Statistical Zero-Knowledge

We use $(P, V)(x)$ to refer to the *transcript* of an execution of an interactive protocol with prover P and verifier V on common input x . The **transcript** includes the input x , all messages sent by P to V in the protocol and the verifier’s random coin tosses. We say that the transcript $\tau = (P, V)(x)$ is **accepting** if at the end of the corresponding interaction, the verifier accepts.

Definition 2 (Interactive proof)

Let $c = c(n) \in [0, 1]$ and $s = s(n) \in [0, 1]$. An interactive proof with completeness error c and soundness error s for a promise problem $\Pi = (\Pi_{\text{YES}}, \Pi_{\text{NO}})$, consists of a probabilistic polynomial-time verifier V and a computationally unbounded prover P such that following properties hold:

- **Completeness:** For any $x \in \Pi_{\text{YES}}$:

$$\Pr [(P, V)(x) \text{ is accepting}] \geq 1 - c(|x|).$$

- **Soundness:** For any (computationally unbounded) cheating prover P^* and any $x \in \Pi_{\text{NO}}$:

$$\Pr [(P^*, V)(x) \text{ is accepting}] \leq s(|x|).$$

We denote this proof system by the pair (P, V) .

In this paper we focus on *public-coin* interactive proofs. An interactive proof (P, V) is **public-coin** if the verifier’s messages are selected independently and uniformly at random (and their lengths are fixed independently of the interaction).

We next define *honest-verifier* statistical zero-knowledge proofs, in which zero-knowledge is only guaranteed wrt the *honest* (i.e., prescribed) verifier’s behavior.

Definition 3 (HVSZK). Let $z = z(n) \in [0, 1]$. An interactive proof system (P, V) is an **Honest Verifier SZK Proof-System (HVSZK)**, with zero-knowledge error z , if there exists a probabilistic polynomial-time algorithm Sim (called the **simulator**) such that for any $x \in \Pi_{\text{YES}}$:

$$\Delta ((P, V)(x), \text{Sim}(x)) \leq z(|x|).$$

For the malicious verifier SZK definition, we allow the verifier access to non-uniform advice. Therefore, we also provide the simulator with the same advice. Let $\text{Sim}_{[a]}$ denote the simulator Sim given access to the some advice string $a \in \{0, 1\}^*$.

Definition 4 (SZK). Let $z = z(n) \in [0, 1]$. An interactive-proof (P, V) is a statistical zero-knowledge proof-system (SZK), with zero-knowledge error z , if for every probabilistic polynomial-time verifier V^* there exists an algorithm Sim (called the simulator) that runs in (strict) polynomial time such that for any $x \in \Pi_{\text{YES}}$ and $a \in \{0, 1\}^*$:

$$\Delta \left((P, V_{[a]}^*)(x), \text{Sim}_{[a]}(x) \right) \leq z(|x|).$$

If the completeness, soundness and zero-knowledge (resp., honest-verifier zero-knowledge) errors are all negligible, we simply say that the interactive proof is an SZK (resp., HVSZK) protocol. We also use SZK (resp., HVSZK) to denote the class of promise problems having such an SZK (resp., HVSZK) protocol.

Non-Interactive Statistical Zero-Knowledge Proofs. We also define the non-interactive variant of SZK as follows.

Definition 5 (NISZK). Let $c = c(n) \in [0, 1]$, $s = s(n) \in [0, 1]$ and $z = z(n) \in [0, 1]$. A non-interactive statistical zero-knowledge proof (NISZK) with completeness error c , soundness error s and zero-knowledge error z for a promise problem $\Pi = (\Pi_{\text{YES}}, \Pi_{\text{NO}})$, consists of a probabilistic polynomial-time verifier V , a computationally unbounded prover P and a polynomial $\ell = \ell(n)$ such that following properties hold:

- **Completeness:** For any $x \in \Pi_{\text{YES}}$:

$$\Pr_{r \in \{0,1\}^{\ell(|x|)}} [V(x, r, \pi) \text{ accepts}] \geq 1 - c(|x|),$$

where $\pi = P(x, r)$.

- **Soundness:** For any $x \in \Pi_{\text{NO}}$:

$$\Pr_{r \in \{0,1\}^{\ell(|x|)}} [\exists \pi^* \text{ s.t. } V(x, r, \pi^*) \text{ accepts}] \leq s(|x|).$$

- **Honest Verifier Statistical Zero Knowledge:** There is a probabilistic polynomial-time algorithm Sim (called the simulator) such that for any $x \in \Pi_{\text{YES}}$:

$$\Delta ((U_\ell, P(x, U_\ell)), \text{Sim}(x)) \leq z(|x|).$$

(Note that the zero-knowledge property in Definition 5 is referred to as “honest-verifier” simply because the verifier does not send any messages to the prover and so it is meaningless to consider malicious behavior.)

As above, if the errors are negligible, we say that Π has a NISZK protocol and use NISZK to denote the class of all such promise problems.

2.2 Many-wise Independence

Definition 6 (δ -almost ℓ -wise Independent Hash Functions). For $\ell = \ell(n) \in \mathbb{N}$, $m = m(n) \in \mathbb{N}$ and $\delta = \delta(n) > 0$, a family of functions $\mathcal{F} = (\mathcal{F}_n)_n$, where $\mathcal{F}_n = \{f : \{0, 1\}^m \rightarrow \{0, 1\}^n\}$ is δ -almost ℓ -wise independent if for every $n \in \mathbb{N}$ and distinct $x_1, x_2, \dots, x_\ell \in \{0, 1\}^m$ the distributions:

- $(f(x_1), \dots, f(x_\ell))$, where $f \leftarrow \mathcal{F}_n$; and
- The uniform distribution over $(\{0, 1\}^n)^\ell$,

are δ -close in statistical distance.

When $\delta = 0$ we simply say that the hash function family is ℓ -wise independent. Constructions of (efficiently computable) many-wise hash function families with a very succinct representation are well known. In particular, when $\delta = 0$ we have the following well-known construction:

Lemma 1 (See, e.g., [Vad12, Section 3.5.5]). For every $\ell = \ell(n) \in \mathbb{N}$ and $m = m(n) \in \mathbb{N}$ there exists a family of ℓ -wise independent hash functions $\mathcal{F}_{n,m}^{(\ell)} = \{f : \{0, 1\}^m \rightarrow \{0, 1\}^n\}$ where a random function from $\mathcal{F}_{n,m}^{(\ell)}$ can be selected using $O(\ell \cdot \max(n, m))$ bits, and given a description of $f \in \mathcal{F}_{n,m}^{(\ell)}$ and $x \in \{0, 1\}^m$, the value $f(x)$ can be computed in time $\text{poly}(n, m, \ell)$.

For $\delta > 0$, the seminal work of Naor and Naor [NN93] yields a highly succinct construction.

Lemma 2 ([NN93, Lemma 4.2]). For every $\ell = \ell(n) \in \mathbb{N}$, $m = m(n) \in \mathbb{N}$ and $\delta = \delta(n) > 0$, there exists a family of δ -almost ℓ -wise independent hash functions $\mathcal{F}_{n,m}^{(\ell)} = \{f : \{0, 1\}^m \rightarrow \{0, 1\}^n\}$ where a random function from $\mathcal{F}_{n,m}^{(\ell)}$ can be selected using $O(\ell \cdot n + \log(m) + \log(1/\delta))$ bits, and given a description of $f \in \mathcal{F}_{n,m}^{(\ell)}$ and $x \in \{0, 1\}^m$, the value $f(x)$ can be computed in time $\text{poly}(n, m, \ell, \log(1/\delta))$.

2.3 Round-By-Round Soundness

In this section we define the notion of *round-by-round soundness* of interactive proofs, as introduced in the recent work of Canetti *et al.* [CCH+19].

Let (P, V) be a public-coin interactive proof. We denote by $V(x, \tau)$ the distribution of the next message (or output) of V on the input x and partial transcript τ .

Definition 7. Let (P, V) be a public-coin interactive proof for the promise problem $\Pi = (\Pi_{\text{YES}}, \Pi_{\text{NO}})$.

We say that (P, V) has a round-by-round soundness error $\epsilon = \epsilon(n)$ if there exists some (possibly inefficient) function State that takes as input the main input x and a partial transcript τ and outputs either *alive* or *doomed* and has the following properties:

1. If $x \in \text{NO}$, then $\text{State}(x, \perp) = \text{doomed}$ (where \perp denotes the empty transcript).
2. For any transcript prefix τ , if $\text{State}(x, \tau) = \text{doomed}$, then for any prover message α it holds that

$$\Pr_{\beta \leftarrow V(x, \tau, \alpha)} [\text{State}(x, \tau, \alpha, \beta) = \text{alive}] \leq \epsilon(n).$$

3. For any full transcript τ (i.e., a transcript in which the verifier halts) such that $\text{State}(x, \tau) = \text{doomed}$, it holds that $V(x, \tau)$ is rejecting.

Canetti *et al.* [CCH+19] also show the following simple fact (which follows from the union bound).

Fact 4. Let (P, V) be a $2r$ -message interactive proof with round-by-round soundness error ϵ . Then, (P, V) has standard soundness error $r \cdot \epsilon$.

3 Load-Balancing Functions

We now define *load-balancing hash functions*, a central tool in our construction. Loosely speaking, a load balancing hash function is a function mapping a set $\{0, 1\}^m$ together with a short auxiliary random string $\{0, 1\}^d$ to a range $\{0, 1\}^n$. The key property that we seek is that for every subset of $\{0, 1\}^m$ of size roughly 2^n it holds that every element $x \in \{0, 1\}^n$ has roughly the same number of preimages $(y, z) \in S \times \{0, 1\}^d$.

Definition 8 (Load Balancing Hash Function Family). Let $m = m(n) \in \mathbb{N}$, $d = d(n) \in \mathbb{N}$ and $\epsilon : \mathbb{N}^4 \rightarrow [0, 1]$. We say that a family of hash functions $G = (G_n)_n$, where $G_n = \left\{ g : \{0, 1\}^m \times \{0, 1\}^d \rightarrow \{0, 1\}^n \right\}$, is (d, ϵ) -load-balancing, if for every $n \in \mathbb{N}$, number of elements $v \in \mathbb{N}$, and set $S \subseteq \{0, 1\}^{m(n)}$ of size $|S| \leq 2^n$ it holds that:

$$\Pr_{g \leftarrow G} \left[\exists x \in \{0, 1\}^n : \left| L_{S,g}(x) - \frac{|S| \cdot 2^d}{2^n} \right| > v \right] \leq \epsilon(n, |S|, v, d),$$

where $L_{S,g}(x) = \left| \left\{ (y, z) \in S \times \{0, 1\}^d : g(y, z) = x \right\} \right|$.

Lemma 3. For any values $n, \lambda \in \mathbb{N}$ and $m = m(n), d = d(n)$, there is an explicit family of hash functions $G = (G_n)_n$ that is (d, λ, ϵ) -load-balancing, where $G_n = \left\{ g : \{0, 1\}^{m(n)} \times \{0, 1\}^{d(n)} \rightarrow \{0, 1\}^n \right\}$ and

$$\epsilon_\lambda(n, |S|, v, d) = 2^n \cdot \left(64 \cdot (n + \lambda) \cdot \frac{\mu + (n + \lambda)}{v^2} \right)^{n + \lambda + 4} + 2^{-\lambda - 1},$$

where $\mu = \frac{|S| \cdot 2^d}{2^n}$, s.t. a random function in the family can be sampled using $O(n^2 + \lambda^2 + d \cdot (n + \lambda))$ uniformly random bits, and each such function can be evaluated in time $\text{poly}(n, m, d, \lambda)$.

Due to space restrictions, the proof of Lemma 3 is deferred to the full version [KRV21].

Corollary 1. *For any $n, m(n), \lambda, \ell, \epsilon' \in (0, 1]$ and $d \geq 3 \log\left(\frac{n+\lambda}{\epsilon'^2}\right) + \ell$, the family of hash functions from Lemma 3 has the following properties:*

1. *For any set $S \subseteq \{0, 1\}^m$ s.t. $2^{n-\ell} \leq |S| \leq 2^n$ it holds that*

$$\Pr_{g \leftarrow G} \left[\exists x \in \{0, 1\}^n : \left| L_{S,g}(x) - \frac{|S| \cdot 2^d}{2^n} \right| > \frac{|S| \cdot 2^d}{2^n} \cdot \epsilon' \right] \leq 2^{-\lambda}.$$

2. *For every ν s.t. $12 \cdot (n + \lambda) \leq \nu$ and set $S \subseteq \{0, 1\}^m$ s.t. $|S| \leq 2^{n-d}$ it holds that*

$$\Pr_{g \leftarrow G} \left[\exists x \in \{0, 1\}^n : L_{S,g}(x) - \frac{|S| \cdot 2^d}{2^n} > \nu \right] \leq 2^{-\lambda}.$$

4 Approximate Injectivity

In this section we analyze the *approximate injectivity* problem, introduced by Kaslasi *et al.* [KRR+20]. In particular, we consider a variant in which NO cases are (approximately) many-to-one and show that it is NISZK-complete.

We say that x' is a *sibling* of x , with respect to the circuit C , if $C(x) = C(x')$. We omit C from the notation if it is clear from the context.

Definition 9. *The problem $\text{AI}_{L,\delta}^{n,m}$ is defined as the promise problem of circuits with n input bits and m output bits, where*

$$(\text{AI}_{L,\delta}^{n,m})^Y = \left\{ \text{circuit } C : \Pr_x [|C^{-1}(C(x))| > 1] < \delta \right\},$$

and

$$(\text{AI}_{L,\delta}^{n,m})^N = \left\{ \text{circuit } C : \Pr_x [|C^{-1}(C(x))| < L] < \delta \right\}.$$

We omit m and n from the notation when they are clear from the context.

To show that $\text{AI}_{L,\delta}$ is NISZK-hard, we rely on the fact that it is known to be NISZK-hard in the special case when $L = 2$.

Lemma 4 ([KRR+20]). *Let $\delta = \delta(n) \in [0, 1]$ be a non-increasing function such that $\delta(n) > 2^{-o(n^{1/4})}$. Then, $\text{AI}_{2,\delta}$ is NISZK-hard.*

Thus, to show that $\text{AI}_{L,\delta}$ is NISZK-hard, it suffices to reduce $\text{AI}_{2,\delta}$ to $\text{AI}_{L,\delta}$.

Lemma 5. *For every parameter $\ell = \text{poly}(n)$, there exists a polynomial time Karp-reduction from $\text{AI}_{2,\delta}^{n,m}$ to $\text{AI}_{2^\ell,\ell,\delta}^{n-\ell,m-\ell}$.*

Before proving Lemma 5, we observe that Lemma 4 together with Lemma 5 immediately implies that $\text{Al}_{L,\delta}$ is NISZK-hard. Since $\text{Al}_{2,\delta}$ is a special case of $\text{Al}_{L,\delta}$, we get also that $\text{Al}_{L,\delta}$ is NISZK-complete.

Corollary 2. *Let $\delta = \delta(n) \in [0, 1]$ be a non-increasing function and $\ell = \text{poly}(n)$ such that $\delta(n) > \frac{2^{-o(n^{1/4})}}{\ell}$. Then, there exist constants $c, d \in \mathbb{N}$ such that $\text{Al}_{2^\ell,\delta}^{n^c, m^d}$ is NISZK-complete.*

Due to space restrictions, the proof of Lemma 5 is deferred to the full version [KRV21].

5 Public-Coin Batch Verification for $\text{Al}_{L,\delta}$

In this section we prove the following lemma by showing a *public-coin* HVSZK protocol for batch verification of $\text{Al}_{L,\delta}$ (as defined in Definition 9).

Lemma 6. *Let $\delta = \delta(n) \in [0, 1]$ and $k = k(n) \in \mathbb{N}$. Also, let $\lambda = \lambda(n) \in \mathbb{N}$ be a security parameter and let $\ell = \ell(n) \in \mathbb{N}$, with $\ell(n) \geq \lambda(n)$ for all n . Set $d = 7 \cdot (\log n + \log k + \lambda) + \ell$ and assume that $\delta \leq 2^{-d}$ and $d < 2\ell - 2\lambda$.*

Then, $\text{Al}_{2^\ell,\delta}^{\otimes k}$ has an HVSZK public-coin protocol with completeness error $2^{-\lambda+1}$, round-by-round soundness error $2^{-\lambda}$ and statistical zero knowledge error $k \cdot (\delta \cdot 2^{d+2} + 2^{-\lambda+6})$. The communication complexity is $O(n^2) + k \cdot \text{poly}(\log n, \log k, \lambda)$ and the verifier running time is $k \cdot \text{poly}(n, \log k, \lambda)$.

Furthermore, the protocol consists of k rounds. The length of the verifier's first message is $O(n^2 + \ell \cdot n \cdot \text{poly}(\log n, \log k, \lambda))$ and the length of all other verifier messages is $\text{poly}(\log n, \log k, \lambda)$.

The protocol establishing Lemma 6 is presented in Fig. 3. Due to space restrictions, the analysis of the protocol is deferred to the full version [KRV21].

6 From Honest to Malicious Verifier

In this section, we show how efficiently to transform an *honest-verifier* SZK protocol with round-by-round soundness, into a *malicious-verifier* SZK protocol. Our transformation builds on the prior work of Goldreich, Sahai and Vadhan [GSV98] who showed a generic transformation from honest to malicious verifiers for SZK, which unfortunately (and as discussed in more detail in the introduction) is not efficient enough for our purposes.

Theorem 5. *Suppose a problem Π has a public-coin honest-verifier SZK proof system. This protocol can be transformed into a public-coin malicious-verifier SZK proof system for Π with the following properties when given security parameter λ :*

Public-coin HVSZK Batching Protocol for $\text{Al}_{2^\ell, \delta}$.

PARAMETERS: input length n , output length m , number of instances k , security parameter λ , arity ℓ and seed length $d = 7 \cdot (\log n + \log k + \lambda) + \ell$.

INPUT: Circuits $C_1, \dots, C_k : \{0, 1\}^n \rightarrow \{0, 1\}^m$, where all circuits^a have size at most N , input length n , and output length $m \leq N$.

INGREDIENTS:

- Let $G = (G_n)_n$, where $G_n = \{g : \{0, 1\}^m \times \{0, 1\}^d \rightarrow \{0, 1\}^n\}$, be the explicit family of load-balancing functions from Lemma 3, with seed length d and accuracy $2^{-\lambda}$ with respect to security parameter $\lambda + \log k + 1$.
- Let $H = (H_n)_n$, where $H_n = \{h : \{0, 1\}^m \times \{0, 1\}^d \rightarrow \{0, 1\}^{d/2}\}$, be the explicit family of $2^{-3d/2}$ -almost pairwise independent hash function from Lemma 2.
- Let $F = (F_n)_n$, where $F_n = \{f : \{0, 1\}^n \rightarrow \{0, 1\}^{2d+\lambda+\log k}\}$, be the explicit family of $2^{-(2d+\lambda+\log k)}$ -almost pairwise independent hash functions from Lemma 2.

THE PROTOCOL:

1. V samples $g \leftarrow G_n$ and $x_{k+1} \leftarrow \{0, 1\}^n$ and sends both to P .
2. For $i = k, \dots, 1$:
 - (a) V samples $\alpha_i \leftarrow \{0, 1\}^{d/2}$, $h_i \leftarrow H_n$, and $f_i \leftarrow F_n$, and sends (α_i, h_i, f_i) to P .
 - (b) P generates the set

$$XZ_i = \left\{ (x, z) \in \{0, 1\}^n \times \{0, 1\}^d : g(C_i(x), z) = x_{i+1} \text{ and } h_i(C_i(x), z) = \alpha_i \right\}.$$

- (c) P samples $(x_i, z_i) \leftarrow XZ_i$, and sends (z_i, β_i) to V , where $\beta_i = f_i(x_i)$. In case the set XZ_i is empty, P sends an arbitrary^b $(z_i, \beta_i) \in \{0, 1\}^d \times \{0, 1\}^{2d+\lambda+\log k}$. We denote the pair received by V by (z'_i, β'_i) (allegedly equal to (z_i, β_i)).
3. P sends x_1 to V .
4. V receives $x'_1 \in \{0, 1\}^n$ (allegedly equal to x_1) and computes:
 - (a) For $i = 1, \dots, k$:
 - i. $y'_i = C_i(x'_i)$
 - ii. $x'_{i+1} = g(y'_i, z'_i)$
5. V checks that $x_{k+1} = x'_{k+1}$ and that $\forall i \in [k]$, $\beta'_i = f_i(x'_i)$ and $\alpha_i = h_i(y'_i, z'_i)$.
6. If all of V 's checks passed then she accepts. Otherwise she rejects.

^a The circuits can be trivially modified to have the same output length $m \leq N$ by padding.

^b Alternatively, we could simply have the prover abort in this case. However, it will be more convenient for our analysis that P send an arbitrary (z_i, β_i) pair rather than sending a special abort symbol.

Fig. 3. A Public-coin HVSZK Batching Protocol for $\text{Al}_{2^\ell, \delta}$

1. Suppose the original protocol has r rounds and the prover and verifier communication in its i^{th} round are s_i and ℓ_i , respectively. Then the transformed protocol has $2r$ rounds, and its total communication is $\left(\sum_{i \in [r]} s_i + O\left(\sum_{i \in [r]} \ell_i^4\right)\right)$.
2. The completeness and statistical zero-knowledge errors are at most $\text{poly}(r, \ell_{\max}) \cdot 2^{-\Omega(\lambda)}$ more (additively) than the respective errors in the original protocol, where $\ell_{\max} = \max_i \ell_i$.
3. If the original protocol has round-by-round soundness error ϵ , then this protocol has soundness error $(\epsilon r 2^\lambda + \frac{1}{2^\lambda})$.
4. The verifier runs in time polynomial in the input length, r , ℓ_{\max} , and λ , as does the prover, if given oracle access to the prover from the original protocol.

The transformation that we use to prove Theorem 5 is almost exactly the same as the one of Goldreich, Sahai and Vadhan [GSV98]. The main difference is that [GSV98] first perform an $O(r)$ -fold parallel repetition of the underlying HVSZK protocol, where r is its round complexity. This increases the communication complexity by a factor of r , which we cannot afford.

In contrast, in our analysis we avoid the use of parallel repetition and instead rely on the underlying protocol satisfying a stronger notion of soundness - namely, round-by-round soundness (Definition 7).

Due to space restrictions, we defer the proof of Theorem 5 to the full version [KRV21].

7 Public-Coin Malicious Verifier SZK Batching for NISZK

In this section we state our main theorems. The proof, which build on results established in the prior sections is deferred to the full version [KRV21]. We first state our public-coin HVSZK batch verification protocol for NISZK.

Theorem 6. *Let $\Pi \in \text{NISZK}$ and $k = k(n) \in \mathbb{N}$ such that $k(n) \leq 2^{n^{0.01}}$ and let $\lambda = \lambda(n) \in \mathbb{N}$ be a security parameter such that $\lambda(n) \leq n^{0.1}$. Then, $\Pi^{\otimes k}$ has a public-coin HVSZK protocol with completeness, zero-knowledge, and round-by-round soundness errors of $2^{-\lambda}$.*

The communication complexity is $(k + \text{poly}(n)) \cdot \text{poly}(\log n, \log k, \lambda)$ and the verifier running time is $k \cdot \text{poly}(n, \log k, \lambda)$.

Furthermore, the protocol consists of k rounds. The length of the verifier's first message is $\text{poly}(n)$ and the length of each of the verifier's other messages is $\text{poly}(\log(n, k, \lambda))$.

Combining Theorem 6 with Theorem 5, we get a *malicious-verifier* SZK batch verification protocol.

Theorem 7. *Let $\Pi \in \text{NISZK}$ and $k = k(n) \in \mathbb{N}$ such that $k(n) \leq 2^{n^{0.01}}$ and let $\lambda = \lambda(n) \in \mathbb{N}$ be a security parameter such that $\lambda(n) \leq n^{0.09}$. Then, $\Pi^{\otimes k}$ has a public-coin SZK protocol with completeness, soundness, and zero-knowledge errors of $2^{-\Omega(\lambda)}$, and communication complexity of $(k + \text{poly}(n)) \cdot \text{poly}(\log n, \log k, \lambda)$. The verifier running time is $k \cdot \text{poly}(n, \lambda, \log k)$ and the number of rounds is $O(k \cdot \lambda)$.*

Acknowledgments. We thank the anonymous Eurocrypt 2021 reviewers for useful comments.

Inbar Kaslasi and Ron Rothblum were supported in part by a Milgrom family grant, by the Israeli Science Foundation (Grants No. 1262/18 and 2137/19), and grants from the Technion Hiroshi Fujiwara cyber security research center and Israel cyber directorate.

Prashant Vasudevan was supported in part by AFOSR Award FA9550-19-1-0200, AFOSR YIP Award, NSF CNS Award 1936826, DARPA and SPAWAR under contract N66001-15-C-4065, a Hellman Award and research grants by the Okawa Foundation, Visa Inc., and Center for Long-Term Cybersecurity (CLTC, UC Berkeley). The views expressed are those of the authors and do not reflect the official policy or position of the funding agencies.

References

- [Aar12] Aaronson, S.: Impossibility of succinct quantum proofs for collision-freeness. *Quantum Inf. Comput.* **12**(1–2), 21–28 (2012)
- [ADM+99] Alon, N., Dietzfelbinger, M., Miltersen, P.B., Petrank, E., Tardos, G.: Linear hash functions. *J. ACM* **46**(5), 667–683 (1999)
- [AH91] Aiello, W., Hastad, J.: Statistical zero-knowledge Languages can be recognized in two rounds. *J. Comput. Syst. Sci.* **42**(3), 327–345 (1991)
- [APS18] Alamati, N., Peikert, C., Stephens-Davidowitz, N.: New (and Old) proof systems for lattice problems. In: Abdalla, M., Dahab, R. (eds.) PKC 2018, Part II. LNCS, vol. 10770, pp. 619–643. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-76581-5_21
- [AV19] Applebaum, B., Vasudevan, P.N.: Placing conditional disclosure of secrets in the communication complexity universe. In: Blum, A. (ed.) 10th Innovations in Theoretical Computer Science Conference, ITCS 2019, San Diego, California, USA, 10–12 January 2019. LIPIcs, vol. 124, pp. 4:1–4:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2019)
- [BBD+20] Ball, M., et al.: Cryptography from information loss. In: Vidick, T. (ed.) 11th Innovations in Theoretical Computer Science Conference, ITCS 2020, Seattle, Washington, USA, 12–14 January 2020. LIPIcs, vol. 151, pp. 81:1–81:27. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2020)
- [BCH+20] Bouland, A., Chen, L., Holden, D., Thaler, J., Vasudevan, P.N.: On the power of statistical zero knowledge. *SIAM J. Comput.* **49**(4) (2020)
- [BDRV18] Berman, I., Degwekar, A., Rothblum, R.D., Vasudevan, P.N.: From laconic zero-knowledge to public-key cryptography. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part III. LNCS, vol. 10993, pp. 674–697. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-96878-0_23
- [BGR98] Bellare, M., Garay, J.A., Rabin, T.: Fast batch verification for modular exponentiation and digital signatures. In: Nyberg, K. (ed.) EUROCRYPT 1998. LNCS, vol. 1403, pp. 236–250. Springer, Heidelberg (1998). <https://doi.org/10.1007/BFb0054130>
- [BL13] Bogdanov, A., Lee, C.H.: Limits of provable security for homomorphic encryption. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 111–128. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40041-4_7

- [CCH+19] Canetti, R., et al.: Fiat-Shamir: from practice to theory. In: Charikar, M., Cohen, E. (eds.) Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, 23–26 June 2019, pp. 1082–1090. ACM (2019)
- [CHP12] Camenisch, J., Hohenberger, S., Pedersen, M.Ø.: Batch verification of short signatures. *J. Cryptol.* **25**(4), 723–747 (2011). <https://doi.org/10.1007/s00145-011-9108-z>
- [CP92] Chaum, D., Pedersen, T.P.: Wallet databases with observers. In: Brickell, E.F. (ed.) CRYPTO 1992. LNCS, vol. 740, pp. 89–105. Springer, Heidelberg (1993). https://doi.org/10.1007/3-540-48071-4_7
- [CRSW11] Celis, L.E., Reingold, O., Segev, G., Wieder, U.: Smaller hash families and faster evaluation. In: Electronic Colloquium on Computational Complexity (ECCC), vol. 18, p. 68 (2011)
- [Dru15] Drucker, A.: New limits to classical and quantum instance compression. *SIAM J. Comput.* **44**(5), 1443–1479 (2015)
- [For87] Fortnow, L.: The complexity of perfect zero-knowledge (extended abstract). In: Aho, A.V. (ed.) Proceedings of the 19th Annual ACM Symposium on Theory of Computing, New York, New York, USA, pp. 204–209. ACM (1987)
- [For89] Fortnow, L.J.: Complexity-theoretic aspects of interactive proof systems. Ph.D. thesis, Massachusetts Institute of Technology (1989)
- [FS86] 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
- [GG00] Goldreich, O., Goldwasser, S.: On the limits of nonapproximability of lattice problems. *J. Comput. Syst. Sci.* **60**(3), 540–563 (2000)
- [GK93] Goldreich, O., Kushilevitz, E.: A perfect zero-knowledge proof system for a problem equivalent to the discrete logarithm. *J. Cryptol.* **6**(2), 97–116 (1993). <https://doi.org/10.1007/BF02620137>
- [GMR89] Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof systems. *SIAM J. Comput.* **18**(1), 186–208 (1989)
- [GMR98] Gennaro, R., Micciancio, D., Rabin, T.: An efficient non-interactive statistical zero-knowledge proof system for quasi-safe prime products. In: Gong, L., Reiter, M.K. (eds.) Proceedings of the 5th ACM Conference on Computer and Communications Security, CCS 1998, San Francisco, CA, USA, 3–5 November 1998, pp. 67–72. ACM (1998)
- [GMW87] Goldreich, O., Micali, S., Wigderson, A.: How to play any mental game or A completeness theorem for protocols with honest majority. In: Proceedings of the 19th Annual ACM Symposium on Theory of Computing, New York, New York, USA, pp. 218–229 (1987)
- [GR14] Goldwasser, S., Rothblum, G.N.: On best-possible obfuscation. *J. Cryptol.* **27**(3), 480–505 (2014)
- [GS89] Goldwasser, S., Sipser, M.: Private coins versus public coins in interactive proof systems. *Adv. Comput. Res.* **5**, 73–90 (1989)
- [GSV98] Goldreich, O., Sahai, A., Vadhan, S.: Honest-verifier statistical zero-knowledge equals general statistical zero-knowledge. In: STOC (1998)
- [GV99] Goldreich, O., Vadhan, S.P.: Comparing entropies in statistical zero knowledge with applications to the structure of SZK. In: CCC (1999)

- [HRV18] Hubáček, P., Rosen, A., Vald, M.: An efficiency-preserving transformation from honest-verifier statistical zero-knowledge to statistical zero-knowledge. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part III. LNCS, vol. 10822, pp. 66–87. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-78372-7_3
- [KMN+14] Komargodski, I., Moran, T., Naor, M., Pass, R., Rosen, A., Yogev, E.: One-way functions and (im)perfect obfuscation. In: 55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, 18–21 October 2014, pp. 374–383. IEEE Computer Society (2014)
- [KRR+20] Kaslasi, I., Rothblum, G.N., Rothblum, R.D., Sealfon, A., Vasudevan, P.N.: Batch verification for statistical zero knowledge proofs. In: Electronic Colloquium on Computational Complexity (ECCC) (2020)
- [KRV21] Kaslasi, I., Rothblum, R.D., Vasudevan, P.N.: Public-coin statistical zero-knowledge batch verification against malicious verifiers. Cryptology ePrint Archive, Report 2021/233 (2021). <https://eprint.iacr.org/2021/233>
- [KY18] Komargodski, I., Yogev, E.: On distributional collision resistant hashing. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part II. LNCS, vol. 10992, pp. 303–327. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-96881-0_11
- [LV16] Liu, T., Vaikuntanathan, V.: On basing private information retrieval on NP-hardness. In: Kushilevitz, E., Malkin, T. (eds.) TCC 2016, Part I. LNCS, vol. 9562, pp. 372–386. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49096-9_16
- [MV03] Micciancio, D., Vadhan, S.P.: Statistical zero-knowledge proofs with efficient provers: lattice problems and more. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 282–298. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-45146-4_17
- [NMVR94] Naccache, D., M’Raïhi, D., Vaudenay, S., Raphaeli, D.: Can D.S.A. be improved?—complexity trade-offs with the digital signature standard. In: De Santis, A. (ed.) EUROCRYPT 1994. LNCS, vol. 950, pp. 77–85. Springer, Heidelberg (1995). <https://doi.org/10.1007/BFb0053426>
- [NN93] Naor, J., Naor, M.: Small-bias probability spaces: efficient constructions and applications. *SIAM J. Comput.* **22**(4), 838–856 (1993)
- [NV06] Nguyen, M.-H., Vadhan, S.P.: Zero knowledge with efficient provers. In: Proceedings of the 38th Annual ACM Symposium on Theory of Computing, Seattle, WA, USA, 21–23 May 2006, pp. 287–295 (2006)
- [Oka00] Okamoto, T.: On relationships between statistical zero-knowledge proofs. *J. Comput. Syst. Sci.* **60**(1), 47–108 (2000)
- [Ost91] Ostrovsky, R.: One-way functions, hard on average problems, and statistical zero-knowledge proofs. In: Structure in Complexity Theory Conference, pp. 133–138 (1991)
- [OV08] Ong, S.J., Vadhan, S.: An equivalence between zero knowledge and commitments. In: Canetti, R. (ed.) TCC 2008. LNCS, vol. 4948, pp. 482–500. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78524-8_27
- [OW93] Ostrovsky, R., Wigderson, A.: One-way functions are essential for non-trivial zero-knowledge. In: ISTCS, pp. 3–17 (1993)
- [PS96] Pointcheval, D., Stern, J.: Security proofs for signature schemes. In: Maurer, U. (ed.) EUROCRYPT 1996. LNCS, vol. 1070, pp. 387–398. Springer, Heidelberg (1996). https://doi.org/10.1007/3-540-68339-9_33

- [PV08] Peikert, C., Vaikuntanathan, V.: Noninteractive statistical zero-knowledge proofs for lattice problems. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 536–553. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-85174-5_30
- [SV03] Sahai, A., Vadhan, S.: A complete problem for statistical zero knowledge. *J. ACM (JACM)* **50**(2), 196–249 (2003)
- [Vad99] Vadhan, S.P.: A study of statistical zero-knowledge proofs. PhD thesis, Massachusetts Institute of Technology (1999)
- [Vad12] Vadhan, S.P.: Pseudorandomness. *Found. Trends Theor. Comput. Sci.* **7**(1–3), 1–336 (2012)