

Batch Proofs are Statistically Hiding

Nir Bitansky¹, Chethan Kamath¹, Omer Paneth¹, Ron D. Rothblum², and Prashant Nalini Vasudevan³

¹Tel Aviv University,

nirbitan@tau.ac.il, ckamath@protonmail.com, omerpa@tauex.tau.ac.il

²Technion, rothblum@cs.technion.ac.il

³National University of Singapore, prashant@comp.nus.edu.sg

June 26, 2023

Abstract

Batch proofs are proof systems that convince a verifier that $x_1, \dots, x_t \in \mathcal{L}$, for some NP language \mathcal{L} , with communication that is much shorter than sending the t witnesses. In the case of *statistical soundness* (where the cheating prover is unbounded but the honest prover is efficient given the witnesses), interactive batch proofs are known for UP, the class of *unique witness* NP languages. In the case of computational soundness (a.k.a. arguments, where both honest and dishonest provers are efficient), *non-interactive* solutions are now known for all of NP, assuming standard cryptographic assumptions. We study the necessary conditions for the existence of batch proofs in these two settings. Our main results are as follows.

Statistical Soundness: the existence of a statistically-sound batch proof for \mathcal{L} implies that \mathcal{L} has a *statistically witness indistinguishable (SWI) proof*, with inverse polynomial SWI error, and a non-uniform honest prover. The implication is unconditional for obtaining honest-verifier SWI or for obtaining full-fledged SWI from public-coin protocols, whereas for private-coin protocols full-fledged SWI is obtained assuming one-way functions.

This poses a barrier for achieving batch proofs beyond UP (where witness indistinguishability is trivial). In particular, assuming that NP does not have SWI proofs, batch proofs for all of NP do not exist.

Computational Soundness: the existence of batch arguments (BARGs) for NP, together with one-way functions, implies the existence of statistical zero-knowledge (SZK) arguments for NP with roughly the same number of rounds, an inverse polynomial zero-knowledge error, and non-uniform honest prover. Thus, constant-round interactive BARGs from one-way functions would yield constant-round SZK arguments from one-way functions. This would be surprising as SZK arguments are currently only known assuming constant-round statistically-hiding commitments (which in turn are unlikely to follow from one-way functions).

Non-interactive: the existence of non-interactive BARGs for NP and one-way functions, implies non-interactive statistical zero-knowledge arguments (NISZKA) for NP, with negligible soundness error, inverse polynomial zero-knowledge error, and non-uniform honest prover. Assuming also *lossy public-key encryption*, the statistical zero-knowledge error can be made negligible and the honest prover can be made uniform.

All of our results stem from a common framework showing how to transform a batch protocol for a language \mathcal{L} into an SWI protocol for \mathcal{L} .

Contents

1	Introduction	3
1.1	Our Results	4
1.2	Additional Related Works	6
1.3	Technical Overview	7
2	Definitions	10
2.1	Proof Systems: Soundness, Privacy and Batching	11
3	Statistical Witness Indistinguishability from Batching	16
3.1	Witness Indistinguishability	17
3.2	Non-Interactive Protocols	21
3.3	Corollaries	24
4	Dual-Mode NISZKA from Non-Interactive Batch Arguments	31
4.1	Dual-Mode NISZKA from Dual-Mode Distributional NISZKA	31
4.2	Dual-Mode NISZKA Privacy Amplification	37
A	Proof of Lemma 4.12	50
A.1	Good Distributions	53
B	LPKE via Somewhere Extractable NIBARGs	53
B.1	PIR from Somewhere Extractability	53
B.2	Honest Somewhere Extractability	55

1 Introduction

Batch proofs are interactive proof-systems that enable a prover to convince a verifier that input statements x_1, \dots, x_t all belong to a language $\mathcal{L} \in \text{NP}$, with communication that is much shorter than sending the t witnesses. Batch proofs have been studied recently in two main threads: depending on whether the soundness property is required to hold against arbitrary cheating prover strategies, or only against computationally bounded ones.

The Statistical Setting. In the *statistical setting*, we require that even a computationally unbounded prover cannot convince the verifier to accept a false statement (other than with some bounded probability). On the other hand, we require that there is an *efficient honest prover* strategy (given the witnesses as an auxiliary input) for convincing the verifier of true statements. Such proofs systems are known as *doubly efficient interactive proofs* (see [Gol18] for a recent survey).

A recent sequence of works by Reingold *et al.* [RRR21, RRR18, RR20] construct doubly-efficient batch proofs for any language in the class UP (consisting of NP languages in which YES instances have a *unique* accepting witness). In particular, Rothblum and Rothblum [RR20] give such a protocol with communication $\text{poly}(m, \log(t))$, where m is the length of a single witness and poly is a polynomial that depends only on the UP language. Doubly-efficient batch proofs beyond UP remain unknown, leading to a natural question [RRR21]:

Does every language $\mathcal{L} \in \text{NP}$ have a statistically sound doubly-efficient batch proof? Do there exist other subclasses of NP (beyond UP) that have such proofs?

If we waive the restriction that the honest prover is efficient, there is a simple answer to this question. Specifically, there is a space $\text{poly}(n, m) + \log(t)$ algorithm for deciding whether $x_1, \dots, x_t \in \mathcal{L}$, where n is the instance length and m is the witness length. Thus, via the $\text{IP} = \text{PSPACE}$ theorem [LFKN92, Sha92], there is an interactive proof for this problem with communication $\text{poly}(n, m, \log(t))$. However, this protocol is entirely impractical as the honest prover runs in time $2^{\Omega(n)}$.

The Computational Setting. A natural relaxation of the statistical soundness condition is to only require *computational soundness*, which means that soundness is guaranteed only against *efficient* cheating provers. Such proof systems are commonly called *argument systems*. The seminal work of Kilian [Kil92] gives general-purpose succinct arguments for all of NP, assuming the existence of collision-resistant hash functions (CRH). In more detail, Kilian’s protocol is a four-message argument-system with communication $\text{poly}(\lambda, \log(n))$, where λ is the security parameter, for any language $\mathcal{L} \in \text{NP}$. In particular, for any $\mathcal{L} \in \text{NP}$, we can apply Killian’s result to the related NP language

$$\mathcal{L}^{\otimes t} = \{(x_1, \dots, x_t) \in (\{0, 1\}^n)^t : x_1, \dots, x_t \in \mathcal{L}\}$$

and obtain a *batch argument* (BARG) for \mathcal{L} with communication $\text{poly}(\lambda, \log(n), \log(t))$.

Kilian’s protocol relies on collision-resistant hash functions (or certain relaxations thereof [BKP18, KNY18]). However, it is unclear whether such hash functions are also necessary. This gives rise to the following basic question:

What are the minimal cryptographic assumptions needed for succinct arguments for NP? In particular, can BARGs be constructed based solely on the existence of one-way functions?

We remark that it is not clear that the existence of one-way functions is even *necessary* for general purpose succinct arguments for NP. The only result that we are aware of is by Wee [Wee05], who showed that 2-message succinct arguments imply the existence of a hard on average search problem in NP.

The Non-Interactive (Computational) Setting. As noted above, Kilian’s protocol requires 4 messages. Reducing the number of messages in succinct arguments is a major open question in the field.¹ Restricting to the case of BARGs though, we have a much better understanding due to recent breakthrough works. In particular, a sequence of works [BHK17, CJJ21, CJJ22, WW22, HJKS22, PP22, KLVW22] construct BARGs consisting of a single message, given a common reference string (equivalently, 2-message publicly verifiable arguments in the plain model), assuming specific cryptographic assumptions such as LWE or assumptions related to discrete log.

Still, so far all constructions of non-interactive BARGs rely on specific cryptographic assumptions. This raises the question of whether one can make do with a general assumption as in Kilian’s protocol. More ambitiously:

Can non-interactive BARGs be constructed from collision-resistant hash functions?

1.1 Our Results

In this work we study, and give partial answers, to all of the above questions. Our key idea is a new transformation that compiles a batch protocol² Π , for verifying that $x_1, \dots, x_t \in \mathcal{L}$, into a protocol Π' , for a single instance, which has a hiding property. Here and below, when we say that \mathcal{L} has a batch protocol, we mean that the communication for proving that $x_1, \dots, x_t \in \mathcal{L}$ is $t^{1-\epsilon} \cdot \text{poly}(m)$, for some $\epsilon > 0$, where m is the length of a single witness and poly is some polynomial which may depend on \mathcal{L} but does not depend on t .

More specifically, we transform a batch protocol Π into a protocol Π' for a single instance satisfying a form of *statistical witness indistinguishability* (SWI). Recall that a protocol for an NP relation \mathcal{R} is ϵ -SWI, if for every input x and witnesses $w_1, w_2 \in \{w : \mathcal{R}(x, w) = 1\}$, the view of the verifier when the prover uses w_1 is ϵ -close, in statistical distance, to its view in an interaction in which the prover uses w_2 . We say that the protocol is *honest verifier* SWI if the SWI property only holds in an honest execution of the protocol (but the default notion applies to *malicious verifiers*).

Our main step transforms a batch protocol Π into an *honest-verifier* SWI protocol Π' , where the SWI-error ϵ can be any inverse polynomial. The transformation also preserves the soundness of the original protocol. In other words, if Π is computationally (resp., statistically) sound then the resulting protocol Π' is computationally (resp., statistically) sound. If Π has r -rounds then Π' has $r+1$ rounds. However, the efficient honest prover strategy of Π' is non-uniform, where the non-uniform advice depends on the specification of the protocol Π .

We use this basic step in the different settings described above to reduce batch protocols into protocols satisfying hiding properties, as described next.

The Statistical Setting. Our first application of the above framework is in the statistical setting. In this setting we obtain SWI against malicious verifiers, in which the SWI error is inverse polynomial. In case we start off with a public-coin BARG the result is unconditional. Otherwise we need to assume a one-way function (or settle for honest-verifier SWI).³

Theorem 1 (Informally Stated, see Theorem 3.1 and Corollaries 3.14, 3.15). *Suppose that $\mathcal{L} \in \text{NP}$ has a statistically sound r -round public-coin batch proof. Then, for any polynomial p , the language \mathcal{L} has an $O(r)$ -round SWI proof with $\frac{1}{p}$ -SWI error and a non-uniform honest prover.*

Furthermore, for general (i.e., private-coin) statistically sound batch proofs we achieve the weaker conclusion of honest-verifier SWI, or, assuming the existence of a one-way function, malicious verifier SWI.

¹Here and throughout this work we focus either on the plain model in the interactive setting, or common reference string (CRS) model in the non-interactive setting. In contrast, in the *random oracle model* general purpose succinct arguments (and in particular non-interactive BARGs) are known to exist unconditionally [Mic94].

²We use the terminology of “protocol” where we want to be intentionally vague as to whether soundness is computational or statistical.

³Note that the Goldwasser and Sipser [GS89] transformation from private-coin to public-coin protocols is inapplicable, since it results in an inefficient honest prover (see also [AR21]).

It is worth pointing out that Theorem 1 is also applicable to languages in UP (for which batch proofs are known), but there the conclusion is meaningless since UP has a trivial SWI proof - just send the witness! In contrast though, for general NP languages, the existence of an SWI proof seems extremely surprising. In particular, it is known that NP does not have proofs satisfying the stronger property of *statistical zero-knowledge* (SZK) (assuming the polynomial hierarchy does not collapse [For89, AH91]).⁴ As it seems that the notion of SWI is closely related to that of SZK (modulo the trivial cases arising from unique witnesses) it seems reasonable to expect that NP does not have such proofs. Thus, we derive the following immediate corollary:

Corollary 2 (Informally Stated). *Assume that there exists some $\mathcal{L} \in \text{NP}$ that does not have an SWI proof as in Theorem 1. Then NP does not have statistically sound batch proofs.*

We do not take for granted the fact that NP does not have SWI proofs, and we find this to be an intriguing open question. Indeed, while we have a very deep understanding of the structure of SZK (see [Vad99]), the structure of the class of languages having SWI proofs has, to the best of our knowledge, not been explored. Theorem 1 provides concrete motivation for a similar study of the class SWI, which we leave to future work.

The Computational Setting. We also apply our basic framework in the computational setting. Here though, assuming that one-way functions exist, we are able to derive the stronger hiding property of *statistical zero-knowledge*.

Theorem 3 (Informally Stated, See Theorem 3.1 and Corollary 3.15). *Assume the existence of a one-way function. Suppose that every $\mathcal{L} \in \text{NP}$ has an r -round BARG. Then, for every polynomial p , every $\mathcal{L} \in \text{NP}$ has an $O(r)$ -round statistical zero-knowledge argument-system (SZKA) with $\frac{1}{p}$ -zero-knowledge error and a non-uniform honest prover.*

Recall that constant-round SZKA for NP are only known to exist assuming constant-round statistically-hiding commitments, and the latter seem stronger than one-way functions (and there is a blackbox separation [HHR15]). Thus, Theorem 3 shows that the existence of constant-round BARGs for NP suffices to “lift” one-way functions to a primitive which is only known based on collision-resistant hash functions (or multi-collision resistant hash functions [BKP18, BDRV18]).

We remark that a related positive result was obtained recently by Amit and Rothblum [AR23], who constructed constant-round succinct arguments for deterministic languages (specifically for the class NC) from one-way functions. Thus, a negative interpretation of Theorem 3 is that extending the [AR23] result from succinct arguments for deterministic languages to BARGs for NP seems unlikely as it would have unexpected implications. Alternatively, a positive perspective is that Theorem 3 presents a concrete direction for constructing constant-round SZKA for NP from one-way functions.

The Non-Interactive (Computational) Setting. We apply the basic framework for the third time in the context of *non-interactive* BARGs. Here we face a difficulty, in that our basic framework increases the round complexity of the protocol by one round. We are able to overcome this challenge by relying on BARGs satisfying a weak form of adaptive soundness called *somewhere soundness*, a relaxation of *somewhere extractability* [CJJ22], which is achieved by recent BARG constructions. We obtain the following result:

Theorem 4 (Informally Stated, See Theorem 3.1 and Corollary 3.17). *Assume the existence of one-way functions and that NP has somewhere-sound non-interactive BARGs. Then, for any polynomial p , NP has non-interactive statistical zero-knowledge arguments (NISZKA), with a negligible non-adaptive soundness error, $\frac{1}{p}$ -zero-knowledge error, and a non-uniform honest prover.*

Like non-interactive BARGs, NISZKA are currently only known to exist based on specific cryptographic assumptions (or in the random oracle model). Theorem 4 shows that a construction from a “relatively weak”

⁴Recall that *statistical zero-knowledge* (SZK) requires that for every efficient verifier strategy there is an efficient simulator that generates a view that is statistically close to that in the actual interaction (for instances in the language). SWI can be thought of as a relaxation of SZK in which the simulator can be unbounded.

assumption, such as collision-resistant hash functions, would yield a similar result for NISZKA, which would constitute major progress in the field of zero-knowledge.

We further prove that assuming *lossy public key encryption*, which exist from a variety of assumptions (c.f. [PW08]), we can improve the conclusion in Theorem 4 in two ways: (i) the honest prover is a uniform algorithm and (ii) we can reduce the statistical zero knowledge error to negligible.

Theorem 5 (Informally Stated, See Theorem 4.6). *Assume the existence of lossy public-key encryption and that NP has somewhere-sound non-interactive BARGs. Then NP has non-interactive statistical zero-knowledge arguments (NISZKA), with a negligible non-adaptive soundness error, a negligible zero-knowledge error, and a uniform honest prover.*

We in fact prove a slightly more general result. We show a dual-mode NISZKA where the common-reference string has two indistinguishable modes: the first statistically zero knowledge and non-adaptively sound, and the second (adaptively) computationally zero knowledge and adaptively sound.

We also observe that lossy public-key encryption follows from a variant of somewhere extractable BARGs, which guarantees that it is possible to extract the specific witness that was used in some predefined index in an honest proof. This is in contrast to the standard notion of somewhere extractability guaranteeing that *some* witness can be extracted (even from maliciously generated accepting proofs). As a result of independent interest, we also show that the standard notion of somewhere extractable BARGs imply *private information retrieval* and thus also *statistically sender-private oblivious transfer* and lossy public-key encryption. However the lossy public-key encryption obtained has (negligible) decryption errors (which is not sufficient for our amplification theorem). See further details in Appendix B.

Remark 1 (Hiding for Batch Protocols). *All of the results listed above start with a batch protocol for a language \mathcal{L} and derive a protocol with hiding properties (i.e., either SWI or SZK) for a single instance of \mathcal{L} . We note that all of the results can be used to obtain similar hiding properties also for a batch protocol for \mathcal{L} via the following simple observation: rather than applying the basic result to \mathcal{L} , we can apply it to $\mathcal{L}^{\otimes t'}$ for any $t' \ll t$.*

Remark 2 (On the Possibility of Weak Batching). *All of our results assume a batch protocol for t instances, with communication $t^{1-\epsilon}$.⁵ Thus, our results are inapplicable to very weakly compressing batch protocols that have slightly non-trivial communication such as say, $t \cdot \sqrt{m} + \text{poly}(m)$, where m is the witness length. Such weak batch protocols can nevertheless be quite useful (see [RRR21]) and we leave the study of this setting as an interesting open problem.*

1.2 Additional Related Works

The study of the minimal necessary communication in statistically sound interactive proofs, focusing on the prover to verifier communication, was initiated in [GH98, GVW02]. In particular, Goldreich *et al.* [GVW02] transform interactive proofs with a *single bit* of communication to be SZK. We emphasize that the results in [GH98, GVW02] are inapplicable in the setting of batch proofs. For example, the main result in [GH98] says that proofs with short communication can be emulated in time that is exponential in the communication, but this merely indicates that the communication in batch proofs for NP needs to be $\Omega(m + \log t)$, where m is the witness length.

Kaslasi *et al.* [KRR⁺20, KRV21] consider batch verification of protocols that are *a priori* statistical zero-knowledge, while retaining the zero-knowledge property. The constructions of [KRR⁺20, KRV21] are *not* doubly-efficient and so our results are inapplicable in their context.

Batch verification is also related to the problem of *AND instance compression* [HN10, FS08]. In AND instance compression, the goal is, given formulas ϕ_1, \dots, ϕ_k , to generate in polynomial time a new formula ϕ that is satisfiable if and only if ϕ_1, \dots, ϕ_k are all satisfiable, and so that the length of ϕ is less than k .

⁵We remark that Kalai *et al.* [KLVW22] show how to amplify weak non-interactive BARGs into BARGs with very good compression but they assume the existence of rate-1 OT, whereas we are seeking transformations that rely only on the existence of the weak BARG.

Batch verification considers the dual problem of compressing the *witnesses*. We note that strong infeasibility results for AND instance compression were shown by Drucker [Dru15]. Despite the differences, a main technical lemma used by Drucker (and a subsequent simplification by Dell [Del16]) is a key inspiration for our analysis.⁶ We note that this lemma has previously been used for identifying sufficient conditions for obtaining cryptographic primitives from average-case hardness [BBD⁺20].

The minimax theorem has found several applications in cryptography: see [VZ13] for the references. The work of [CLP15] also establishes (among other results) a result of the form “succinctness implies hiding” using the minimax theorem. To be specific, they showed that for a proof system with a *laconic prover*, i.e. where the communication from the prover is $O(\log(n))$ bits, [GVW02] implies *distributional* computational zero knowledge. It is also worth pointing out that the usage of the minimax theorem there is for a different purpose compared to us: it is used there to switch the order of quantifiers of the simulator and distinguisher to obtain (normal) zero knowledge from weak zero knowledge.

Lastly, we mention a recent result of Kitagawa *et al.* [KMY20], who show how to transform any SNARG (a much stronger notion than non-interactive BARG, and not known based on standard assumptions) into a NIZK, assuming one-way functions. The resulting NIZK argument is only computational zero-knowledge. In contrast, Corollary 3.17 assumes the weaker notion of non-interactive BARG and constructs the stronger notion of statistical zero-knowledge. Still, the results are incomparable as we rely also on lossy public-key encryption, or alternatively, only get an inverse polynomial zero knowledge error and non-uniform prover.

Comparison with [CW23]. In a concurrent and independent work, Champion and Wu [CW23] constructed *computational* NIZK arguments assuming non-interactive BARGs, lossy public key encryption and sub-exponentially-secure *local* pseudorandom generators. Since they build on the result of Kitagawa *et al.* [KMY20], their approach is quite different from ours (although they also rely on lossy public-key encryption).

In the first posting of our paper, which was concurrent to [CW23] we obtained a weaker NIZK result which was technically incomparable to that [CW23] as the prover was *non-uniform*. However, subsequently to the first posting and that of [CW23], we were able to obtain Theorem 5 which subsumes the [CW23] result, since we construct *statistical* NIZK arguments (in fact, dual mode ones) assuming only BARGs and lossy public-key encryption (i.e., a stronger conclusion based on a subset of the assumptions used in [CW23]).

1.3 Technical Overview

Let \mathcal{R} be an NP relation, and let

$$\mathcal{R}^{\otimes t} = \left\{ \left((x_1, \dots, x_t), (w_1, \dots, w_t) \right) : |x_1| = \dots = |x_t| \text{ and } \forall i \in [t], (x_i, w_i) \in \mathcal{R} \right\}$$

be the corresponding batch relation. We start by assuming a batch protocol for $\mathcal{R}^{\otimes t}$ (without specifying yet whether soundness is statistical or computational). For simplicity, let us assume that $\mathcal{R}^{\otimes t}$ has an entirely non-interactive protocol - that is, a single message sent from the prover to the verifier. We view the prover message in this case as a “compression function” f that takes as input $(x_1, \dots, x_t, w_1, \dots, w_t)$ and outputs a short proof string π that convinces the verifier. Note that f is an efficiently computable function, since we assume the honest prover strategy is efficient (given also the witnesses).

Since f outputs a short string, of length less than t , its output cannot contain all of the witnesses. Thus, intuitively at least, a large portion of the information about the witnesses must be lost. This leads us to the following natural idea for a protocol, for a single⁷ instance of \mathcal{R} , that has hiding properties.

P(x, w) : (where x is an input and w is a corresponding witness)

1. Choose a random index $i^* \in [t]$.

⁶We note that a closely related lemma was established earlier in the context of constructing an oblivious transfer protocol from any private information retrieval scheme [DMO00].

⁷By this we mean for an instance corresponding to $\mathcal{L}(\mathcal{R})$, the language corresponding to the relation \mathcal{R} .

2. Select input/witness pairs $(x_i, w_i) \in \mathcal{R}$ for all $i \in [t] \setminus \{i^*\}$, in some yet-to-be-specified way.
3. Generate $\pi = f(x_1, \dots, x_t, w_1, \dots, w_t)$, where we fix $x_{i^*} = x$ and $w_{i^*} = w$.
4. Send $(x_1, \dots, x_t, i^*, \pi)$ to the verifier.

The verifier V accepts if (1) $x_{i^*} = x$ and (2) the batch verifier accepts the input (x_1, \dots, x_t) with the proof π . Completeness and soundness of this protocol follow immediately from the completeness and soundness of the batch protocol (notice that for soundness, it suffices that x is a NO instance for \mathcal{R} to make (x_1, \dots, x_t) a NO instance for $\mathcal{R}^{\otimes t}$).

The key question is how to choose the instance-witness pairs in Step 2 in such a way that π hides w_{i^*} . This choice is crucial. To see this, consider a contrived compression function whose goal is to be maximally non-hiding for some specific input x^* . For example, the compression function, in addition to outputting a convincing proof, might check if one of the t inputs is equal to x^* . If so, it also outputs the corresponding witness as part of the proof. Notice that this strategy is still highly compressing. While this is clearly a contrived strategy, since we seek to give a general result, that compiles *any* batch proof, we have to consider such strategies as well.

The above contrived strategy is a major concern for SWI as there exists a specific input, namely x^* , for which the prover always reveals the witness. A natural approach to circumvent this attack is to consider a distributional notion of SWI. That is, consider some efficiently sampleable distribution \mathcal{D} supported on triples (x, w_0, w_1) , where $(x, w_0), (x, w_1) \in \mathcal{R}$. Suppose we only want SWI to hold for random instance/witness pairs sampled from \mathcal{D} . In such a case, P can choose each (x_i, w_i) from \mathcal{D} independently. Now, for inputs (x, w_0, w_1) that are also generated from \mathcal{D} , by symmetry, the function f will be unable to discover whether w_0 or w_1 was guessed (other than with inverse polynomial probability). Intuitively, and this can be formalized, this leads to a distributional-SWI protocol (with an SWI error that decreases polynomially with t).

While the distributional approach described above works, it is weaker than what we aim to achieve in two ways. First, it is restricted to NP languages that have a solved instance generator (recall that if the language is also hard with respect to this distribution then the sampler is a one-way function). Second, the SWI property is distributional - it only holds with respect to instance-witness pairs sampled from \mathcal{D} (rather than the usual worst-case guarantee).

At this point we face a problem. If we aim to get a worst-case SWI guarantee, the contrived compression function f that targets some specific x^* seems like a non-starter. Indeed, using f as a blackbox, it is hopeless to try to discover x^* . Still, if we happened to know that the compression function is precisely the contrived one described above, we could fix the same x^* as part of prover P and then use x^* (with corresponding random witnesses that are also hardwired) in *all* of the coordinates of f . Doing so would hide the specific witness that P uses in the i -th coordinate. But what about a general compression function f ? Can we somehow fix specific instance/witness pairs that are specifically good for fooling f ? Somewhat surprisingly the answer turns out to be yes.

How to find instance-witness pairs. Our main technical result shows that for every compression function f there exists a polynomial-size multiset $S \subseteq \mathcal{R}^{\otimes t}$ (i.e. a polynomial number of instance-witness t -tuples), so that if the tuple $((x_1, w_1), \dots, (x_t, w_t))$ used in the above protocol is sampled uniformly from S , then the resulting protocol is SWI (with error that depends on how compressing f is).

Central to our approach is a lemma of Dell [Dell16] (building on work by Drucker [Dru15] and related to a result of [DMO00]) about information lost by compressing functions. Consider a function $g : \{0, 1\}^t \rightarrow \{0, 1\}^{\rho t}$ for some $\rho < 1$. Intuitively, as the function is compressing, it must be losing information about some of its input bits. Dell formalized this by showing that the output distribution of g when its input bits are chosen uniformly at random is not affected much by arbitrarily fixing the bit at a randomly chosen location. Let B be the uniform distribution over $\{0, 1\}^t$, and denote by $B|_{j \leftarrow b}$ the variable corresponding to sampling B and setting the j^{th} co-ordinate to b . Dell showed that in terms of statistical distance:

$$(j, g(B|_{j \leftarrow 0})) \approx_{\sqrt{\rho}} (j, g(B|_{j \leftarrow 1})).$$

Suppose g is a function parameterized by triples (x_i, w_i^0, w_i^1) , where $(x_i, w_i^0), (x_i, w_i^1) \in \mathcal{R}$, and uses its input bits b_i to select witness $w_i^{b_i}$, and outputs f computed with these instance-witness pairs $(x_i, w_i^{b_i})$. The above lemma would then say that picking a random $j \in [t]$ and fixing the witness used for x_j to be either of w_j^0 or w_j^1 would not make much of a difference to the output distribution of f . Denoting (x_1, \dots, x_t) by \mathbf{x} and (w_1, \dots, w_t) by \mathbf{w} , with $j \leftarrow [t]$ and each w_i sampled uniformly from $\{w_i^0, w_i^1\}$, this implies that:

$$\left(j, \mathbf{x}, f(\mathbf{x}, \mathbf{w}|_{j \leftarrow w_j^0})\right) \approx \left(j, \mathbf{x}, f(\mathbf{x}, \mathbf{w}|_{j \leftarrow w_j^1})\right)$$

This is already reminiscent of witness-indistinguishability, though the property here only holds for a randomly chosen instance among a set of t instances. We can, in fact, use this to get the distributional version of SWI discussed above. Consider any distribution D over (x, w_0, w_1) such that $(x, w_0), (x, w_1) \in \mathcal{R}$. Now, with (x, w_0, w_1) and all the (x_i, w_i^0, w_i^1) sampled from \mathcal{D} , we have:

$$(j, \mathbf{x}|_{j \leftarrow x}, f(\mathbf{x}|_{j \leftarrow x}, \mathbf{w}|_{j \leftarrow w_0})) \approx (j, \mathbf{x}|_{j \leftarrow x}, f(\mathbf{x}|_{j \leftarrow x}, \mathbf{w}|_{j \leftarrow w_1}))$$

Note that in the protocol above, when the prover inserts the given (x, w) at location j and uses instances x_i and witnesses w_i in the remaining locations, the view of the verifier is precisely $(j, \mathbf{x}|_{j \leftarrow x}, f(\mathbf{x}|_{j \leftarrow x}, \mathbf{w}|_{j \leftarrow w}))$. So the above implies that the *expected SWI error* for the protocol when everything is sampled as specified is small.

In other words, for every distribution D over (x, w_0, w_1) , there is a distribution over $((x_i, w_i^0, w_i^1))$ such that with samples from these, the expected SWI error in our protocol is small. We can view this process as a 2-player zero-sum game: the row player chooses (x, w_0, w_1) and the column player chooses a distribution D over all such tuples. The payoff is the expected SWI error in our protocol. The above argument shows that for every strategy D for the column player there is a mixed strategy for the row player (specifically, the strategy D), for which we can bound the expected payoff. The minimax theorem now implies that there is a *single* distribution D' over tuples $((x_i, w_i^0, w_i^1))$ such that *for every* (x, w_0, w_1) , if the prover uses a sample from D' to populate the other inputs to f , the SWI error is small. Using a sparse minimax theorem [LY94] now implies the existence of a polynomial-sized multiset of $((x_i, w_i^0, w_i^1))$'s such that sampling from this leads to almost the same SWI error. This implies the existence of the set we want, which we hard-code into the prover's algorithm as a non-uniform advice.⁸

Remark 3. *The $O(\sqrt{\rho})$ error in our analysis is tight for some functions (e.g., if g is the majority function). However, "natural" compression functions might not exhibit such a behavior and could potentially give rise to a negligible SWI error.*

Handling Multi-round Protocols. To handle multi-round protocols we follow the same basic strategy, running the underlying batch protocol using tailor-made instance/witness pairs. While we are unable to show that this approach satisfies *malicious-verifier* SWI, we manage to show that it is *honest-verifier* SWI. We do so by first extending the above analysis to 2-message protocols (i.e. a verifier message followed by a prover message). To handle protocols with more messages, we observe that when analyzing *honest-verifier* SWI, we can imagine that the verifier sends to the prover all of its randomness in advance and reduce back to the 2-message case.

Augmenting the Basic Result in The Interactive Setting. At this point we have a transformation from any batch protocol into an honest-verifier SWI protocol with inverse polynomial SWI error. We can improve this state of affairs in the different settings as follows:

1. In the case of statistical soundness, if the batch proof is *public-coin*, we can rely on an information-theoretic coin-flipping protocol due to Goldreich *et al.* [GSV98] which leads to malicious verifier SWI.⁹

⁸It seems tempting to try to use a uniform minimax theorem, as in [VZ13], to obtain a uniform honest prover. A key bottleneck however is that our payoff function does not seem to be efficiently computable. See also Remark 7.

⁹Note that we cannot use the honest-to-malicious transformation of Hubáček *et al.* [HRV18] (which works also in the private-coin setting) because that result relies on the connection of SZK to instance dependent commitments. Thus, it is not clear how to apply their result in the setting of SWI.

For the case of private-coin protocols, following an approach of [BMO90, OY93, Oka96], we show that assuming the existence of a one-way function, we can transform any honest-verifier SWI protocol to be malicious verifier SWI. We emphasize that despite the usage of a one-way function, both soundness and hiding properties are *statistical*.

2. In the case of computational soundness, assuming the existence of a one-way function, we can rely on the celebrated “FLS trick” of Feige *et al.* [FLS90] to bootstrap the honest-verifier SWI argument to an honest-verifier SZK argument.¹⁰ Then, using the [GMW86] compiler from honest-verifier to malicious verifier we obtain a full-fledged malicious verifier zero-knowledge argument (using the [FS90] constant-round private-coin argument-system as the underlying zero-knowledge proof).

The Non-Interactive Setting. Recall that in this setting the prover sends a single message, that may depend on a previously chosen common random string (CRS). One challenge that we have to deal with is that in the basic protocol, the prover needs to send its choice of (x_1, \dots, x_t) before starting the batch protocol (i.e., before the CRS is chosen), whereas in the protocol we construct this happens after the CRS is chosen. As mentioned earlier, we handle this reversing by relying on somewhere soundness, a weak form of adaptive soundness for BARGs. Given the resulting non-interactive SWI argument, we can use the FLS trick in a similar way to obtain a NISZKA protocol with inverse polynomial error. Assuming lossy public-key encryption, we then enhance the above in two ways: (a) we obtain a uniform prover, (b) we reduce the statistical zero-knowledge error. For both, lossy public-key encryption is used as a dual-mode commitment — for computationally indistinguishable public keys we get either statistical hiding or statistical binding.

To obtain a uniform prover, our key observation is that we can upgrade distributional NISZKA (with inverse polynomial error), to worst-case NISZK (with inverse polynomial error), while retaining a uniform honest prover, as follows.¹¹ The transformation is inspired by the NIZK of [GOS06] — we use a lossy encryption scheme to commit to all the wires of the NP verification circuit (in lossy mode), and prove local consistency of each gate using distributional NISZKA. Specifically, for each gate, we consider the distribution of *random* encryptions of *random* wire values satisfying the gate. While the prover actually uses a specific (worst-case) wire assignment for the gate, this only skews the distribution by a constant factor, as the total number of assignments is constant.

As our last step, we show a statistical zero-knowledge amplification theorem similar to the one by [GJS19] for computational zero knowledge. Like their transformation, we construct a combiner based on MPC-in-the-head (in our case, an information-theoretic one, such as [BGW88]). We show based on a coupling proof, similar to the one in [LM20], that the combiner is in fact also an amplifier.

2 Definitions

We rely on the standard computational concepts and notation:

- A binary relation $\mathcal{R} \subseteq \{0, 1\}^* \times \{0, 1\}^*$ is said to be *polynomially balanced* if there is a polynomial p such that for any strings $x, y \in \{0, 1\}^*$, if $(x, y) \in \mathcal{R}$ then $|y| \leq p(|x|)$. For a (polynomially-balanced) relation \mathcal{R} and $\lambda \in \mathbb{N}$, \mathcal{R}_λ denotes $\mathcal{R} \cap (\{0, 1\}^\lambda \times \{0, 1\}^*)$. For a (polynomially-balanced) relation \mathcal{R} , we use $\mathcal{L}(\mathcal{R})$ to denote the language defined by \mathcal{R} , i.e., $\{x \in \{0, 1\}^* : \exists w \in \{0, 1\}^* \text{ s.t. } (x, w) \in \mathcal{R}\}$. $\mathcal{L}(\mathcal{R}_\lambda) \subseteq \{0, 1\}^\lambda$ is defined similarly, but with respect to \mathcal{R}_λ .
- We say that a function $f : \mathbb{N} \rightarrow \mathbb{R}$ is *negligible* if for all constants $c > 0$, there exists $N \in \mathbb{N}$ such that for all $n > N$, $f(n) < n^{-c}$. We sometimes denote negligible functions by *negl.* We say that a function $f : \mathbb{N} \rightarrow \mathbb{R}$ is *overwhelming* if $1 - f$ is negligible.

¹⁰In a nutshell, the verifier sends to the prover $z = G(s)$, where G is a PRG and s is a random seed, and the prover then proves that either $x \in \mathcal{L}$ or z is in the image of the PRG. Computational soundness can be argued by switching to a truly random z , and SWI by having the simulator use s as the witness.

¹¹This is in contrast to the transformation described above which uses the minmax theorem and obtains a non-uniform honest prover.

- A PPT algorithm is a probabilistic polynomial-time algorithm. A family of circuits $A = (A_\lambda)_{\lambda \in \mathbb{N}}$ is *polynomial-sized* if there exists a polynomial p such that for all $\lambda \in \mathbb{N}$, $|A_\lambda| \leq p(\lambda)$. We follow the common practice of modelling any efficient adversary as a family of polynomial-size circuits $A = (A_\lambda)_{\lambda \in \mathbb{N}}$. We also say that such an A runs in *non-uniform polynomial time*.
- For a distribution X over a set Ω and $x \in \Omega$, we use $x \leftarrow X$ to denote the result of sampling according to X . For a random variable X over Ω and $x \in \Omega$, we use $X(x)$ to denote the probability that the value of the random variable is x . We denote *statistical distance* by SD . For two random variables X, Y and $\varepsilon \in [0, 1]$, we write $X \approx_\varepsilon Y$ to denote the fact that $\text{SD}(X, Y) \leq \varepsilon$ and say that X is ε -*statistically indistinguishable* from Y . For two ensembles $\mathcal{X} = (X_\lambda)_{\lambda \in \mathbb{N}}$ and $\mathcal{Y} = (Y_\lambda)_{\lambda \in \mathbb{N}}$ and function ε , we write $\mathcal{X} \approx_\varepsilon \mathcal{Y}$ if for all large enough λ , $X_\lambda \approx_{\varepsilon(\lambda)} Y_\lambda$. For two ensembles $\mathcal{X} = (X_\lambda)_{\lambda \in \mathbb{N}}$ and $\mathcal{Y} = (Y_\lambda)_{\lambda \in \mathbb{N}}$, we write $\mathcal{X} \approx_c \mathcal{Y}$ to denote the fact that \mathcal{X} is *computationally indistinguishable* from \mathcal{Y} , i.e., for every polynomial-sized family of distinguishers $D = (D_\lambda)_{\lambda \in \mathbb{N}}$, there exists a negligible function μ such that for all $\lambda \in \mathbb{N}$

$$\left| \Pr_{x \leftarrow X_\lambda} [D_\lambda(x) = 1] - \Pr_{y \leftarrow Y_\lambda} [D_\lambda(y) = 1] \right| \leq \mu(\lambda).$$

- We say that an ensemble of distributions $\mathcal{D} = (D_\lambda)_{\lambda \in \mathbb{N}}$ is *efficiently sampleable* if there exists a polynomial-sized family of circuits $S = (S_\lambda)_{\lambda \in \mathbb{N}}$ such that the distribution of the outputs of S_λ , on input a uniformly random string, is D_λ . We also use the stronger *uniform* notion of efficient sampleability, where there exists a PPT algorithm S such that the distribution of outputs of $S(1^\lambda)$ is D_λ .
- We denote by $\langle P \rightleftharpoons V \rangle$ an interactive protocol between two parties P and V . For common input x and auxiliary input w for P , by $\langle P(w) \rightleftharpoons V \rangle(x)$ we denote the *view* of V in the protocol, including all received messages and random coins (if V is randomized). We abuse notation and write $\langle P(w) \rightleftharpoons V \rangle(x) = 1$ to denote the fact that V accepts. A similar convention applies also to non-interactive protocols, which are denoted by $\langle P \rightarrow V \rangle$.

2.1 Proof Systems: Soundness, Privacy and Batching

Interactive Protocols. In this paper, we are mainly interested in *doubly-efficient* interactive protocols [Gol18] as defined below.

Definition 2.1 (Doubly-Efficient Interactive Protocols). *Let $\langle P \rightleftharpoons V \rangle$ be an interactive protocol for an NP relation \mathcal{R} , i.e., the common input is a statement $x \in \mathcal{R}$ and the auxiliary input for the (honest) prover P is a witness w for x . The protocol $\langle P \rightleftharpoons V \rangle$ is doubly efficient if both P and V are PPT algorithms.*

Remark 4. *We will also consider doubly-efficient interactive protocols where the prover is non-uniform (i.e., a polynomial-sized family of circuits), in which case we will always explicitly point this out. In the rest of the paper, we drop “doubly efficient”, and by interactive protocols we always refer to doubly-efficient interactive protocols.*

We next define the relevant notions of completeness, soundness and privacy. In the following definitions $\langle P \rightleftharpoons V \rangle$ is a (doubly-efficient) protocol for an NP relation \mathcal{R} .

Definition 2.2 (Completeness). *$\langle P \rightleftharpoons V \rangle$ is complete with (completeness) error ε if for all large-enough $\lambda \in \mathbb{N}$ and every $(x, w) \in \mathcal{R}_\lambda$*

$$\Pr [\langle P(w) \rightleftharpoons V \rangle(x) = 1] \geq 1 - \varepsilon(\lambda) .$$

Definition 2.3 (Statistical Soundness). *The protocol $\langle P \rightleftharpoons V \rangle$ is statistically sound with (soundness) error ε if for every (unbounded) prover P^* , all large enough $\lambda \in \mathbb{N}$ and every $x \in \{0, 1\}^\lambda \setminus \mathcal{L}(\mathcal{R})$,*

$$\Pr [\langle P^* \rightleftharpoons V \rangle(x) = 1] \leq \varepsilon(\lambda) .$$

A statistically sound protocol is also called a proof.

Definition 2.4 (Computational Soundness). $\langle P \rightleftharpoons V \rangle$ is computationally sound if for every polynomial-size circuit family of provers $P^* = (P_\lambda^*)_{\lambda \in \mathbb{N}}$, there exists a negligible function μ , such that for every $\lambda \in \mathbb{N}$ and every $x \in \{0, 1\}^\lambda \setminus \mathcal{L}(\mathcal{R})$,

$$\Pr [\langle P_\lambda^* \rightleftharpoons V \rangle(x) = 1] \leq \mu(\lambda) .$$

A computationally-sound protocol is also called an argument.

Definition 2.5 (Statistical Witness Indistinguishability). $\langle P \rightleftharpoons V \rangle$ is statistically witness-indistinguishable with error ε if for every polynomial-size circuit family $V^* = (V_\lambda^*)_{\lambda \in \mathbb{N}}$,

$$\{\langle P(w_0) \rightleftharpoons V_\lambda^* \rangle(x)\}_{\substack{(x, w_0, w_1) \in \mathcal{R} \\ |x| = \lambda}} \approx_\varepsilon \{\langle P(w_1) \rightleftharpoons V_\lambda^* \rangle(x)\}_{\substack{(x, w_0, w_1) \in \mathcal{R} \\ |x| = \lambda}} ,$$

where $(x, w_0, w_1) \in \mathcal{R}$ is an abuse of notation to be interpreted as $(x, w_0), (x, w_1) \in \mathcal{R}$. If the above indistinguishability is only guaranteed for the honest verifier V , then $\langle P \rightleftharpoons V \rangle$ is honest-verifier statistically witness-indistinguishable.

Definition 2.6 (Statistical Zero Knowledge). $\langle P \rightleftharpoons V \rangle$ is statistically zero-knowledge with error ε if there exists an expected PPT simulator S such that for every polynomial-size circuit family $V^* = (V_\lambda^*)_{\lambda \in \mathbb{N}}$,

$$\{\langle P(w) \rightleftharpoons V_\lambda^* \rangle(x)\}_{\substack{(x, w) \in \mathcal{R} \\ |x| = \lambda}} \approx_\varepsilon \{S(x, V_\lambda^*)\}_{\substack{(x, w) \in \mathcal{R} \\ |x| = \lambda}} .$$

The protocol is honest-verifier statistical zero-knowledge if the above is only guaranteed for the honest verifier V .

Definition 2.7 (Computational Zero Knowledge). $\langle P \rightleftharpoons V \rangle$ is computationally zero-knowledge if there exists an expected PPT simulator S such that for every polynomial-size circuit family $V^* = (V_\lambda^*)_{\lambda \in \mathbb{N}}$,

$$\{\langle P(w) \rightleftharpoons V_\lambda^* \rangle(x)\}_{\substack{(x, w) \in \mathcal{R} \\ |x| = \lambda}} \approx_c \{S(x, V_\lambda^*)\}_{\substack{(x, w) \in \mathcal{R} \\ |x| = \lambda}} .$$

The protocol is honest-verifier computational zero-knowledge if the above is only guaranteed for the honest verifier V .

Interactive Batching. Here we define interactive batch protocols.

Definition 2.8 (Interactive Batch Protocol). A batch protocol for \mathcal{R} is an interactive protocol for $\bigcup_{t \in \mathbb{N}} \mathcal{R}^{\otimes t}$, where:

$$\mathcal{R}^{\otimes t} := \{(x_1, \dots, x_t), (w_1, \dots, w_t) : |x_1| = \dots = |x_t|, (x_1, w_1), \dots, (x_t, w_t) \in \mathcal{R}\} . \quad (1)$$

- The protocol's completeness and soundness errors ($\delta(\lambda, t)$ and $\varepsilon(\lambda, t)$) are defined to be its largest completeness and soundness errors, respectively, on any t instances (and any of their witnesses) of size λ
- The protocol has compression rate $\rho = \rho(\lambda, t)$, for instance length λ and number of instances t , if maximum total length of prover messages (over all such sets of instances) is ρt

Non-Interactive Protocols. We now define non-interactive protocols in the *common reference string* (CRS) model and the different notions of soundness and privacy that we need. This includes *distributional* formulations of NISWI and NISZK [Go193] (also see [CLP15, DNRS99]).

Definition 2.9 (Non-Interactive Protocol in CRS model). A non-interactive protocol $\langle P \rightarrow V \rangle$ in CRS model for an NP relation \mathcal{R} is a triple of algorithms (Gen, P, V) , with the following syntax:

- $\text{crs} \leftarrow \text{Gen}(1^\lambda)$: Given the instance size λ , the randomised set-up algorithm Gen outputs a CRS crs .

- $\pi \leftarrow P(crs, x, w)$: Given CRS crs , instance x , and witness w , the randomised prover outputs a proof π
- $b := V(crs, x, \pi)$: Given CRS crs , instance x , and proof π , the deterministic verifier returns a bit b representing accept or reject

In the following definitions, $\langle P \rightarrow V \rangle$ is a non-interactive protocol in CRS model for an NP relation \mathcal{R} .

Definition 2.10 (Completeness for Non-Interactive Protocols). $\langle P \rightarrow V \rangle$ has completeness error δ if for all large enough $\lambda \in \mathbb{N}$ and $(x, w) \in \mathcal{R}_\lambda$,

$$\Pr_{\substack{crs \leftarrow \text{Gen}(1^\lambda) \\ \pi \leftarrow P(crs, x, w)}} [V(crs, x, \pi) = 1] \geq 1 - \delta(\lambda).$$

For soundness, we will need the following notions.

Definition 2.11 (Non-Adaptive Computational Soundness). $\langle P \rightarrow V \rangle$ is non-adaptively computationally sound if for every polynomial-size circuit family of provers $P^* = (P_\lambda^*)_{\lambda \in \mathbb{N}}$, there is a negligible function μ , such that for all $\lambda \in \mathbb{N}$ and all $x \in \{0, 1\}^\lambda \setminus \mathcal{L}(\mathcal{R})$:

$$\Pr_{crs \leftarrow \text{Gen}(1^\lambda)} [V(crs, x, P_\lambda^*(crs, x)) = 1] \leq \mu(\lambda).$$

Definition 2.12 (Adaptive Computational Soundness). $\langle P \rightarrow V \rangle$ is adaptively computationally sound if, for every polynomial-size circuit family of provers $P^* = (P_\lambda^*)_{\lambda \in \mathbb{N}}$, there is a negligible function μ , such that for all $\lambda \in \mathbb{N}$:

$$\Pr_{crs \leftarrow \text{Gen}(1^\lambda)} [x \notin \mathcal{L}(\mathcal{R}_\lambda) \wedge V(crs, x, \pi) = 1] \leq \mu(\lambda),$$

where $(x, \pi) := P_\lambda^*(crs)$.

Definition 2.13 (SWI for Non-Interactive Protocols). $\langle P \rightarrow V \rangle$ is statistically witness-indistinguishable with (WI) error ε if, for all large enough λ and every $(x, w_0), (x, w_1) \in \mathcal{R}_\lambda$, we have:

$$(crs, \pi_0) \approx_\varepsilon (crs, \pi_1)$$

where $crs \leftarrow \text{Gen}(1^\lambda)$, $\pi_0 \leftarrow P(crs, x, w_0)$, and $\pi_1 \leftarrow P(crs, x, w_1)$.

Definition 2.14 (SZK for Non-Interactive Protocols). $\langle P \rightarrow V \rangle$ is statistically zero-knowledge with (ZK) error ε if there exists a PPT simulator S such that:

$$\{(crs, \pi)\}_{\substack{(x, w) \in \mathcal{R} \\ |x| = \lambda}} \approx_\varepsilon \{S(x)\}_{\substack{(x, w) \in \mathcal{R} \\ |x| = \lambda}}.$$

where $crs \leftarrow \text{Gen}(1^\lambda)$, $\pi \leftarrow P(crs, x, w)$.

Definition 2.15 (CZK for Non-Interactive Protocols). $\langle P \rightarrow V \rangle$ is computationally zero-knowledge if there exists a PPT simulator S such that:

$$\{(crs, \pi)\}_{\substack{(x, w) \in \mathcal{R} \\ |x| = \lambda}} \approx_c \{S(x)\}_{\substack{(x, w) \in \mathcal{R} \\ |x| = \lambda}}.$$

where $crs \leftarrow \text{Gen}(1^\lambda)$, $\pi \leftarrow P(crs, x, w)$.

Definition 2.16 (Distributional SWI for Non-Interactive Protocols). Let $\mathcal{D} = (D_\lambda)_{\lambda \in \mathbb{N}}$ be an efficiently-sampleable distribution, where D_λ is supported on triples (x, w_0, w_1) such that $(x, w_0), (x, w_1) \in \mathcal{R}_\lambda$. $\langle P \rightarrow V \rangle$ is distributionally SWI for D with error ε if for all large enough $\lambda \in \mathbb{N}$

$$\mathbb{E}_{(x, w_0, w_1) \leftarrow D_\lambda} [\text{SD}((crs, \pi_0), (crs, \pi_1))] \leq \varepsilon(\lambda),$$

where $crs \leftarrow \text{Gen}(1^\lambda)$, $\pi_0 \leftarrow P(crs, x, w_0)$, and $\pi_1 \leftarrow P(crs, x, w_1)$.

Definition 2.17 (Distributional SZK for Non-Interactive Protocols). Let $\mathcal{D} = (D_\lambda)_{\lambda \in \mathbb{N}}$ be an efficiently-sampleable distribution, where D_λ is supported on pairs $(x, w) \in \mathcal{R}_\lambda$. $\langle P \rightarrow V \rangle$ is distributionally SZK for D with error ε if there exists a PPT simulator S such that for all large enough $\lambda \in \mathbb{N}$

$$\mathbb{E}_{(x,w) \leftarrow D_\lambda} [\text{SD}((\text{crs}, \pi), S(x))] \leq \varepsilon(\lambda),$$

where $\text{crs} \leftarrow \text{Gen}(1^\lambda)$ and $\pi \leftarrow P(\text{crs}, x, w)$.

Dual-Mode Non-Interactive Protocols. It will be useful to define non-interactive protocols where the set-up algorithm has two different, indistinguishable modes [GOS12]. Looking ahead, the two modes that we work with are *statistically-hiding* mode and *adaptively-sound* mode (see Remark 5).

Definition 2.18 (Dual-Mode Non-Interactive Protocol). A dual-mode non-interactive protocol (in the CRS model) is a four-tuple of algorithms $(\text{Gen}, \text{AGen}, P, V)$, such that both (Gen, P, V) and (AGen, P, V) are non-interactive protocols (in the CRS model) with syntax as defined in Definition 2.9. In addition, we require mode indistinguishability, i.e., the CRSs generated by the two set-up modes Gen and AGen to be computationally indistinguishable.

Definition 2.19 (Dual-Mode (Distributional) NISZKA and NISWIA). A dual-mode ε -NISZKA (resp., ε -NISWI) $\Pi = \langle P \rightarrow V \rangle$ is a dual mode non-interactive protocol with the following requirements:

1. (Gen, P, V) is statistically ZK (resp., WI) with error ε . Therefore, we refer to this mode as the statistically-hiding mode.
2. (AGen, P, V) is adaptively computationally sound with negligible soundness error. Therefore, we refer to this mode as the adaptively-sound mode.

In dual-mode distributional ε -NISZKA (resp., ε -NISWIA), the hiding requirement in Item 1 is relaxed to its distributional formulation, i.e., Definition 2.17 (resp., Definition 2.16).

Remark 5 (On the Two Modes). Below we focus on dual-mode NISZKA, but analogous remarks apply also to dual-mode NISWIA, and the distributional counterparts.

- Whenever we say that a dual-mode NISZKA Π is statistically ZK (resp., adaptively sound), we are implicitly referring to its statistically-hiding (resp., adaptively-sound) mode.
- As a consequence of mode indistinguishability of dual-mode protocols, we can infer that: (i) the statistically-hiding mode has non-adaptive computational soundness (with negligible soundness error) and (ii) adaptive-soundness mode is computationally ZK with error $\varepsilon + \text{negl}(\lambda)$.

Remark 6 (On Adaptive Hiding). Note that the definition of statistical ZK in Definition 2.19 (via Definition 2.14) is non-adaptive (with respect to the CRS). However, since non-adaptive statistical ZK with error ε implies adaptive statistical ZK with error $\varepsilon \cdot 2^n$ and, then if ε can be made arbitrarily small, one gets adaptive statistical ZK (this will be the case in our constructions). By mode indistinguishability of dual-mode protocols, this also implies adaptive computational ZK in the adaptively-sound mode. An analogous observation also applies for dual-mode NISWIA.

Non-Interactive Batching. Here we define non-interactive batch protocols in the CRS model and the notion of soundness we need.

Definition 2.20 (Non-Interactive Batch Protocol in the CRS model). A non-interactive batch protocol in the CRS model is a four-tuple of algorithms $(\text{Gen}, \text{TGen}, P, V)$ with the following syntax::

- $\text{crs} \leftarrow \text{Gen}(1^\lambda, 1^t)$: Given the instance size λ and the number of instances t , the randomised set-up algorithm outputs a CRS crs

- $(crs^*, td) \leftarrow \text{TGen}(1^\lambda, 1^t, i^*)$: Given in addition an index $i^* \in [t]$, the trapdoored set-up algorithm outputs a CRS crs^* together with a trapdoor td
- $\pi \leftarrow \text{P}(crs, (x_1, \dots, x_t), (w_1, \dots, w_t))$: Given CRS crs , instances x_i , and witnesses w_i , the randomised prover outputs a proof π
- $b := \text{V}(crs, (x_1, \dots, x_t), \pi)$: Given CRS crs , instances x_i , and proof π , the deterministic verifier outputs a bit b representing accept or reject

Here, the prover's communication is just the proof π , and the compression rate is defined with respect to this.

The following definitions of soundness properties are adapted from [CJJ22], though they have been simplified and slightly weakened as this is sufficient for our purposes.

Definition 2.21 (CRS Indistinguishability). *A batch protocol $(\text{Gen}, \text{TGen}, \text{P}, \text{V})$ is CRS-indistinguishable if for every polynomial t and every $i(\lambda) \in [t(\lambda)]$, the distributions of $\text{Gen}(1^\lambda, 1^{t(\lambda)})$ and crs^* sampled from $\text{TGen}(1^\lambda, 1^{t(\lambda)}, i(\lambda))$ are computationally indistinguishable.*

Definition 2.22 (Somewhere Soundness). *A batch protocol $(\text{Gen}, \text{TGen}, \text{P}, \text{V})$ for a relation \mathcal{R} is somewhere computationally sound if it satisfies CRS indistinguishability, and for every polynomial t and polynomial-size circuit family of provers $\text{P}^* = (\text{P}_\lambda^*)_{\lambda \in \mathbb{N}}$, there is a negligible function μ such that for all $\lambda \in \mathbb{N}$, letting $t = t(\lambda)$, and for every $i^* \in [t]$:*

$$\Pr_{crs} [x_{i^*} \notin \mathcal{L}(\mathcal{R}_\lambda) \wedge \text{V}(crs, (x_1, \dots, x_t), \pi) = 1] \leq \mu(\lambda),$$

where $(crs, td) \leftarrow \text{TGen}(1^\lambda, 1^t, i^*)$, and $((x_1, \dots, x_t), \pi) := \text{P}_\lambda^*(crs, i^*)$.

List of Proof Systems. The proof systems that are relevant to this paper are listed below.

- Statistical Zero-Knowledge Argument (SZKA): Interactive protocol (Definition 2.1) satisfying Definitions 2.2, 2.4 and 2.6.
- Statistical Witness-Indistinguishable Argument (SWIA): Interactive protocol satisfying Definitions 2.2, 2.4 and 2.5.
- Honest-Verifier SZK (HVSZK) proof: Interactive protocol satisfying Definition 2.2, Definition 2.6 restricted to *honest verifiers* and Definition 2.3.
- Honest-Verifier SWI (HVSWI) proof (resp., argument): Interactive protocol satisfying Definition 2.2, Definition 2.5 restricted to *honest verifiers* and Definition 2.3 (resp., Definition 2.4).
- Non-interactive SZKA (NISZKA) with non-adaptive (resp., adaptive) soundness: Non-interactive protocol (Definition 2.9) satisfying Definitions 2.10 and 2.14, and Definition 2.11 (resp., Definition 2.12).
- Non-interactive SWIA (NISWIA) with non-adaptive (resp., adaptive) soundness: Non-interactive protocol satisfying Definitions 2.10 and 2.13 and Definition 2.11 (resp., Definition 2.12).
- Distributional NISWIA (resp., NISZKA) with adaptive soundness: Non-interactive protocol satisfying Definition 2.10, Definition 2.16 (resp., Definition 2.17) and Definition 2.12.
- Batch proof: Interactive batch protocol (Definition 2.8) satisfying Definitions 2.2 and 2.3.
- Batch argument (BARG): Interactive batch protocol satisfying Definitions 2.2 and 2.4.
- Non-interactive BARG (NIBARG): Non-interactive batch protocol (Definition 2.20) satisfying Definitions 2.10 and 2.22.

3 Statistical Witness Indistinguishability from Batching

In this section, we prove that a sufficiently shrinking batch protocol for a relation can be used to construct an honest-verifier statistically witness-indistinguishable protocol for it with the same soundness properties. This is captured by the following theorem. In Section 3.2, we prove a related theorem that preserves non-interactivity and stronger notions of computational soundness, which is required for our results for non-interactive BARGs. Recall that for a relation \mathcal{R} and polynomial t , $\mathcal{R}^{\otimes t}$ denotes the product relation (as in Definition 2.8).

Theorem 3.1. *Consider an NP relation \mathcal{R} . Suppose it has a batch protocol $\Pi = \langle P \leftrightarrow V \rangle$ that, when run on some polynomial $t = t(\lambda)$ instances of size λ , has compression rate $\rho = \rho(\lambda) < 1$. Then, \mathcal{R} has a protocol $\Pi_{\text{WI}} = \langle P_{\text{WI}} \leftrightarrow V_{\text{WI}} \rangle$ with the following properties (on instances of size λ):*

- Π_{WI} is HVSWI with error $O(\sqrt{\rho})$.
- Π_{WI} has the same completeness error as Π run on t instances.
- If Π is statistically sound, then so is Π_{WI} , with the same soundness error as Π run on t instances.
- If Π is computationally sound, then so is Π_{WI} .
- If P is computed by a family of polynomial-sized circuits, then so is P_{WI} ; and V_{WI} runs in uniform polynomial-time given blackbox access to V .
- The communication and round complexity in Π_{WI} are the same as those of Π , plus an additional message sent by P_{WI} at the start that is $(\lambda \cdot t + \log t)$ bits long.

Fix some relation \mathcal{R} for which there is a batch protocol $\langle P \leftrightarrow V \rangle$ with compression rate ρ as hypothesized. We will show how to construct from this a protocol $\langle P_{\text{WI}} \leftrightarrow V_{\text{WI}} \rangle$ for \mathcal{R} that inherits its soundness properties and is, in addition, HVSWI. This protocol follows the template in Fig. 1, which is parameterized by an ensemble of distributions \mathcal{D} and a function t , which we will instantiate later.

Given a batch protocol $\langle P \leftrightarrow V \rangle$, a function $t : \mathbb{N} \rightarrow \mathbb{N}$, and an ensemble of distributions $\mathcal{D} = (D_\lambda)_{\lambda \in \mathbb{N}}$, where the support of D_λ is contained in $(\{0, 1\}^\lambda \times \{0, 1\}^*)^{t(\lambda)}$, the protocol $\langle P_{\text{WI}} \leftrightarrow V_{\text{WI}} \rangle^{\langle \mathcal{D}, t \rangle}$ works as follows given an instance $x \in \{0, 1\}^\lambda$ and a witness $w \in \{0, 1\}^*$:

1. P_{WI} generates a sample $\{(x_i, w_i)\}_{i \in [t(\lambda)]}$ from D_λ , and samples $j \leftarrow [t(\lambda)]$
2. P_{WI} sends all the x_i 's and j to V_{WI}
3. P_{WI} and V_{WI} run the protocol $\langle P \leftrightarrow V \rangle$ on the input $(x_1, \dots, x_{j-1}, x, x_{j+1}, \dots, x_{t(\lambda)})$, with P_{WI} using $(w_1, \dots, w_{j-1}, w, w_{j+1}, \dots, w_{t(\lambda)})$ as the witnesses
4. V_{WI} accepts iff the verifier V in the above execution accepts

Figure 1: Template for constructing HVSWI protocols from batch protocols

We next state lemmas capturing the properties of this protocol, and use them to prove Theorem 3.1. The proof of Lemma 3.2 is included below, and Lemma 3.3 is proven in Section 3.1.

Lemma 3.2 (Completeness and Soundness). *Suppose $\langle P \leftrightarrow V \rangle$ is a batch protocol for a relation \mathcal{R} . Let t be any polynomial and $\mathcal{D} = (D_\lambda)_{\lambda \in \mathbb{N}}$ be such that the support of D_λ is contained within $\mathcal{R}_\lambda^{\otimes t(\lambda)}$. Then, the protocol $\langle P_{\text{WI}} \leftrightarrow V_{\text{WI}} \rangle$ in Fig. 1, when instantiated with $\langle P \leftrightarrow V \rangle$, \mathcal{D} and t , is a protocol for \mathcal{R} that satisfies the following:*

1. If $\langle P \rightleftharpoons V \rangle$ has completeness error $\delta(\lambda)$ when run with $t(\lambda)$ instances of size λ , then $\langle P_{\text{WI}} \rightleftharpoons V_{\text{WI}} \rangle^{(\mathcal{D}, t)}$ has completeness error $\delta(\lambda)$.
2. If $\langle P \rightleftharpoons V \rangle$ has statistical soundness error $\epsilon(\lambda)$ when run with $t(\lambda)$ instances of size λ , then $\langle P_{\text{WI}} \rightleftharpoons V_{\text{WI}} \rangle^{(\mathcal{D}, t)}$ has statistical soundness error $\epsilon(\lambda)$.
3. If $\langle P \rightleftharpoons V \rangle$ is computationally sound, then $\langle P_{\text{WI}} \rightleftharpoons V_{\text{WI}} \rangle^{(\mathcal{D}, t)}$ is also computationally sound.

Proof. Fix any x such that $|x| = \lambda$, and denote $t(\lambda)$ by t . As all the (x_i, w_i) 's sampled from D_λ are contained in \mathcal{R} , the input $(x_1, \dots, x_{j-1}, x, x_{j+1}, \dots, x_t)$ is contained in $\mathcal{R}^{\otimes t}$ if and only if there is some w such that $(x, w) \in \mathcal{R}$. The completeness and statistical soundness errors of $\langle P \rightleftharpoons V \rangle$ thus carry over immediately to $\langle P_{\text{WI}} \rightleftharpoons V_{\text{WI}} \rangle$ as stated in the theorem.

For computational soundness, suppose there is a malicious prover P_{WI}^* that can make V_{WI} accept with probability μ given an $x \notin \mathcal{R}$. Then, without loss of generality, there exists a $j \in [t]$ and (x_1, \dots, x_t) such that P_{WI}^* can make V_{WI} accept with probability μ with the first message being $(x_1, \dots, x_{j-1}, x, x_{j+1}, \dots, x_t)$ and j . As V_{WI} is just emulating the verifier V , this means there is a P^* that emulates P_{WI}^* and makes V accept on this input with probability μ . Further, if P_{WI}^* is polynomial-time, so is P^* , as t is a polynomial. If $\mu(\lambda)$ is non-negligible, this breaks computational soundness of $\langle P \rightleftharpoons V \rangle$, proving the theorem. \square

Lemma 3.3 (Witness Indistinguishability). *Consider a batch protocol $\langle P \rightleftharpoons V \rangle$ for a relation \mathcal{R} that has polynomial-sized witnesses. For a polynomial t , when the protocol is run with $t(\lambda)$ instances of size λ , suppose the total communication from the prover is at most $\rho(\lambda)t(\lambda)$ bits for some function ρ . Then, there is an efficiently sampleable ensemble of distributions $\mathcal{D} = (D_\lambda)_{\lambda \in \mathbb{N}}$, where D_λ is supported in $\mathcal{R}_\lambda^{\otimes t(\lambda)}$, such that the protocol $\langle P_{\text{WI}} \rightleftharpoons V_{\text{WI}} \rangle$ in Fig. 1, when instantiated with $\langle P \rightleftharpoons V \rangle$, \mathcal{D} , and t , is HVSWI with error $O(\sqrt{\rho(\lambda)})$.*

Using Lemmas 3.2 and 3.3 (the latter is proved in Section 3.1), we are ready to prove Theorem 3.1.

Proof of Theorem 3.1. Consider a relation \mathcal{R} with polynomial-sized witnesses and a batch protocol $\langle P \rightleftharpoons V \rangle$ that, for some polynomial t , when run on $t(\lambda)$ instances of size λ , has completeness error $\delta(\lambda)$, statistical soundness error $\epsilon(\lambda)$, and at most $\rho(\lambda)t(\lambda)$ bits of communication from the prover for some function ρ . Let \mathcal{D} be the ensemble guaranteed by Lemma 3.3, and consider the protocol $\langle P_{\text{WI}} \rightleftharpoons V_{\text{WI}} \rangle$ as described in Fig. 1 instantiated with this \mathcal{D} and t . This protocol has the following properties:

- Lemma 3.3 implies that this protocol is HVSWI with WI error $O(\sqrt{\rho(\lambda)})$.
- Lemma 3.2 implies that its completeness and statistical soundness errors are $\delta(\lambda)$ and $\epsilon(\lambda)$, respectively.
- Lemma 3.2 implies that if $\langle P \rightleftharpoons V \rangle$ is computationally sound, then so is $\langle P_{\text{WI}} \rightleftharpoons V_{\text{WI}} \rangle$.
- All V_{WI} does is run V on an input provided by P_{WI} and accept iff it accepts. P_{WI} also simply runs P on an input and witnesses, and in addition computes samples from D_λ and $[t(\lambda)]$, which can be done in non-uniform polynomial time since \mathcal{D} is efficiently sampleable.
- In addition to the messages of $\langle P \rightleftharpoons V \rangle$, the only additional communication in $\langle P_{\text{WI}} \rightleftharpoons V_{\text{WI}} \rangle$ is the initial prover message consisting of $t(\lambda)$ instances and an element of $[t(\lambda)]$.

The above arguments prove the respective properties of the protocol promised by the theorem. \square

3.1 Witness Indistinguishability

In this section, we prove Lemma 3.3 about the witness indistinguishability of the protocol from Fig. 1. We will first come up with an ensemble of distributions \mathcal{D} that, when used to instantiate this protocol, will make the protocol witness-indistinguishable. Fix any batch protocol $\langle P \rightleftharpoons V \rangle$ for a relation \mathcal{R} , an instance length λ , witness length m , and the number of batch instances t . Suppose that when $\langle P \rightleftharpoons V \rangle$ is run on t instances of length λ , each with witness of length m , the total prover communication is at most ρt , where the compression rate ρ is less than 1.

Compressing Functions. We will use the fact that compressing functions necessarily lose information to make such a prover lose information about the witness we want to hide. This property of compression is captured by the following lemma by Dell, building on the work of Drucker [Dru15]. Similar consequences of compression have been used in the context of cryptography in the past, for instance to construct Oblivious Transfer from Private Information Retrieval protocols [DMO00, Lemma 1].

Lemma 3.4 ([Dell16, Lemma 9]). *Let $t \in \mathbb{N}$, $\rho \in [0, 1)$, and B be the uniform distribution over $\{0, 1\}^t$. For any randomized mapping $f : \{0, 1\}^t \rightarrow \{0, 1\}^{\rho t}$, with $j \leftarrow [t]$, we have:*

$$\mathbb{E}_{j \leftarrow [t]} [\text{SD}(f(B|_{j \leftarrow 0}), f(B|_{j \leftarrow 1}))] \leq \sqrt{2 \ln 2 \cdot \rho},$$

where $B|_{j \leftarrow b}$ is the result of drawing a sample $(b_1, \dots, b_t) \leftarrow B$ and then replacing b_j with b .

We now define a function that captures the knowledge gained by the honest verifier by interacting with the honest prover in the protocol $\langle \text{P} \leftrightarrow \text{V} \rangle$. Its input consists of t instances $x_1, \dots, x_t \in \{0, 1\}^\lambda$, potential witnesses $w_1, \dots, w_t \in \{0, 1\}^m$, and potential random string r of V . We use \mathbf{x} to denote (x_1, \dots, x_t) for brevity.

$f((x_1, \dots, x_t), (w_1, \dots, w_t), r)$:

1. Run $\langle \text{P} \leftrightarrow \text{V} \rangle$ with input (x_1, \dots, x_t) , using r as randomness for V , and with (w_1, \dots, w_t) as the witnesses provided to P
2. Output the sequence of prover messages in the above execution

In addition, for any pair of tuples of t potential witnesses $y_1, \dots, y_t \in \{0, 1\}^m$ and $z_1, \dots, z_t \in \{0, 1\}^m$, we define the following function on bits b_i .

$g_{\mathbf{x}, \mathbf{y}, \mathbf{z}, r}(b_1, \dots, b_t)$:

1. For each $i \in [t]$, set $w_i = y_i$ if $b_i = 0$, and $w_i = z_i$ if $b_i = 1$
2. Output $f(\mathbf{x}, \mathbf{w}, r)$

The proposition below follows immediately from Lemma 3.4 and the compression of the protocol.

Proposition 3.5. *For any tuple of $x_i \in \{0, 1\}^\lambda$, $y_i, z_i \in \{0, 1\}^m$, and any r of the appropriate length, letting B be the uniform distribution over $\{0, 1\}^t$,*

$$\mathbb{E}_{j \leftarrow [t]} [\text{SD}(g_{\mathbf{x}, \mathbf{y}, \mathbf{z}, r}(B|_{j \leftarrow 0}), g_{\mathbf{x}, \mathbf{y}, \mathbf{z}, r}(B|_{j \leftarrow 1}))] \leq \sqrt{2 \ln 2 \cdot \rho}.$$

Interpreting the function g in terms of the function f then gives the following.

Proposition 3.6. *Consider any t -tuple of $x_i \in \{0, 1\}^\lambda$, $y_i, z_i \in \{0, 1\}^m$, and any r of the appropriate length. For $i \in [t]$, let W_i be set to y_i or z_i uniformly at random. Then,*

$$\mathbb{E}_{j \leftarrow [t]} [\text{SD}(f(\mathbf{x}, \mathbf{W}|_{j \leftarrow y_j}, r), f(\mathbf{x}, \mathbf{W}|_{j \leftarrow z_j}, r))] \leq \sqrt{2 \ln 2 \cdot \rho}.$$

Two-Player Zero-Sum Games. Consider a two-player zero-sum game $G = (R, C, p)$, where R is the set of pure strategies for the “row” player, C the same for the “column” player, and $p : R \times C \rightarrow \mathbb{R}$ is the payoff function. Let ρ and κ denote mixed strategies for the two players, which are distributions over R and C , respectively. The *value* of this game is defined as:

$$\text{val}(G) = \min_{\rho} \max_{\kappa} \mathbb{E}_{\substack{r \leftarrow \rho \\ c \leftarrow \kappa}} [p(r, c)] = \max_{\kappa} \min_{\rho} \mathbb{E}_{\substack{r \leftarrow \rho \\ c \leftarrow \kappa}} [p(r, c)].$$

where the equality follows from von Neumann’s minimax theorem [vN28]. Lipton and Young prove the following sparse minimax theorem that will be useful for us to infer sampleable mixed strategies.

Lemma 3.7 ([LY94]). *Consider any two-player zero-sum game $G = (R, C, p)$ such that $p(r, c) \in [0, 1]$ for any (r, c) . For any $\epsilon > 0$, there is multiset $S \subseteq R$ of size $\Theta(\log |C|/\epsilon^2)$ such that for every $c \in C$:*

$$\mathbb{E}_{r \leftarrow S} [p(r, c)] \leq \text{val}(G) + \epsilon.$$

That is, there is a sparse mixed strategy that is almost as good as the optimal strategy over R . We will now define a game that captures the witness indistinguishability of the protocol described in Fig. 1, and use the above lemma to find a distribution \mathcal{D}_λ with which to instantiate the protocol. Note that this is the first point in the proof where we involve the relation \mathcal{R} that the protocols are for.

The game $G_W = (R, C, p)$ is defined with the following sets of pure strategies:

- $R = \{(\mathbf{x}, \mathbf{y}, \mathbf{z})\}$, where each vector is of length t , $x_i \in \{0, 1\}^\lambda$, $y_i, z_i \in \{0, 1\}^m$, and $(x_i, y_i), (x_i, z_i) \in \mathcal{R}$
- $C = \{(x, y, z)\}$, where $x \in \{0, 1\}^\lambda$, $y, z \in \{0, 1\}^m$, and $(x, y), (x, z) \in \mathcal{R}$

Given $r = (\mathbf{x}, \mathbf{y}, \mathbf{z}) \in R$, for each $i \in [t]$, define a random variable W_i that is set to y_i or z_i uniformly at random. The payoff function $p : R \times C \rightarrow [0, 1]$ is then defined as follows, with r distributed uniformly over the appropriate domain:

$$p((\mathbf{x}, \mathbf{y}, \mathbf{z}), (x, y, z)) = \mathbb{E}_{j \leftarrow [t], r} [\text{SD}(f(\mathbf{x}|_{j \leftarrow x}, \mathbf{W}|_{j \leftarrow y}, r), f(\mathbf{x}|_{j \leftarrow x}, \mathbf{W}|_{j \leftarrow z}, r)))]$$

Proposition 3.8. *The value of the game G_W defined above is at most $\sqrt{2 \ln 2} \cdot \rho$.*

Proof. It is sufficient to show that for any distribution (X, Y, Z) over C , there is a distribution $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ over R such that the expected payoff under these strategies is at most the required bound. Given such a distribution (X, Y, Z) , consider $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ defined by $(x_i, y_i, z_i) \leftarrow (X, Y, Z)$ for $i \in [t]$. The expected payoff is then as follows, with each W_i set to y_i or z_i at random:

$$\mathbb{E}_{\substack{(x_1, y_1, z_1) \leftarrow (X, Y, Z) \\ \vdots \\ (x_t, y_t, z_t) \leftarrow (X, Y, Z)}} \mathbb{E}_{(x, y, z) \leftarrow (X, Y, Z)} \mathbb{E}_{j \leftarrow [t], r} [\text{SD}(f(\mathbf{x}|_{j \leftarrow x}, \mathbf{W}|_{j \leftarrow y}, r), f(\mathbf{x}|_{j \leftarrow x}, \mathbf{W}|_{j \leftarrow z}, r)))] .$$

Noting that r and j are sampled independently of all the other quantities¹², by linearity of expectation, the above is the same as:

$$\mathbb{E}_{j \leftarrow [t], r} \left[\mathbb{E}_{\substack{(x_1, y_1, z_1) \leftarrow (X, Y, Z) \\ \vdots \\ (x_t, y_t, z_t) \leftarrow (X, Y, Z)}} \mathbb{E}_{(x, y, z) \leftarrow (X, Y, Z)} [\text{SD}(f(\mathbf{x}|_{j \leftarrow x}, \mathbf{W}|_{j \leftarrow y}, r), f(\mathbf{x}|_{j \leftarrow x}, \mathbf{W}|_{j \leftarrow z}, r)))] \right] .$$

¹²This requirement of independence, specifically between r and x , is why this proof only provides honest-verifier SWI and does not work for a malicious verifier. The WI of our protocol could potentially be broken by a malicious verifier that chooses r based on x .

As (x, y, z) and (x_j, y_j, z_j) are identically distributed and are independent of all other variables, this is the same as:

$$\mathbb{E}_{j \leftarrow [t], r} \left[\mathbb{E}_{\substack{(x_1, y_1, z_1) \leftarrow (X, Y, Z) \\ \vdots \\ (x_t, y_t, z_t) \leftarrow (X, Y, Z)}} \left[\text{SD} \left(f(\mathbf{x}, \mathbf{W}|_{j \leftarrow y_j}, r), f(\mathbf{x}, \mathbf{W}|_{j \leftarrow z_j}, r) \right) \right] \right].$$

By Proposition 3.6 and linearity of expectation, the above is at most $\sqrt{2 \ln 2 \cdot \rho}$, which proves the proposition. \square

By Lemma 3.7 and Proposition 3.8, we have the following proposition.

Proposition 3.9. *For every $\epsilon > 0$, there is a multiset $S = \{(\mathbf{x}, \mathbf{y}, \mathbf{z})\}$ of size $\Theta((\lambda + m)/\epsilon^2)$ such that:*

- for every $i \in [t]$, both (x_i, y_i) and (x_i, z_i) are in \mathcal{R}
- for every $x \in \{0, 1\}^\lambda$ and $y, z \in \{0, 1\}^m$ such that $(x, y), (x, z) \in \mathcal{R}$,

$$\mathbb{E}_{\substack{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \leftarrow S \\ w_i \leftarrow \{y_i, z_i\} \\ j \leftarrow [t], r}} \left[\text{SD} \left(f(\mathbf{x}|_{j \leftarrow x}, \mathbf{w}|_{j \leftarrow y}, r), f(\mathbf{x}|_{j \leftarrow x}, \mathbf{w}|_{j \leftarrow z}, r) \right) \right] \leq \sqrt{2 \ln 2 \cdot \rho} + \epsilon.$$

Proof of Lemma 3.3. We can now describe the distribution D_λ that we will instantiate the protocol in Fig. 1 with. Recall that $\rho(\lambda)$ is the compression rate of the batch protocol we started with when run on $t(\lambda)$ instances of size λ .

Let S be the multiset guaranteed by Proposition 3.9 for $\epsilon = \sqrt{\rho(\lambda)}$. The distribution D_λ is sampled as follows:

1. Sample $(\mathbf{x}, \mathbf{y}, \mathbf{z}) \leftarrow S$.
2. For each $i \in [t(\lambda)]$, set w_i to y_i or z_i uniformly at random.
3. Output $\{(x_i, w_i)\}_{i \in [t(\lambda)]}$.

As S is of size $\Theta((\lambda + m(\lambda))/\epsilon^2) = \Theta((\lambda + m(\lambda))/\rho(\lambda))$, which is polynomial in λ , the distribution D_λ can be sampled non-uniformly in $\text{poly}(\lambda)$ time. In any element $(\mathbf{x}, \mathbf{y}, \mathbf{z})$ of S , we are guaranteed that each (x_i, y_i) and (x_i, z_i) is in \mathcal{R} . So the support of D_λ is contained in $\mathcal{R}_\lambda^{\otimes t(\lambda)}$, as required.

To argue HVSWI of the protocol $\langle P_{\text{WI}} \rightleftharpoons V_{\text{WI}} \rangle$ when instantiated with this distribution, we need to show that for every possible pair $(x, y), (x, z) \in \mathcal{R}_\lambda$, the views of the verifier V_{WI} on input x when P_{WI} uses y or z as the witness are statistically close. Fix any such pair.

Note that for any (\mathbf{x}, \mathbf{w}) sampled from D_λ , the view of V_{WI} on input x , when P uses witness w , is completely determined by the following quantities: \mathbf{x}, j, r , and $f(\mathbf{x}|_{j \leftarrow x}, \mathbf{w}|_{j \leftarrow w}, r)$ – all this is missing is the sequence of verifier messages in the protocol, which can be reconstructed efficiently given the verifier randomness r and the prover messages $f(\cdot)$. Thus, by the data processing inequality, the statistical distance between the views of V_{WI} in the cases where P_{WI} uses witness y or z is at most the following, where $(\mathbf{x}, \mathbf{w}) \leftarrow D_\lambda, j \leftarrow [t(\lambda)]$, and r is over the appropriate domain:

$$\text{SD} \left((\mathbf{x}, j, r, f(\mathbf{x}|_{j \leftarrow x}, \mathbf{w}|_{j \leftarrow y}, r)), (\mathbf{x}, j, r, f(\mathbf{x}|_{j \leftarrow x}, \mathbf{w}|_{j \leftarrow z}, r)) \right).$$

Taking into account the definition of D_λ , this is equal to:

$$\mathbb{E}_{\substack{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \leftarrow \mathcal{S} \\ w_i \leftarrow \{y_i, z_i\} \\ j \leftarrow [t], r}} [\text{SD}(f(\mathbf{x}|_{j \leftarrow x}, \mathbf{w}|_{j \leftarrow y}, r), f(\mathbf{x}|_{j \leftarrow x}, \mathbf{w}|_{j \leftarrow z}, r))],$$

which, by Proposition 3.9, is at most $\sqrt{2 \ln 2 \cdot \rho(\lambda)} + \epsilon = O(\sqrt{\rho(\lambda)})$, proving the lemma. \square

Remark 7. *The prover P_{WI} in protocol $\langle P_{\text{WI}} \leftrightarrow V_{\text{WI}} \rangle$ we construct is non-uniform even if the prover P from the original batch protocol is uniform. This is because the minimax theorem we use (Lemma 3.7), while constructive, is not uniform. An interesting question here is whether a uniform version of the minimax theorem can be used instead to preserve uniformity of the prover. As far as we can tell, existing uniform minimax theorems ([VZ13], for instance) do not seem useful for this purpose. They require the payoff of the game to be efficiently computable given the strategies, which does not seem to be the case here as it involves computing the statistical distance between two rather arbitrary distributions.*

Remark 8. *The bound of $O(\sqrt{\rho})$ in the statements above (and particularly in Lemma 3.4) is optimal upto constant factors. In the case of Lemma 3.4, a function g that splits its input into blocks of size $\Theta(1/\rho)$ and outputs the majority of the bits in each block witnesses this optimality. This can then be extended to proof systems, where the bits may represent predicates that distinguish between two witnesses.*

3.2 Non-Interactive Protocols

In this section, we prove a version of Theorem 3.1 for non-interactive protocols. It preserves non-interactivity and considers adaptive notions of soundness. Along the way, we show in Theorem 3.13 that if we weaken the requirement to *distributional* SWI (Definition 2.16), then it is possible to avoid non-uniformity of the prover in Theorem 3.10.

Theorem 3.10. *Suppose an NP relation \mathcal{R} has a non-interactive batch protocol $\Pi = (\text{Gen}, \text{TGen}, P, V)$ that, when run on some polynomial $t = t(\lambda)$ instances of size λ , has compression rate $\rho = \rho(\lambda) < 1$. Then, \mathcal{R} has a non-interactive protocol $\Pi_{\text{WI}} = (\text{Gen}_{\text{WI}}, P_{\text{WI}}, V_{\text{WI}})$ with the following properties (on instances of size λ):*

- Π_{WI} is SWI with error $O(\sqrt{\rho})$.
- If Π is CRS-indistinguishable, then Π_{WI} has completeness error negligibly close to that of Π run on t instances.
- If Π is somewhere computationally sound, then Π_{WI} is adaptively computationally sound.
- If P is computed by a family of polynomial-sized circuits, then so is P_{WI} ; and V_{WI} and Gen_{WI} run in uniform polynomial-time given blackbox access to V and TGen , respectively.
- The length of the proof in Π_{WI} is that in Π plus an additional $\lambda \cdot t$ bits. The length of the CRS is the same.

Fix some relation \mathcal{R} for which there is a non-interactive batch protocol $(\text{Gen}, \text{TGen}, P, V)$ with compression rate ρ as hypothesized. We will show how to construct from this a non-interactive SWI protocol $(\text{Gen}_{\text{WI}}, P_{\text{WI}}, V_{\text{WI}})$ for \mathcal{R} in a manner similar to that earlier in this section for general interactive protocols. This protocol follows the template in Fig. 1, which is parametrised by an ensemble of distributions \mathcal{D} and a function t , which we will instantiate later.

Given a non-interactive batch protocol $(\text{Gen}, \text{TGen}, \text{P}, \text{V})$, a function $t : \mathbb{N} \rightarrow \mathbb{N}$, and an ensemble of distributions $\mathcal{D} = (D_\lambda)_{\lambda \in \mathbb{N}}$, where the support of D_λ is contained in $(\{0, 1\}^\lambda \times \{0, 1\}^*)^{t(\lambda)}$, the non-interactive protocol $(\text{Gen}_{\text{WI}}, \text{P}_{\text{WI}}, \text{V}_{\text{WI}})$ are as follows.

$\text{Gen}_{\text{WI}}(1^\lambda)$:

- Sample $j \leftarrow [t(\lambda)]$, and $\text{crs} \leftarrow \text{TGen}(1^\lambda, 1^{t(\lambda)}, j)$.
- Output (j, crs) .

$\text{P}_{\text{WI}}((j, \text{crs}), x, w)$:

- Sample $\{(x_i, w_i)\}_{i \in [t(\lambda)]}$ from D_λ .
- Compute $\pi \leftarrow \text{P}(\text{crs}, (x_1, \dots, x_{j-1}, x, x_{j+1}, \dots, x_{t(\lambda)}), (w_1, \dots, w_{j-1}, w, w_{j+1}, \dots, w_{t(\lambda)}))$.
- Output (\mathbf{x}, π) .

$\text{V}_{\text{WI}}((j, \text{crs}), x, (\mathbf{x}, \pi))$:

- Accepts iff $\text{V}(\text{crs}, (x_1, \dots, x_{j-1}, x, x_{j+1}, \dots, x_{t(\lambda)}), \pi)$ accepts.

Figure 2: Template for constructing non-interactive SWI protocols from non-interactive batch protocols

We next state lemmas capturing the properties of this protocol, and use them to prove Theorem 3.10.

Lemma 3.11 (Completeness and Soundness). *Suppose $\Pi = (\text{Gen}, \text{P}, \text{V})$ is a non-interactive batch protocol for a relation \mathcal{R} . Let t be any polynomial and $\mathcal{D} = (D_\lambda)_{\lambda \in \mathbb{N}}$ be such that the support of D_λ is contained within $\mathcal{R}_\lambda^{\otimes t(\lambda)}$. Then, the protocol $\Pi_{\text{WI}} = (\text{Gen}_{\text{WI}}, \text{P}_{\text{WI}}, \text{V}_{\text{WI}})$ in Fig. 2, when instantiated with Π , \mathcal{D} and t , is a non-interactive protocol for \mathcal{R} that satisfies the following:*

1. *If Π has completeness error $\delta(\lambda)$ when run with $t(\lambda)$ instances of size λ and is CRS-indistinguishable, then Π_{WI} has completeness error at most $\delta(\lambda) + \text{negl}(\lambda)$.*
2. *If Π is somewhere computationally sound, then Π_{WI} is adaptively computationally sound.*

Proof. Let $t = t(\lambda)$. If in $\text{Gen}_{\text{WI}}(1^\lambda)$, the crs had been sampled from $\text{Gen}(1^\lambda, 1^t)$ instead of $\text{TGen}(1^\lambda, 1^t, j)$, then the completeness of Π_{WI} follows that of Π , with the same error $\delta(\lambda)$ (by the same arguments as in Lemma 3.2). By the CRS-indistinguishability of Π , and as both P_{WI} and V_{WI} are polynomial-time algorithms, making this change in Gen_{WI} can only change the completeness error by a negligible amount.

For soundness, suppose there is a malicious prover P_{WI}^* and a non-negligible function μ such that, with $\text{crs} \leftarrow \text{Gen}_{\text{WI}}(1^\lambda)$ and $(x, \pi) \leftarrow \text{P}_{\text{WI}}^*(\text{crs})$ we have:

$$\Pr [x \notin \mathcal{L}(\mathcal{R}_\lambda) \wedge \text{V}_{\text{WI}}(\text{crs}, x, \pi) \text{ accepts}] \geq \mu(\lambda).$$

By the definition of the protocol, the above is the same as the following: with $j \leftarrow [t]$, $\text{crs} \leftarrow \text{TGen}(1^\lambda, 1^t, j)$, $(x, \mathbf{x}, \pi) \leftarrow \text{P}_{\text{WI}}^*(j, \text{crs})$,

$$\Pr [x \notin \mathcal{L}(\mathcal{R}_\lambda) \wedge \text{V}(\text{crs}, \mathbf{x}|_{j \leftarrow x}, \pi) \text{ accepts}] \geq \mu(\lambda),$$

which immediately contradicts the somewhere computational soundness of $(\text{Gen}, \text{TGen}, \text{P}, \text{V})$ if μ is non-negligible. This proves the lemma. \square

Lemma 3.12 (Witness Indistinguishability). *Consider a non-interactive batch protocol $\Pi = (\text{Gen}, \text{TGen}, \text{P}, \text{V})$ for a relation \mathcal{R} that has polynomial-sized witnesses. For a polynomial t , when the protocol is run with $t(\lambda)$ instances of size λ , suppose the length of the proof is at most $\rho(\lambda)t(\lambda)$ bits for some function ρ . Then, there is an efficiently sampleable ensemble of distributions $\mathcal{D} = (D_\lambda)_{\lambda \in \mathbb{N}}$, where D_λ is supported in $\mathcal{R}_\lambda^{\otimes t(\lambda)}$, such that the protocol $(\text{Gen}_{\text{WI}}, \text{P}_{\text{WI}}, \text{V}_{\text{WI}})$ in Fig. 2, when instantiated with Π , \mathcal{D} , and t , is SWI with error $O(\sqrt{\rho(\lambda)})$.*

Proof Sketch. The proof of this lemma is identical to that of Lemma 3.3, with the only difference being that instead of the verifier's random string r , here we use the CRS sampled by Gen_{WI} . \square

Proof of Theorem 3.10. Consider a relation \mathcal{R} with polynomial-sized witnesses and a non-interactive batch protocol $\Pi = (\text{Gen}, \text{TGen}, \text{P}, \text{V})$ that, for some polynomial t , when run on $t(\lambda)$ instances of size λ , has completeness error $\delta(\lambda)$, statistical soundness error $\epsilon(\lambda)$, and proofs of length at most $\rho(\lambda)t(\lambda)$ bits for some function ρ . Let \mathcal{D} be the ensemble guaranteed by Lemma 3.12, and consider the protocol $\Pi_{\text{WI}} = (\text{Gen}_{\text{WI}}, \text{P}_{\text{WI}}, \text{V}_{\text{WI}})$ as described in Fig. 2 instantiated with this \mathcal{D} and t . This protocol has the following properties:

- Lemma 3.12 implies that this protocol is SWI with WI error $O(\sqrt{\rho(\lambda)})$.
- Lemma 3.11 implies that its completeness error is $\delta(\lambda) + \text{negl}(\lambda)$.
- By Lemma 3.11, if Π is somewhere computationally sound, then Π_{WI} is adaptively computationally sound.
- All V_{WI} does is run V on an input provided by P_{WI} and accept iff it accepts. Gen_{WI} similarly only samples from $[t(\lambda)]$ and runs TGen once. P_{WI} also simply runs P on an input and witnesses, and in addition computes samples from D_λ and $[t(\lambda)]$, which can be done in non-uniform polynomial time since \mathcal{D} is efficiently sampleable.
- In addition to the proof from Π , the proof in Π_{WI} consists only of the $t(\lambda)$ instances of length λ sampled by P_{WI} .

The above arguments prove the respective properties of the protocol promised by the theorem. \square

Distributional NISWIA with uniform prover. If we weaken the hiding requirement to *distributional* SWI (Definition 2.16), then it is possible to avoid the non-uniformity of prover in Theorem 3.10.

Theorem 3.13 (Distributional NISWIA from NIBARG). *Consider an NP relation \mathcal{R} and let $\mathcal{D}' = (D'_\lambda)_{\lambda \in \mathbb{N}}$ be a distribution where D'_λ is supported on triples (x, w^0, w^1) such that $(x, w^0), (x, w^1) \in \mathcal{R}_\lambda$. Consider the distribution $\mathcal{D} = (D_\lambda)_{\lambda \in \mathbb{N}}$, where D_λ is supported on pairs $(x, w) \in \mathcal{R}_\lambda$ and defined via the following sampling procedure. On input 1^λ ,*

1. Sample $(x, w^0, w^1) \leftarrow D'_\lambda$ and $b \leftarrow \{0, 1\}$
2. Output (x, w^b) .

Let $\Pi = (\text{Gen}, \text{TGen}, \text{P}, \text{V})$ be a non-interactive batch protocol for \mathcal{R} that, when run on some polynomial $t = t(\lambda)$ instances of size λ , has compression rate $\rho = \rho(\lambda) < 1$. Then, the non-interactive protocol $\Pi_{\text{WI}} = \langle \text{P}_{\text{WI}} \rightarrow \text{V}_{\text{WI}} \rangle$ from Fig. 2 instantiated with Π and \mathcal{D} has the following properties (on instances of size λ):

- If \mathcal{D}' is uniformly efficiently-sampleable, then the prover P_{WI} is uniform.
- Π_{WI} is distributional SWI for \mathcal{D}' with error $O(\sqrt{\rho})$.

- If Π is CRS-indistinguishable, then Π_{WI} has completeness error negligibly close to that of Π run on t instances.
- If Π is somewhere computationally sound, then Π_{WI} is adaptively computationally sound.
- The length of the proof in Π_{WI} is that in Π plus an additional $\lambda \cdot t$ bits. The length of the CRS is the same.

Proof Sketch. We focus on the first two claims since the rest of the claims can be proved as in Theorem 3.10. From the description of the sampling procedure, it is clear that if \mathcal{D} is uniformly efficiently samplable then \mathcal{D}' also is, and as a result the prover P_{WI} is uniform. To prove the second claim, observe that

$$\begin{aligned}
\mathbb{E}_{(x,w^0,w^1) \leftarrow D'_\lambda} [\text{SD}(crs_{\text{WI}}, \pi^0), (crs_{\text{WI}}, \pi^1)] &= \mathbb{E}_{(x,w^0,w^1) \leftarrow D'_\lambda} \left[\mathbb{E}_{j \leftarrow [t], crs} [\text{SD}(\pi^0, \pi^1)] \right] \\
&= \mathbb{E}_{(x,w^0,w^1) \leftarrow D'_\lambda} \left[\mathbb{E}_{\substack{(x_1, w_1) \leftarrow D_\lambda \\ (x_t, w_t) \leftarrow D_\lambda}} \left[\mathbb{E}_{j \leftarrow [t], crs} [\text{SD}(\pi^0, \pi^1)] \right] \right] \\
&= \mathbb{E}_{(x,w^0,w^1) \leftarrow D'_\lambda} \left[\mathbb{E}_{\substack{(x_1, w_1^0, w_1^1) \leftarrow D'_\lambda \\ (x_t, w_t^0, w_t^1) \leftarrow D'_\lambda}} \left[\mathbb{E}_{j \leftarrow [t], crs} [\text{SD}(\pi^0, \pi^1)] \right] \right] \\
&\leq O(\sqrt{\rho}),
\end{aligned}$$

where

- $crs_{\text{WI}} = (j, crs) \leftarrow \text{Gen}_{\text{WI}}(1^\lambda)$ and $\pi^b \leftarrow \text{P}_{\text{WI}}(crs_{\text{WI}}, x, w^b)$ in the first equality;
- π^b is generated using P as described within P_{WI} in Fig. 2 in the second and third equalities;
- the third equality follows by the definition of \mathcal{D} ; and
- the inequality follows by Proposition 3.8 with $f(\mathbf{x}, \mathbf{w}, r) := \text{P}(r, \mathbf{x}, \mathbf{w})$ and column player's mixed strategy set to \mathcal{D}' .

□

3.3 Corollaries

In this section we state some of the known results on transforming HVSWI protocols into SWI and SZK protocols against malicious verifiers. Starting from a public-coin HVSWI proof, Corollary 3.14 gives an SWI proof against malicious verifiers without any additional assumptions. Even if the original HVSWI proof is not public coin, we can use it to obtain an SWI proof under computational assumptions. This transformation is given by Corollary 3.15. Moving on to the setting of computational soundness, assuming OWFs exist, Corollary 3.15 shows that any HVSWI argument can be transformed into an SWI argument. Under the same assumption, Corollary 3.16 gives a transformation from an SWI argument to an SZK argument. Finally, Corollary 3.17 gives a similar transformation from an SWI argument to an SZK argument for non-interactive protocols; Corollary 3.18 is its distributional counterpart.

Corollary 3.14. *If there exists a public-coin HVSWI proof Π for an NP relation \mathcal{R} then there exists a public-coin SWI proof Π_M for \mathcal{R} with the following properties:*

- Π_M has negligible completeness and soundness error.
- Π_M has WI error $\text{poly}(\lambda) \cdot \varepsilon + 2^{-\Theta(\lambda)}$ where λ is the instance length and ε is WI error of Π .

- If Π has d rounds then Π_M has $2d$ rounds.
- If the honest prover in Π is non-uniform then so is the honest prover in Π_M .

Proof Sketch. The proof is based on [Vad99, Theorem 6.3.5] which gives a transformation from any public-coin HVZK proof to a public-coin SZK proof with the following properties:

- Π_M has negligible completeness error $2^{-\lambda}$ and soundness error $1/\lambda$.
- Π_M has ZK error $\text{poly}(\lambda) \cdot \varepsilon + 2^{-\Theta(\lambda)}$ where λ is the instance length and ε is ZK error of Π .
- If Π has d rounds then Π_M has $2d$ rounds.
- If the honest prover in Π is non-uniform then so is the honest prover in Π_M .

We observe that the same transformation also transforms any HVSWI proof to an SWI proof with the same properties. To see that, recall that a protocol $\langle P \leftrightarrow V \rangle$ is SWI with error ε if and only if there exists an (unbounded) simulator S such that for every polynomial-size circuit family $V^* = (V_\lambda^*)_{\lambda \in \mathbb{N}}$,

$$\{\langle P(w) \leftrightarrow V_\lambda^* \rangle(x)\}_{\substack{(x,w) \in \mathcal{R} \\ |x|=\lambda}} \approx_\varepsilon \{S(x, V_\lambda^*)\}_{\substack{(x,w) \in \mathcal{R} \\ |x|=\lambda}} .$$

Similarly, the protocol is HVSWI if and only if the above is guaranteed for the honest verifier V . The proof of [Vad99, Theorem 6.3.5] uses the polynomial time simulator of the original SZK proof to construct a polynomial-time simulator for the new proof system. By inspecting the transformation and its analysis, we conclude that if the original protocol has an unbounded simulator instead of a polynomial-time one, the simulator constructed for the new proof system is also unbounded with the same error as in the efficient case.

The SWI proof resulting from the [Vad99] transformation has a non-negligible soundness error. To get an SWI proof with the claimed properties we can repeat the resulting SWI proof in parallel $\text{poly}(\lambda)$ times. This reduces the soundness error to negligible and only increases the WI error by a factor of $\text{poly}(\lambda)$. \square

Corollary 3.15. *Let Π be an HVSWI protocol for an NP relation \mathcal{R} with d rounds.*

- Assume there exists a statistically hiding commitment with d^* rounds. If Π is statistically sound then there exists an SWI proof for \mathcal{R} with $O(d \cdot d^*)$ rounds and soundness error that is negligibly close to that of Π .
- Assuming OWFs exist, if Π is computationally sound then there exists an SWI argument for \mathcal{R} with $O(d)$ rounds.

The WI error of the new protocol is negligibly close to that of Π . If the honest prover in Π is non-uniform then so is the honest prover in the new protocol.

Statistically-hiding commitment can be constructed in two rounds from CRHFs [DPP97, FS90, HM96], in a constant number of rounds from multi-collision resistant hash functions [BDRV18, KNY18] or distributional CRHFs [BHKY19], and in $O(\lambda)$ rounds from OWFs [HNO+09].

The proof is based on the compiler of [GMW86]. We start with the case of computational soundness and then explain how to modify the protocol to obtain statistical soundness. The verifier starts by committing to a random string r_V using a statistically-binding commitment and the prover responds with a random string r_P . Then the prover and verifier execute the HVSWI protocol where the verifier uses the randomness $r_V \oplus r_P$. After each message, the verifier proves using a computational ZK argument that the message was generated correctly. The SWI argument $\langle P_M \leftrightarrow V_M \rangle$ is described in Fig. 3. The construction uses a two-message statistically binding commitment and a constant-round computational ZK argument, both of which can be constructed from OWFs [Nao91, FS90]. Next, we sketch the proof of soundness and SZK.

Let COM be a two-message statistically binding commitment. Let $\langle \text{P}_{\text{CZK}} \leftrightarrow \text{V}_{\text{CZK}} \rangle$ be a constant-round computational ZK argument. Let $\langle \text{P}_{\text{WI}} \leftrightarrow \text{V}_{\text{WI}} \rangle$ be a d -round HVSWI argument where verifier's randomness is of length ℓ . Assume WLOG that d is even.

The SWI argument $\langle \text{P}_M \leftrightarrow \text{V}_M \rangle$ is as follows. The prover and verifier are given an instance $x \in \{0, 1\}^\lambda$. The prover is also given a witness $w \in \{0, 1\}^*$.

Proof Sketch. 1. P_M samples a first message k for COM and sends it to V_M

2. V_M samples $r_V \leftarrow \{0, 1\}^\ell$ and a commitment $c \leftarrow \text{COM}_k(r_V)$, and sends c to P_M

3. P_M samples $r_P \leftarrow \{0, 1\}^\ell$ sends it to V_M

4. For $i = 1, \dots, d$:

(a) V_M computes the next message of V_{WI} using randomness $r = r_V \oplus r_P$:

$$\alpha_i \leftarrow \text{V}_{\text{WI}}(x, \alpha_1, \beta_1, \dots, \alpha_{i-1}, \beta_{i-1}; r) ,$$

and sends α_i to P_M

(b) V_M and P_M execute the protocol $\langle \text{P}_{\text{CZK}} \leftrightarrow \text{V}_{\text{CZK}} \rangle$ where V_M proves to P_M that there exist strings \tilde{r}_V and σ such that:

$$c = \text{COM}_k(\tilde{r}_V; \sigma) \wedge \alpha_i = \text{V}_{\text{WI}}(x, \alpha_1, \beta_1, \dots, \alpha_{i-1}, \beta_{i-1}; \tilde{r}_V \oplus r_P)$$

(c) If V_{CZK} rejects then P_M aborts. Otherwise, P_M computes the next message of P_{WI} :

$$\beta_i \leftarrow \text{P}_{\text{WI}}(x, w, \alpha_1, \beta_1, \dots, \alpha_i) ,$$

and sends β_i to V_M

Figure 3: Malicious-verifier SWI argument from an HVSWI argument and OWFs

Computational soundness. Assume towards contradiction that there exists a polynomial-size cheating prover P^* that can prove a false statement with non-negligible probability ϵ . We use P^* to break the computational soundness of the HVSWI argument $\langle P_{WI} \rightleftharpoons V_{WI} \rangle$. First we consider a hybrid experiment where we emulate an execution P^* with the verifier V_M , but each execution of the ZK argument $\langle P_{CZK} \rightleftharpoons V_{CZK} \rangle$ is simulated. By the zero-knowledge property of the ZK argument, P^* will continue to produce accepting proofs with probability that is negligibly close to ϵ . In the next hybrid, we modify the value in the initial commitment sent by V_M from r_V to 0^ℓ . By the computational hiding property of the commitment, P^* will continue to produce accepting proofs with probability that is negligibly close to ϵ . Now we can break the soundness of the HVSWI argument $\langle P_{WI} \rightleftharpoons V_{WI} \rangle$ by emulating this final hybrid experiment and forwarding the messages of the external verifier V_{WI} to P^* instead of computing them using the randomness $r = r_V \oplus r_P$. Since the string r is uniformly distributed, we convince the external verifier of a false statement with probability that is negligibly close to ϵ .

SWI. Fix any polynomial-size cheating verifier V^* and statement-witness pairs $(x, w_0), (x, w_1) \in \mathcal{R}$. Let ϵ denote the distance between the views $\text{View}_0 = \langle P_M(w_0) \rightleftharpoons V^* \rangle(x)$ and $\text{View}_1 = \langle P_M(w_1) \rightleftharpoons V^* \rangle(x)$. Since the commitment COM is statistically binding, we can fix the first commitment message k sampled by P_M such that COM_k is perfectly binding and the distance between View_0 and View_1 remains negligibly close to ϵ . Let \tilde{r}_V be the string that V^* commits to in its first message.

For $b \in \{0, 1\}$ we consider the view of the honest verifier V_{WI} in the interaction of $\langle P_{WI}(w_b) \rightleftharpoons V_{WI} \rangle(x)$ which consist of the verifier's randomness r and the prover's messages $(\beta_1, \dots, \beta_d)$. We argue that given this view we can efficiently sample from a distribution that is negligibly close to View_b (with the first commitment message fixed to k). Therefore, it follows that ϵ must be negligibly close to the SWI error of the HVSWI argument.

Given the view $r, (\beta_1, \dots, \beta_d)$ we sample from a distribution close to View_b as follows. We emulate the execution of V^* , setting the first prover message to k and the second prover message to $r_P = \tilde{r}_V \oplus r$. Since r is uniform, r_P is distributed exactly as in View_b . In every one of the remaining d rounds, starting from $i = 1$ to d we interact with V^* emulating the verifier of the ZK argument $\langle P_{CZK} \rightleftharpoons V_{CZK} \rangle$. If the ZK argument is accepted then we set the next prover message to β_i , otherwise the prover aborts.

Let E be the event that the verifier V^* proves a true statement in each of the accepting executions of the ZK argument $\langle P_{CZK} \rightleftharpoons V_{CZK} \rangle$. Conditioned on E , the view sampled above is distributed exactly the same as View_b . By the computational soundness property of the ZK argument, E occurs with all but negligible probability. Therefore, the sampled view is negligibly close to View_b .

An SWI proof. If the original SWI protocol $\langle P_{WI} \rightleftharpoons V_{WI} \rangle$ has statistical soundness we can modify the protocol $\langle P_M \rightleftharpoons V_M \rangle$ described in Fig. 3 and obtain an SWI proof. We make the following modifications:

- We replace the two-message statistically-binding commitment with a statistically-hiding commitment.
- After the verifier sends the commitment c and before the prover sends r_P , have the verifier prove that it knows an opening of c using a SZK argument of knowledge where SZK holds even against an unbounded malicious verifier.¹³ (We describe how this SZK argument of knowledge is constructed below.)
- Replace each invocation of the computational ZK argument with a SZK argument of knowledge against an unbounded malicious verifier.

Since the verifier's commitment and ZK arguments are all statistical, we can show statistical soundness following the same argument as in the computational case. To prove SWI, modify the above proof as follows. Since the commitment c is statistically-hiding, the string \tilde{r}_V that V^* commits to is not well defined. Instead, we invoke the knowledge extractor of the SZK argument of knowledge and extract an opening to a string \tilde{r}_V . To prove SWI we need to show that, with all but negligible probability, all the messages β_1, \dots, β_d are computed according to the strategy of the honest verifier in the HVSWI argument V_{WI} using the randomness

¹³This is in contrast to the weaker notion of SZK in Definition 2.6 that only considers polynomial-size malicious verifiers.

$\tilde{r}_V \oplus r_P$. If this is not the case for some β_i , we can use the knowledge extractor of the SZK argument of knowledge and obtain an opening of the commitment c to a value other than r_V with non-negligible probability, contradicting the computational binding property of the commitment.

Using a statistically-hiding commitment with d^* rounds, a SZK argument of knowledge against an unbounded malicious verifier with $O(d^*)$ rounds can be constructed following the outline of [FLS90, GK96]: Start from an SWI argument of knowledge against an unbounded verifier in $O(d^*)$ rounds. Such a protocol can be obtained by taking the parallel repetition of the ZK protocol of [Blu81, GMW86] and instantiating the commitment scheme with the statistically hiding commitment. Next, the SWI argument of knowledge is transformed into an SZK argument of knowledge using the compiler of [FLS90]. In more details, the verifier starts by committing to trapdoor statement using a statistically hiding commitment and proving that the committed statement is true using a computational ZK proof of knowledge. Then the prover uses the SWI argument of knowledge to prove that either the original statement or the committed trapdoor statement is true. The required computational ZK proof of knowledge can be constructed by combining the computational ZK proof of [GK96] (instantiated with the statistically-hiding commitment) with a computational WI proof of knowledge (given by parallel repetition of the ZK protocol of [Blu81, GMW86], instantiated with a statistically binding commitment) via the [FLS90] compiler. \square

Corollary 3.16. *Assuming one-way functions exist, if there exists an SWI argument Π for an NP-complete relation \mathcal{R} with d rounds, then there exists an SZK argument Π_{ZK} for \mathcal{R} with $O(d)$ rounds and ZK error that is negligibly close to the WI error of Π . If the honest prover in Π is non-uniform then so is the honest prover in Π_{ZK} .*

Proof Sketch. The proof is based on the compiler of [FLS90]. The verifier starts by sending a random image y of a length-doubling PRG and proving that it knows a corresponding preimage using a computational ZK argument of knowledge. Then, the prover and verifier execute the SWI protocol proving that either the original statement is true *or* that y is in the image of the PRG. In other words, the prover and verifier execute SWI protocol to prove that $(x, y) \in \mathcal{R}_{\text{OR}}$, which is the relation obtained by “ORing” \mathcal{R} and the relation defined by the PRG: i.e,

$$\mathcal{R}_{\text{OR}} := \{((x, y), z) : (x, z) \in \mathcal{R} \vee \text{PRG}(z) = y\} \quad (2)$$

The SZK argument $\langle P_{\text{ZK}} \rightleftharpoons V_{\text{ZK}} \rangle$ is described in Fig. 4. The construction uses a PRG and a constant-round computational ZK argument of knowledge, both of which can be constructed from OWFs [HILL99, FS90]. Next, we sketch the proof of soundness and SZK.

Soundness. Assume towards contradiction that there exists a polynomial-size cheating prover P^* that can prove a false statement with non-negligible probability ϵ . We use P^* to break the computational soundness of the SWI argument $\langle P_{\text{WI}} \rightleftharpoons V_{\text{WI}} \rangle$. First we consider a hybrid experiment where we emulate an execution P^* with the verifier V_{ZK} , but the execution of the computational ZK argument $\langle P_{\text{CZK}} \rightleftharpoons V_{\text{CZK}} \rangle$ is simulated. By the zero-knowledge property of the ZK argument, P^* will continue to produce accepting proofs with probability that is negligibly close to ϵ . In the next hybrid, we sample a uniform $y \leftarrow \{0, 1\}^{2\lambda}$ instead of sampling y as a random image of the PRG. By the pseudorandomness of the generator, P^* will continue to produce accepting proofs with probability that is negligibly close to ϵ . Now, the statement for the SWI argument $\langle P_{\text{WI}} \rightleftharpoons V_{\text{WI}} \rangle$ is false with probability $1 - 2^{-\lambda}$. Therefore, we can break the soundness of $\langle P_{\text{WI}} \rightleftharpoons V_{\text{WI}} \rangle$ with probability that is negligibly close to ϵ .

SZK. We describe a simulator S . The simulator is given an instance $x \in \mathcal{L}(\mathcal{R})$ and the description of a cheating verifier V^* . S emulates an interaction with V^* . If the verifier V_{CZK} rejects in the execution of the computational ZK argument of knowledge $\langle P_{\text{CZK}} \rightleftharpoons V_{\text{CZK}} \rangle$ then S outputs the transcript of the interaction with V^* up to that point. Otherwise, S invokes the knowledge extractor of $\langle P_{\text{CZK}} \rightleftharpoons V_{\text{CZK}} \rangle$ on the description of the residual verifier V^* after sending its first message, right before the execution of $\langle P_{\text{CZK}} \rightleftharpoons V_{\text{CZK}} \rangle$. If the extractor fails to output a string r such that $y = \text{PRG}(r)$ then S aborts. Otherwise, S continues to emulate

Let $\text{PRG} : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{2\lambda}$ be a length-doubling PRG. Let $\langle \text{P}_{\text{CZK}} \rightleftharpoons \text{V}_{\text{CZK}} \rangle$ be a constant-round computational ZK argument of knowledge. Let $\langle \text{P}_{\text{WI}} \rightleftharpoons \text{V}_{\text{WI}} \rangle$ be an SWI argument. The SZK argument $\langle \text{P}_{\text{ZK}} \rightleftharpoons \text{V}_{\text{ZK}} \rangle$ is as follows. The prover and verifier are given an instance $x \in \{0, 1\}^\lambda$. The prover is also given a witness $w \in \{0, 1\}^*$.

1. V_{ZK} samples a string $r \leftarrow \{0, 1\}^\lambda$ and sends $y = \text{PRG}(r)$ to P_{ZK}
2. V_{ZK} proves to P using $\langle \text{P}_{\text{CZK}} \rightleftharpoons \text{V}_{\text{CZK}} \rangle$ that there exists a string \tilde{r} such that $y = \text{PRG}(\tilde{r})$
3. If V_{CZK} rejects then P_{ZK} aborts. Otherwise, P_{ZK} proves to V_{ZK} using $\langle \text{P}_{\text{WI}} \rightleftharpoons \text{V}_{\text{WI}} \rangle$ that there exist strings \tilde{w} and \tilde{r} such that:

$$(x, \tilde{w}) \in \mathcal{R} \vee y = \text{PRG}(\tilde{r}) ,$$

using the witness $\tilde{w} = w$ and $\tilde{r} = \perp$

Figure 4: SZK argument from an SWI argument and PRG

an interaction with V^* by executing the honest prover of the SWI argument $\langle \text{P}_{\text{WI}} \rightleftharpoons \text{V}_{\text{WI}} \rangle$ using the witness $\tilde{w} = \perp$ and $\tilde{r} = r$. Finally, S outputs the transcript of the entire interaction with V^* . Fix any $(x, w) \in \mathcal{R}$ and a polynomial-size cheating verifier V^* . Since the extractor runs in time that is inverse polynomial to the probability that the proof given by V^* is accepted, it follows that S runs in expected polynomial time. By the knowledge soundness property of the computational ZK argument of knowledge the extractor fails to find a witness and the simulation aborts only with negligible probability. Conditioned on the fact that S does not abort, the only difference between the simulated view generated by $\text{S}(x, \text{V}^*)$ and the real view $\langle \text{P}_{\text{ZK}}(w) \rightleftharpoons \text{V}^* \rangle(x)$ is the witness used by the honest prover in the execution of $\langle \text{P}_{\text{WI}} \rightleftharpoons \text{V}_{\text{WI}} \rangle$. Therefore, the ZK error of the SWI argument $\langle \text{P}_{\text{ZK}} \rightleftharpoons \text{V}_{\text{ZK}} \rangle$ is negligibly close to the WI error of the SWI argument $\langle \text{P}_{\text{WI}} \rightleftharpoons \text{V}_{\text{WI}} \rangle$. \square

Non-Interactive SZK. Finally, we describe the transformation for non-interactive protocols. To be specific, building on [FLS90], we transform a (distributional) NISWIA to a dual-mode (distributional) NISZKA.

Corollary 3.17. *Assuming one-way functions exist, if there exists an adaptively-sound non-interactive SWI argument Π for an NP-complete relation \mathcal{R} in the CRS model, then there exists a dual-mode non-interactive SZK argument Π_{ZK} for \mathcal{R} such that:*

- The ZK error of Π_{ZK} is the same as the WI error of Π .
- If the honest prover in Π is non-uniform then so is the honest prover in Π_{ZK} (otherwise Π_{ZK} is uniform).

We skip the proof of this corollary and prove below the more general corollary.

Corollary 3.18. *Let PRG be a length-doubling PRG. Consider an NP relation \mathcal{R} and let $\mathcal{D} = (D_\lambda)_{\lambda \in \mathbb{N}}$ be any uniformly efficiently-sampleable distribution such that D_λ is supported on pairs $(x, w) \in \mathcal{R}_\lambda$. Moreover, consider the relation \mathcal{R}_{OR} defined in Eq. (2), and the (uniformly) efficiently-sampleable distribution $\mathcal{D}_{\text{OR}} = (D_{\text{OR}, \lambda})_{\lambda \in \mathbb{N}}$ on \mathcal{R}_{OR} , where $D_{\text{OR}, \lambda}$ is defined via the following sampling procedure. On input 1^λ ,*

1. Sample $(x, w) \leftarrow D_\lambda$,
2. Sample $s \leftarrow \{0, 1\}^\lambda$ and set $y := \text{PRG}(s)$,
3. Output $((x, y), w, s)$.

Let PRG be a length-doubling PRG. Let $\Pi_{\text{WI}} := \langle P_{\text{WI}} \rightarrow V_{\text{WI}} \rangle$ be a NISWI argument for any NP-complete relation. The NISZK argument $\langle P_{\text{ZK}} \rightarrow V_{\text{ZK}} \rangle$ for any NP relation \mathcal{R} is as follows.

$crs_{\text{ZK}} \leftarrow \text{Gen}_{\text{ZK}}(1^\lambda)$

1. Sample a random element $y \in \{0, 1\}^{2\lambda}$ from the *image* of PRG.
2. Generate a CRS of Π_{WI} : $crs_{\text{WI}} \leftarrow \text{Gen}_{\text{WI}}(1^\lambda)$.
3. Output $crs_{\text{ZK}} := (crs_{\text{WI}}, y)$ as the CRS.

$crs_{\text{ZK}} \leftarrow \text{AGen}_{\text{ZK}}(1^\lambda)$

1. Sample a random element $y \in \{0, 1\}^{2\lambda}$ from the *codomain* of PRG.
2. Generate a CRS of Π_{WI} : $crs_{\text{WI}} \leftarrow \text{Gen}_{\text{WI}}(1^\lambda)$.
3. Output $crs_{\text{ZK}} := (crs_{\text{WI}}, y)$ as the CRS.

$\pi \leftarrow P_{\text{ZK}}(crs_{\text{ZK}}, x, w)$

1. Use witness w to generate a proof that $(x, y) \in \mathcal{R}_{\text{OR}}$ (see Eq. (2)): $\pi \leftarrow P_{\text{WI}}(crs_{\text{WI}}, (x, y), w)$.
2. Output π .

$b := V_{\text{ZK}}(crs_{\text{ZK}}, x, \pi)$

1. Accept ($b := 1$) if and only if $V_{\text{WI}}(crs_{\text{ZK}}, (x, y), \pi)$ accepts.

Figure 5: Dual-mode NISZKA from an NISWIA and PRG.

If Π is an adaptively-sound distributional NISWIA for \mathcal{D}_{OR} then Π_{ZK} , defined in Fig. 5, is dual-mode distributional NISZKA for \mathcal{D} such that

- The ZK error of Π_{ZK} is the same as the WI error of Π .
- If the honest prover in Π is uniform then so is the honest prover in Π_{ZK} .

Proof Sketch. Completeness follows readily; below we argue adaptive soundness and statistical ZK.

Adaptive soundness. Note that the y sampled as part of (crs_{WI}, y) in AGen lies outside the image of PRG with overwhelming probability. As a result $(x, y) \notin \mathcal{R}_{\text{OR}}$ implies that $x \notin \mathcal{R}$, and we end up reducing to adaptive soundness of Π .

Distributional statistical ZK. Recall the definitions of the distributions \mathcal{D} and \mathcal{D}_{OR} . With the simulator described in Fig. 6, we have

$$\begin{aligned} \mathbb{E}_{(x,w) \leftarrow D_\lambda} [\text{SD}((crs_{\text{ZK}}, \pi_{\text{ZK}}), S(x))] &= \mathbb{E}_{(x,w) \leftarrow D_\lambda} [\text{SD}((crs_{\text{WI}}, y, \pi_w), (crs_{\text{WI}}, y, \pi_s))] \leq \\ &= \mathbb{E}_{(x,w) \leftarrow D_\lambda} [\text{SD}((crs_{\text{WI}}, y, s, \pi_w), (crs_{\text{WI}}, y, s, \pi_s))] = \\ &= \mathbb{E}_{((x,y),w,s) \leftarrow D_{\text{OR},\lambda}} [\text{SD}((crs_{\text{WI}}, \pi_w), (crs_{\text{WI}}, \pi_s))] \quad , \end{aligned}$$

Let PRG and Π_{WI} be as in Fig. 5. The simulator for Π_{ZK} is described below.

$(\text{crs}_{\text{ZK}}, \pi) \leftarrow \text{S}_{\text{ZK}}(x)$

1. Sample a random image as follows: $y := \text{PRG}(s)$, where $s \leftarrow \{0, 1\}^\lambda$ is the trapdoor information.
2. Generate a CRS of Π_{WI} : $\text{crs}_{\text{WI}} \leftarrow \text{Gen}_{\text{WI}}(1^\lambda)$.
3. Use the trapdoor information to generate proof that $(x, y) \in \mathcal{R}_{\text{OR}}$: $\pi \leftarrow \text{P}_{\text{WI}}(\text{crs}_{\text{WI}}, x, s)$.
4. Output $(\text{crs}_{\text{ZK}}, \pi)$, where $\text{crs}_{\text{ZK}} := (\text{crs}_{\text{WI}}, y)$.

Figure 6: Simulator for NISZKA from Fig. 5.

where in the first equation $\text{crs}_{\text{WI}} \leftarrow \text{Gen}_{\text{WI}}(1^\lambda)$, $s \leftarrow \{0, 1\}^\lambda$, $y = \text{PRG}(s)$, $\pi_w \leftarrow \text{P}_{\text{WI}}(\text{crs}_{\text{WI}}, (x, y), w)$ and $\pi_s \leftarrow \text{P}_{\text{WI}}(\text{crs}_{\text{WI}}, (x, y), s)$. □

4 Dual-Mode NISZKA from Non-Interactive Batch Arguments

The WI and ZK errors for the protocols obtained in the previous section were only inverse-polynomially small. Moreover, the prover in those protocols was non-uniform. In this section, we remedy both these issues for the non-interactive setting. First, in Section 4.1, we construct (worst-case) dual-mode NISZKA from dual mode *distributional* NISZKA assuming *lossy public-key encryption* (LPKE). Since the prover in our dual-mode distributional NISZKA is uniform (see Corollary 3.18), the prover in the resulting dual-mode NISZKA is also uniform. Then, in Section 4.2 we boost the ZK error of this protocol using LPKE and semi-honest *multi-party computation* (MPC) to obtain dual-mode NISZKA with all the desired properties.

We further show in Appendix B that LPKE can be generically constructed (Theorem B.3) from NIBARGs that are *honestly* somewhere extractable.

Remark 9 (On the Restriction to the Non-interactive Setting). *We note that both steps described above crucially rely on lossy encryption and are hence only interesting in the non-interactive setting. Indeed, since lossy encryption implies two-message statistically-hiding commitments, in the interactive setting they alone would allow the construction of constant-round statistical zero knowledge arguments with a uniform prover, negligible ZK error (and no reliance on BARGs).*

4.1 Dual-Mode NISZKA from Dual-Mode Distributional NISZKA

In this section, we construct dual-mode NISZKA for circuit satisfiability given dual-mode *distributional* NISZKA and LPKE. First, we recall the definition of LPKE.

Definition 4.1 (LPKE). *A lossy public-key encryption scheme Λ with message-space \mathcal{M} and ciphertext-space \mathcal{C} is a tuple of polynomial-time algorithms $(\text{KGen}, \text{LGen}, \text{E}, \text{D})$ with following syntax:*

- $(pk, sk) \leftarrow \text{KGen}(1^\lambda)$. *The randomised binding-key-generation algorithm, on input a security parameter $\lambda \in \mathbb{N}$, outputs a public-private key-pair (pk, sk) . We refer to a public key generated by KGen as a binding key.*
- $pk \leftarrow \text{LGen}(1^\lambda)$. *The randomised lossy-key-generation algorithm, on input a security parameter $\lambda \in \mathbb{N}$, outputs a public key pk , which we refer to as a lossy key.*

- $c \leftarrow \mathbf{E}(pk, m)$. The randomised encryption algorithm takes as input a public key pk and a message $m \in \mathcal{M}$, and outputs a ciphertext $c \in \mathcal{C}$.
- $m := \mathbf{D}(sk, c)$. The deterministic decryption algorithm takes a secret key sk and a ciphertext $c \in \mathcal{C}$ as input and outputs a message $m \in \mathcal{M}$.

We require the following properties from Λ :

1. Binding public keys are almost-all-keys perfectly correct [DNR04]. With overwhelming probability over the choice of binding keys, perfect correctness of decryption must hold. That is, with overwhelming probability over $(pk, sk) \leftarrow \mathbf{KGen}(1^\lambda)$, for every $m \in \mathcal{M}$

$$\Pr_{c \leftarrow \mathbf{E}(pk, m)} [\mathbf{D}(sk, c) \neq m] = 0.$$

2. Lossy keys are statistically hiding. For a random lossy key, the distribution of ciphertexts of any two messages must be statistically close. To be specific, we say that the lossy keys are δ -statistically-hiding if for $(pk, sk) \leftarrow \mathbf{LGen}(1^\lambda)$ and $m_0, m_1 \in \mathcal{M}$:

$$\mathbf{SD}(\mathbf{E}(pk, m_0), \mathbf{E}(pk, m_1)) \leq \delta(\lambda).$$

3. Mode indistinguishability. We require that the binding and lossy keys are computationally indistinguishable. That is, for every polynomial-size circuit family $\mathbf{A} = (\mathbf{A}_\lambda)_{\lambda \in \mathbb{N}}$, there is a negligible function μ , such that for all $\lambda \in \mathbb{N}$:

$$\left| \Pr_{(pk, sk) \leftarrow \mathbf{KGen}(1^\lambda)} [1 \leftarrow \mathbf{A}_\lambda(pk)] - \Pr_{pk \leftarrow \mathbf{LGen}(1^\lambda)} [1 \leftarrow \mathbf{A}_\lambda(pk)] \right| \leq \mu(\lambda).$$

Remark 10 (Amplifying hiding). Using standard amplification (XOR Lemma), δ can be made as small as $2^{-\text{poly}(\lambda)}$ for any poly at the cost of polynomially increasing the size of commitments (c.f. [LM20]).

Remark 11 (Perfectly-binding commitment from LPKE). With overwhelming probability over the choice of a random binding key pk , the encryption algorithm $\mathbf{E}(pk, \cdot)$ acts as a perfectly-binding (non-interactive) commitment, with the random coins used for encryption serving as opening (see [LS19]).

Local consistency. Before presenting our construction, we define an NP relation \mathcal{R}_g that captures *local consistency* of commitments to gates and then describe a distribution $\mathcal{D}_g = (D_{g, \lambda})_{\lambda \in \mathbb{N}}$ supported on \mathcal{R}_g . Looking ahead, for every Boolean gate g , we will require the underlying NISZKA to be *distributionally SZK* for \mathcal{D}_g . The following notation will also be useful.

Notation 1 (Notation for circuits.). Let C be a Boolean circuit with wires labeled by $[n]$. We use the same label for wires fanning out from the same gate or input bit. We say that an assignment $\mathbf{w} = (w_1, \dots, w_n) \in \{0, 1\}^n$ satisfies a gate g with input wires i, j and output wire k if $g(w_i, w_j) = w_k$. An assignment \mathbf{w} satisfies C if it satisfies all of its gates and the value of the output wire is 1.

For a Boolean gate g , a four-tuple $(pk, c_1, c_2, c_3) \in \mathcal{L}(\mathcal{R}_g)$ if the ciphertexts (c_1, c_2, c_3) encrypt values consistent with g . That is,

1. For $i \in [3]$, there exists openings $(w_i, r_i) \in \{0, 1\} \times \{0, 1\}^{\text{poly}(n)}$ such that $\mathbf{E}(pk, w_i; r_i) = c_i$; and
2. The wire values (w_1, w_2, w_3) are consistent with g , i.e., $g(w_1, w_2) = w_3$.

The distribution $D_{g, \lambda}$ is defined via the following (uniform) sampling procedure. On input 1^λ ,

1. Sample a random lossy key: $pk \leftarrow \mathbf{LGen}(1^\lambda)$
2. Sample a random assignment (w_1, w_2, w_3) consistent with g .

3. For $i \in [3]$: sample $r_i \in \{0, 1\}^{\text{poly}(n)}$ and compute commitment $c_i := E(pk, w_i; r_i)$ to w_i .
4. Output (pk, c_1, c_2, c_3) as the statement and $((w_1, w_2, w_3), (r_1, r_2, r_3))$ as the witness.

Our protocol Π_{ZK} is described formally in Fig. 7. The construction is similar in spirit to [GOS12] in that the prover commits to the wire values and then, for each gate in the circuit, proves *local consistency* of the commitments to that gate's wires. The main difference is that we rely on dual-mode distributional NISZKA instead of homomorphic proof commitments to prove local consistency.

Theorem 4.2 (Dual-Mode NISZKA from Dual-Mode Distributional NISZKA). *Consider the protocol $\Pi_{\text{ZK}} = \langle P_{\text{ZK}} \rightarrow V_{\text{ZK}} \rangle$ described in Fig. 7. Let*

- Π be a dual-mode distributional NISZKA for \mathcal{D}_g with ZK error at most ε ; and
- Λ be any LPKE that is δ -statistically-hiding.

Then Π_{ZK} is a dual-mode $4n(\delta + \varepsilon)$ -NISZKA for circuit satisfiability, where n denotes the number of wires in the circuit. If Π_{ZK} has negligible (resp., 0) completeness error then so does Π

Proof. Completeness follows readily. The computational indistinguishability of statistically-hiding and adaptively-sound modes of Π_{ZK} follows directly by computational indistinguishability of (i) lossy and binding keys of Λ and (ii) statistically-hiding and adaptively-sound modes of Π . We prove soundness and zero knowledge next.

Adaptive Soundness. Note that in the adaptively-sound mode of Π_{ZK} , the public key of Λ is sampled to be binding and set-ups of all instances of Π are carried out in adaptively-sound mode using AGen. Since ciphertexts generated using binding keys of Λ act as perfectly-binding commitments with overwhelming probability (see Remark 10), we are able to show that adaptive soundness of Π implies adaptive soundness of Π_{ZK} , as explained next.

Assume for contradiction that $P^* = (P_n^*)_{n \in \mathbb{N}}$, a polynomial-sized family of malicious provers, breaks Π_{ZK} 's adaptive soundness with non-negligible probability. Given a challenge CRS crs^* , the reduction generates

$$crs_{\text{ZK}}^* := (pk, (crs_1, \dots, crs_n)),$$

by sampling pk using KGen, setting $crs_{g^*} := crs^*$ for $g^* \leftarrow [n]$, and independently and randomly sampling the remaining CRSs using AGen. The reduction then runs P^* on crs_{ZK}^* , which outputs (C^*, π_{ZK}^*) such that C^* is unsatisfiable but $\pi_{\text{ZK}}^* = ((c_1^*, \dots, c_n^*), (\pi_1^*, \dots, \pi_{|C^*|}^*), r_n^*)$ is accepted by V_{ZK} . Since the commitments (c_1^*, \dots, c_n^*) are perfectly binding, they open to certain wire values $\mathbf{w}^* \in \{0, 1\}^n$. Since C^* is unsatisfiable, there must exist a gate g in C^* with input wires $i, j \in [n]$ and output wire $k \in [n]$ such that $(pk, c_i^*, c_j^*, c_k^*) \notin \mathcal{L}(\mathcal{R}_g)$, i.e., the commitments are locally *inconsistent* with respect to g (recall that $C(\mathbf{w}^*) = 1$ for a proof to be accepted). Since V_{ZK} accepts π_{ZK}^* if and only if V accepts all the underlying proofs and $g = g^*$ holds with probability $1/n$, the reduction can output $((pk, c_i^*, c_j^*, c_k^*), \pi_{g^*}^*)$ to break Π 's adaptive soundness with a $1/n$ loss. Since this is non-negligible, we contradict the assumption that Π is adaptively sound.

Statistical Zero-Knowledge. Recall from Remark 5 that when referring to statistical ZK of a dual-mode NISZKA, we are implicitly referring to its statistically-hiding mode. The simulator for Π_{ZK} is formally described in Fig. 8, and to prove that

$$\text{SD}(\langle P_{\text{ZK}}(\mathbf{w}) \rightarrow V_{\text{ZK}} \rangle(C), S_{\text{ZK}}(C)) \leq 4n(\delta + \varepsilon) \quad (3)$$

for any satisfiable circuit C with satisfying assignment \mathbf{w} , we proceed in two steps.

In the first step, we switch to an intermediate distribution H_{ZK} , described in Fig. 9, where all CRS-proof pairs (crs_g, π_g) of the underlying NISZKA Π are simulated but using *actual* commitments to \mathbf{w} . We claim

Consider circuit satisfiability on instances C with n wires (see Notation 1). Let $\Pi = \langle P \rightarrow V \rangle$ be the base non-interactive protocol for the relation \mathcal{R}_g and let $\Lambda = (\text{KGen}, \text{LGen}, \text{E}, \text{D})$ be an LPKE. The non-interactive protocol $\Pi_{\text{ZK}} = \langle P_{\text{ZK}} \rightarrow V_{\text{ZK}} \rangle$ for circuit satisfiability is described below.

$crs_{\text{ZK}} \leftarrow \text{Gen}_{\text{ZK}}(1^n)$

1. Run Λ 's *lossy* key-generation algorithm to generate a lossy key $pk \leftarrow \text{LGen}(1^n)$.
2. For $i \in [n]$: Run set-up of Π in *statistically-hiding* mode to generate CRS: $crs_i \leftarrow \text{Gen}(1^n)$.
3. Output $crs_{\text{ZK}} := (pk, (crs_1, \dots, crs_n))$ as the CRS.

$crs_{\text{ZK}} \leftarrow \text{AGen}_{\text{ZK}}(1^n)$

1. Run Λ 's *binding* key-generation algorithm to generate a binding key $pk \leftarrow \text{KGen}(1^n)$.
2. For $i \in [n]$: Run set-up of Π in *adaptively-sound* mode to generate CRS: $crs_i \leftarrow \text{AGen}(1^n)$.
3. Output $crs_{\text{ZK}} := (pk, (crs_1, \dots, crs_n))$ as the CRS.

$\pi \leftarrow P_{\text{ZK}}(crs_{\text{ZK}}, C, \mathbf{w})$

1. Use the LPKE to commit to the value of each wire. That is, for $i \in [n]$, compute $c_i := \text{E}(pk, w_i; r_i)$ for $r_i \leftarrow \{0, 1\}^{\text{poly}(n)}$. (Recall that wires fanning out from the same gate or input bit have the same label and hence a single commitment.)
2. For each gate g in C with input wires $i, j \in [n]$ and output wire $k \in [n]$, use the opening $((w_i, w_j, w_k), (r_i, r_j, r_k))$ to generate a proof π_g that $(pk, c_i, c_j, c_k) \in \mathcal{L}(\mathcal{R}_g)$. That is,

$$\pi_g \leftarrow P(crs_g, (pk, c_i, c_j, c_k), ((w_i, w_j, w_k), (r_i, r_j, r_k))).$$

3. Output $\pi := ((c_1, \dots, c_n), (\pi_1, \dots, \pi_{|C|}), r_n)$.

$b := V_{\text{ZK}}(crs_{\text{ZK}}, C, \pi)$

1. For each gate g in C with input wires $i, j \in [n]$ and output wire $k \in [n]$, use π_g to verify that $(pk, c_i, c_j, c_k) \in \mathcal{L}(\mathcal{R}_g)$. That is, run

$$b_i := V(crs_g, (pk, c_i, c_j, c_k), \pi_g).$$

2. Verify that C evaluates to 1, by ensuring that c_m opens to 1: $c_m = \text{E}(pk, 1; r_n)$.
3. Accept π (return $b := 1$) if and only if all the checks above pass.

Figure 7: Dual-mode NISZKA Π_{ZK} for circuit satisfiability.

Let C , n , $\Pi = \langle \mathcal{P} \rightarrow \mathcal{V} \rangle$ and $\Lambda = (\text{KGen}, \text{LGen}, \text{E}, \text{D})$ be as defined in Fig. 7. Moreover, let \mathcal{S} denote the simulator for Π . The simulator \mathcal{S}_{ZK} for the non-interactive protocol $\Pi_{\text{ZK}} = \langle \mathcal{P}_{\text{ZK}} \rightarrow \mathcal{V}_{\text{ZK}} \rangle$ from Fig. 7 is described below.

$(\text{crs}_{\text{ZK}}, \boldsymbol{\pi}) \leftarrow \mathcal{S}_{\text{ZK}}(C)$

1. Run Λ 's *lossy* key-generation algorithm to generate a lossy key $pk \leftarrow \text{LGen}(1^n)$.
2. Generate *dummy* commitments: for $i \in [n]$, compute $c_i := \text{E}(pk, 1; r_i)$ for $r_i \leftarrow \{0, 1\}^{\text{poly}(n)}$.
3. Simulate underlying proofs using \mathcal{S} . That is, for each gate g in C with input wires $i, j \in [n]$ and output wire $k \in [n]$ generated simulated proof that $(pk, c_i, c_j, c_k) \in \mathcal{L}(\mathcal{R}_g)$ along with a CRS:

$$(\text{crs}_g, \pi_g) \leftarrow \mathcal{S}((pk, c_i, c_j, c_k)).$$

4. Set $\text{crs}_{\text{ZK}} := (pk, (\text{crs}_1, \dots, \text{crs}_n))$ and $\boldsymbol{\pi} := ((c_1, \dots, c_n), (\pi_1, \dots, \pi_{|C|}), r_n)$.
5. Output $(\text{crs}_{\text{ZK}}, \boldsymbol{\pi})$.

Figure 8: Simulator for Π_{ZK} , the dual-mode NISZKA from Fig. 7.

Let C , n , $\Pi = \langle \mathcal{P} \rightarrow \mathcal{V} \rangle$, $\Lambda = (\text{KGen}, \text{LGen}, \text{E}, \text{D})$ and \mathcal{S} be as defined in Fig. 7. The hybrid distribution H_{ZK} for proof of Theorem 4.2 is described below.

$(\text{crs}_{\text{ZK}}, \boldsymbol{\pi}) \leftarrow H_{\text{ZK}}(C, \boldsymbol{w})$

1. Run Λ 's *lossy* key-generation algorithm to generate a lossy key $pk \leftarrow \text{LGen}(1^n)$
2. Use the LPKE to commit to the value of each wire: for $i \in [n]$, compute $c_i := \text{E}(pk, w_i; r_i)$, where $r_i \leftarrow \{0, 1\}^{\text{poly}(n)}$.
3. Simulate underlying proofs using \mathcal{S} . That is, for each gate g in C with input wires $i, j \in [n]$ and output wire $k \in [n]$ generated a simulated proof that $(pk, c_i, c_j, c_k) \in \mathcal{L}(\mathcal{R}_g)$ along with a CRS:

$$(\text{crs}_g, \pi_g) \leftarrow \mathcal{S}(pk, c_i, c_j, c_k).$$

4. Set $\text{crs}_{\text{ZK}} := (pk, (\text{crs}_1, \dots, \text{crs}_n))$ and $\boldsymbol{\pi} := ((c_1, \dots, c_n), (\pi_1, \dots, \pi_{|C|}), r_n)$
5. Output $(\text{crs}_{\text{ZK}}, \boldsymbol{\pi})$

Figure 9: Intermediate hybrid distribution in proof of Theorem 4.2.

Let g denote the gate that is switched from real to simulated in hybrid ℓ , and let (w_i, w_j, w_k) denote its wire values. The distributions *local* to g in the $\ell - 1$ -th and ℓ -th hybrids are described below.

Real local distribution $R_g(1^n, w_i, w_j, w_k)$

1. Sample a random lossy key: $pk \leftarrow \text{LGen}(1^n)$
2. For $l \in \{i, j, k\}$: $r_l \leftarrow \{0, 1\}^{\text{poly}(n)}$ and $c_l := \mathbf{E}(pk, w_l; r_l)$
3. Sample $crs_g \leftarrow \text{Gen}(1^n)$ and $\pi_g \leftarrow \mathbf{P}(crs_g, (c_i, c_j, c_k), ((w_i, w_j, w_k), (r_i, r_j, r_k)))$
4. Output $(crs_g, \pi_g, (r_i, r_j, r_k), pk)$

Simulated local distribution $S_g(1^n, w_i, w_j, w_k)$

1. Sample a random lossy key: $pk \leftarrow \text{LGen}(1^n)$
2. For $l \in \{i, j, k\}$: $r_l \leftarrow \{0, 1\}^{\text{poly}(n)}$ and $c_l := \mathbf{E}(pk, w_l; r_l)$
3. Sample $(crs_g, \pi_g) \leftarrow \mathbf{S}(c_i, c_j, c_k)$
4. Output $(crs_g, \pi_g, (r_i, r_j, r_k), pk)$

Figure 10: R_g and S_g , the local distributions in the $\ell - 1$ -th and ℓ -th hybrids, respectively.

that Π being distributionally SZK for \mathcal{D}_g with error ε implies for all (C, \mathbf{w}) as above,

$$\text{SD}(\langle \text{P}_{\text{ZK}}(\mathbf{w}) \rightarrow \text{V}_{\text{ZK}} \rangle(C), H_{\text{ZK}}(C, \mathbf{w})) \leq 4|C|\varepsilon. \quad (4)$$

To argue this, we proceed via a hybrid argument consisting of $|C|$ steps, where in the ℓ -th hybrid $H_{\text{ZK}, \ell} = H_{\text{ZK}, \ell}(C, \mathbf{w})$, $\ell \in [0, |C|]$, the CRS-proof pair corresponding to the first ℓ gates (say, according to topological order) are simulated using \mathbf{S} . Hence $H_{\text{ZK}, 0}$ corresponds to $\langle \text{P}_{\text{ZK}} \rightarrow \text{V}_{\text{ZK}} \rangle(C, \mathbf{w})$ and $H_{\text{ZK}, |C|}$ corresponds to the intermediate distribution H_{ZK} . We aim to prove that for every $\ell \in [1, |C|]$

$$\text{SD}(H_{\text{ZK}, \ell-1}, H_{\text{ZK}, \ell}) \leq 4\varepsilon. \quad (5)$$

To this end, for the gate g that is switched from real to simulated in the ℓ -th hybrid, let us consider the “local distributions” R_g and S_g for g formally defined in Fig. 10.

In the following claim, we prove statistical indistinguishability of R_g and S_g assuming distributionally SZK of Π for \mathcal{D}_g . Note that the R_g and S_g are worst-case with respect to consistent assignments, and the crucial observation is that a worst-case to average-case reduction is possible since the number of consistent assignments for a Boolean gate is bounded.

Claim 4.3. *Fix a gate g . For any assignment $\mathbf{w} = (w_1, w_2, w_3) \in \{0, 1\}^3$ consistent with g ,*

$$\text{SD}(R_g(1^n, \mathbf{w}), S_g(1^n, \mathbf{w})) \leq 4\varepsilon .$$

Proof. Since there are at most four consistent assignments, and \mathcal{D}_g samples a uniformly random one, we

have that for every consistent assignment \mathbf{w} ,

$$\begin{aligned}
\frac{1}{4} \cdot \text{SD}(R_g(1^n, \mathbf{w}), S_g(1^n, \mathbf{w})) &= \frac{1}{4} \cdot \mathbb{E}_{\mathbf{r}, pk} [\text{SD}((crs, \pi_{\mathbf{w}}), S(\mathbf{c}))] \\
&\leq \mathbb{E}_{\mathbf{w}'} \left[\mathbb{E}_{\mathbf{r}, pk} [\text{SD}((crs, \pi_{\mathbf{w}'}) , S(\mathbf{c}))] \right] \\
&= \mathbb{E}_{((pk, \mathbf{c}), (\mathbf{w}', \mathbf{r})) \leftarrow \mathcal{D}_g(1^n)} [\text{SD}((crs, \pi_{\mathbf{w}'}) , S(\mathbf{c}))] \\
&\leq \varepsilon,
\end{aligned}$$

where $crs \leftarrow \text{Gen}(1^n)$, $r_i \leftarrow \{0, 1\}^{\text{poly}(n)}$ and $c_i := E(pk, w_i; r_i)$ for $i \in [3]$, for any consistent \mathbf{w}' , $\pi_{\mathbf{w}'} \leftarrow P(crs, (pk, \mathbf{c}) (\mathbf{w}', \mathbf{r}))$, and the last inequality follows from distributional SWI of Π for \mathcal{D}_g . \square

Since the proofs and CRSs for the rest of the gates in $H_{\text{ZK}, \ell-1}$ and $H_{\text{ZK}, \ell}$ can be generated from $(crs_g, \pi_g, (r_i, r_j, r_k))$ using the same random process, Claim 4.3 implies Eq. (5). By applying the triangle inequality to Eq. (5) with every $\ell \in [0, |C|]$, we get $\text{SD}(H_{\text{ZK}, 0}, H_{\text{ZK}, |C|}) \leq 4|C|\varepsilon$, completing the first step.

In the second step, all commitments (except that of the output wire) in H_{ZK} are switched to *dummy* commitments. The resulting distribution is the same as the output distribution of the simulator in Fig. 8. Since Λ is used in hiding mode, this switch is $(n-1)\delta$ -statistically-indistinguishable by Λ 's δ statistical hiding, i.e.:

$$\text{SD}(H_{\text{ZK}}(C, \mathbf{w}), S_{\text{ZK}}(C)) \leq (n-1)\delta. \quad (6)$$

Using the fact that $|C| \leq n$, Eq. (3) follows by applying the triangle inequality to Eqs. (4) and (6), completing the proof of SZK. \square

4.2 Dual-Mode NISZKA Privacy Amplification

Finally, we show how to reduce the ZK error of any NISZKA from inverse polynomial to negligibly small assuming LPKE. To be specific, given a dual-mode NISZKA with a sufficiently-small inverse-polynomial ZK error, we use LPKE and semi-honest MPC to obtain a dual-mode NISWIA with negligible WI error. By (a dual-mode variant of) Corollary 3.17, this (together with the fact that LPKE implies OWF) implies dual-mode NISZKA with negligible ZK error.

We start with the definition of semi-honest MPC, following the convention from [IKOS07].

Definition 4.4 (Multi-Party Computation (MPC) [IKOS07]). *For $n \in \mathbb{N}$, an MPC protocol is a protocol involving n parties P_1, \dots, P_n that takes place in $\rho = \rho(n)$ rounds of communication. The public input is denoted by x , while the private input and random coins of P_i are denoted by $w_i \in \{0, 1\}^{\text{poly}(|x|)}$ and $r_i \in \{0, 1\}^{\text{poly}(|x|)}$, respectively. The protocol is specified by its next-message function*

$$(m_{i,j,k})_{j \in [n]} := M(i, x, w_i, r_i, ((m_{j,i,1})_{j \in [n]}, \dots, (m_{j,i,k-1})_{j \in [n]})),$$

where, for $i \neq j \in [n]$ and $k \in [\rho]$, $m_{i,j,k}$ denotes the returns the message sent by P_i to P_j in round k . The view of party P_i , denoted by $v_i = v_i(x, w_1, \dots, w_n; r_1, \dots, r_n)$, consists of the private input and randomness, and the messages it receives over all rounds: i.e.,

$$v_i := \left(w_i, r_i, (m_{j,i,k})_{j \in [n], k \in [\rho]} \right). \quad (7)$$

Two views v_i and v_j are said to be consistent if the outgoing messages implicit in v_i are identical to the incoming messages in v_j and vice versa.

Definition 4.5 (t -Perfectly-Secure MPC in the Semi-Honest Model [IKOS07]). *For $t \leq n \in \mathbb{N}$, an MPC protocol M realises an n -party functionality $f = f(x, w_1, \dots, w_n)$ with t -perfect-security in the semi-honest model if the following properties hold:*

1. M realises f with perfect correctness. That is, for any input (x, w_1, \dots, w_n) , the probability (over the choice of r_1, \dots, r_n) that the output of some player is different from the value of f is 0.
2. M realises f with perfect t -privacy. That is, there is a PPT simulator S_M such that for any inputs x, w_1, \dots, w_n and any set of corrupted players $\mathcal{C} \subseteq [n]$ with $|\mathcal{C}| < t$, the distribution of joint views of players in \mathcal{C} , denoted $V_{\mathcal{C}} = V_{\mathcal{C}}(x, w_1, \dots, w_n)$, is identical to $S_M(\mathcal{C}, x, (w_i)_{i \in \mathcal{C}}, f_{\mathcal{C}}(x, w_1, \dots, w_n))$. Here $f_{\mathcal{C}}(x, w_1, \dots, w_n) := (f_i(x, w_1, \dots, w_n))_{i \in \mathcal{C}}$ and f_i denotes the i -th output of f .

Remark 12. An $n/2$ -perfectly-secure semi-honest MPC protocol was constructed in [BGW88] (also see [AL17]), which is what we rely on in Theorem 4.6.

Amplification Theorem. Our amplification protocol Π_{WI} is described formally in Fig. 11, followed by the amplification theorem in Theorem 4.6. The approach is similar in spirit to that in [GJS19] in the sense that the prover executes an MPC protocol “in its head” [IKOS07], commits to the view of each party in the execution and then proves consistency of each pair of views using the underlying proof system.

There are some key differences though:

1. We use LPKE instead of a commitment scheme. LPKE has two (indistinguishable) modes of operation: its binding keys act as perfectly-binding commitments with overwhelming probability (see Remark 11), whereas its lossy keys act as statistically-hiding commitments. Since we use lossy keys in the protocol (see Fig. 11, Line 1), we are able to show amplification of privacy using statistical tools. In particular, we build on the approach from [LM20] based on statistical coupling [Ald83]. On the other hand, when arguing soundness we first switch the protocol to a binding key and exploit the fact that it acts as a perfectly-binding commitment. Thus, we are able to avoid the argument based on (computational) hardcore lemmas [Imp95, Hol05], which [GJS19] rely on.
2. We commit to the views as a whole (as in [IKOS07]) instead of the fine-grained way of committing in [GJS19]. In more detail, [GJS19] commit to the private inputs of parties, their private coins and each message in the transcript using separate commitments; we only commit to the view of each party as a whole.
3. Since we start from a dual-mode NISZKA with a negligible soundness error, MPC correctness in our protocol is guaranteed and therefore semi-honest privacy suffices. This is in contrast to [GJS19] who rely on malicious security to deal also with a soundness error of the underlying proof system.

Before stating and proving the theorem that Π_{WI} amplifies hiding, we define an NP relation that captures *local consistency* of commitments to views. Consider any $(i, j) \in \binom{[n]}{2}$, any pk in the support of Λ 's key generation algorithms, and any instance $x \in \mathcal{L}(\mathcal{R}_n)$. We denote the NP relation for local consistency by $\mathcal{R}'_{i,j} := \mathcal{R}'_{M,x,pk,i,j}$, and $((c_i, c_j), (v_i, q_i, v_j, q_j)) \in \mathcal{R}'_{i,j}$ if

1. c_i (respectively, c_j) is the encryption of v_i (respectively, v_j) under pk using random coins q_i (respectively q_j); and
2. v_i and v_j are consistent (with respect to public input x and protocol M) and locally accepting (i.e., the local outputs are 1).

Theorem 4.6 (Amplification Theorem). *Consider the protocol $\Pi_{\text{WI}} = \langle P_{\text{WI}} \rightarrow V_{\text{WI}} \rangle$ described in Fig. 11 obtained by instantiating:*

1. M using an $n/2$ -perfectly-secure semi-honest MPC protocol.
2. Π_{ZK} using a dual-mode NISZKA with ZK error $\varepsilon = 1/100n$.
3. Λ using an LPKE with statistical hiding error δ .

Then Π_{WI} is a dual-mode NISWIA with following properties:

Given a base non-interactive protocol $\Pi_{\text{ZK}} = \langle \text{P}_{\text{ZK}} \rightarrow \text{V}_{\text{ZK}} \rangle$, an LPKE $\Lambda = (\text{KGen}, \text{LGen}, \text{E}, \text{D})$ and an MPC protocol M , the non-interactive protocol $\Pi_{\text{WI}} = \langle \text{P}_{\text{WI}} \rightarrow \text{V}_{\text{WI}} \rangle$ for $(x, w) \in \mathcal{R}_n$, where \mathcal{R} is any NP relation, is described below.

$\text{crs}_{\text{WI}} \leftarrow \text{Gen}_{\text{WI}}(1^n)$

1. Run Λ 's *lossy* key-generation algorithm to generate a lossy key $pk \leftarrow \text{LGen}(1^n)$
2. For $(i, j) \in \binom{[n]}{2}$, run set-up of Π_{ZK} in *statistically-hiding* mode to generate CRS: $\text{crs}_{i,j} \leftarrow \text{Gen}_{\text{ZK}}(1^n)$
3. Output $\text{crs}_{\text{WI}} := (pk, \text{crs}_{1,2}, \dots, \text{crs}_{n-1,n})$ as the CRS

$\text{crs}_{\text{WI}} \leftarrow \text{AGen}_{\text{WI}}(1^n)$

1. Run Λ 's *binding* key-generation algorithm to generate a binding key $pk \leftarrow \text{KGen}(1^n)$
2. For $(i, j) \in \binom{[n]}{2}$, run set-up of Π_{ZK} in *adaptively-sound* mode to generate CRS: $\text{crs}_{i,j} \leftarrow \text{AGen}_{\text{ZK}}(1^n)$
3. Output $\text{crs}_{\text{WI}} := (pk, \text{crs}_{1,2}, \dots, \text{crs}_{n-1,n})$ as the CRS

$\pi \leftarrow \text{P}_{\text{WI}}(\text{crs}_{\text{WI}}, x, w)$

1. Execute M “in the head”, using pk to *commit* to the views:
 - (a) Generate shares w_1, \dots, w_n of the witness w : sample $w_1, \dots, w_{n-1} \leftarrow \{0, 1\}^{|w|}$ and then set $w_n := w \oplus w_1 \oplus \dots \oplus w_{n-1}$.
 - (b) Samples random coins r_1, \dots, r_n for the n parties P_1, \dots, P_n .
 - (c) Set x as the public input, w_i and r_i , respectively, as P_i 's private input and random coins, and run M for the functionality $f(x, w_1, \dots, w_n) := \mathcal{R}(x, \bigoplus_{i=1}^n w_i)$. Let v_i denote P_i 's view in the above execution (see Eq. (7)).
 - (d) Commit to the views: for $i \in [n]$, sample random coins $q_i \leftarrow \{0, 1\}^{\text{poly}(|x|)}$ and compute $c_i := \text{E}(pk, v_i; q_i)$.

2. Prove pairwise local consistency. That is, for all $(i, j) \in \binom{[n]}{2}$, generate a proof that $(c_i, c_j) \in \mathcal{L}(\mathcal{R}'_{i,j})$:

$$\pi_{i,j} \leftarrow \text{P}_{\text{ZK}}(\text{crs}_{i,j}, (c_i, c_j), (v_i, q_i, v_j, q_j))$$

3. Output $\pi := ((c_1, \dots, c_n), (\pi_{1,2}, \dots, \pi_{n-1,n}))$

$b := \text{V}_{\text{WI}}(\text{crs}_{\text{WI}}, x, \pi)$

1. Accept (output $b := 1$) if and only if the underlying NISZKA verifier accepts all proofs. That is, for all $(i, j) \in \binom{[n]}{2}$: $\text{V}_{\text{ZK}}(\text{crs}_{i,j}, (c_i, c_j), \pi_{i,j}) = 1$.

Figure 11: Non-interactive protocol Π_{WI} .

- If Π_{ZK} has negligible (resp., 0) completeness error then so does Π_{WI} .
- The soundness error is negligible in the (computational) security parameter λ of the LPKE.
- The WI error is $2^{-n+1} + 2^{n+2}\delta n$. In particular, taking $\delta \leq 2^{-2n}n$, the WI error is at most $O(2^{-n})$.
- If P_{ZK} is non-uniform then so is P_{WI} .

Proof. The fact that Π_{WI} is non-uniform if Π_{ZK} is follows by construction. Mode indistinguishability follows from mode indistinguishability of the LPKE and Π_{ZK} . We prove the rest of the properties in Propositions 4.7 to 4.9.

Proposition 4.7 (Completeness). *If Π_{ZK} has completeness error ε_c in statistically-hiding (resp., adaptively-sound) mode, then Π_{WI} has completeness error at most $n^2 \cdot \varepsilon_c$ in statistically-hiding (resp., adaptively-sound) mode.*

Proof Sketch. Let us focus on statistically-hiding mode – the proof for adaptively-sound mode is similar. Completeness of Π_{WI} in statistically-hiding mode reduces to completeness of Π_{ZK} in statistically-hiding mode thanks to the perfect correctness of M as we argue next. If $(x, w) \in \mathcal{R}$, then by correctness of M all pairs of views (v_i, v_j) are consistent and locally accepting. This implies that $((c_i, c_i), (v_i, q_i, v_j, q_j)) \in \mathcal{R}'_{i,j}$ for every execution (i, j) .¹⁴ Since V_{WI} accepts if V_{ZK} accepts all the underlying proofs, and since the CRSs of Π_{ZK} are sampled independently, the correctness error is at most $(1 - (1 - \varepsilon_c)^{\binom{n}{2}}) \leq \varepsilon_c \cdot \binom{n}{2} \leq n^2 \cdot \varepsilon_c$. \square

Proposition 4.8 (Adaptive Soundness Preserved). *Π_{WI} is adaptively sound with a negligible soundness error.*

Proof Sketch. We reduce from adaptive soundness of the underlying NISZKA Π_{ZK} . Let $\text{P}^* = (\text{P}_n^*)_{n \in \mathbb{N}}$ be any polynomial-sized family of malicious provers that breaks Π_{WI} 's adaptive soundness with a non-negligible probability. Given a challenge CRS crs^* , the reduction generates

$$\text{crs}_{\text{WI}}^* := (pk, (\text{crs}_{1,2}, \dots, \text{crs}_{n-1,n})),$$

where pk is a binding key sampled using KGen , $\text{crs}_{i^*,j^*} := \text{crs}^*$ for $(i^*, j^*) \leftarrow \binom{[n]}{2}$, and the remaining CRSs are sampled independently using AGen_{ZK} . The reduction runs P^* on crs_{WI}^* , which outputs (x^*, π_{WI}^*) such that $x^* \notin \mathcal{L}(\mathcal{R})$ but $\pi_{\text{WI}}^* = ((c_1^*, \dots, c_n^*), (\pi_{1,2}^*, \dots, \pi_{n-1,n}^*))$ is accepted by V_{WI} . Recall that ciphertexts generated using binding keys of Λ act as perfectly-binding commitments with overwhelming probability (see Remark 10). Therefore, the commitments (c_1^*, \dots, c_n^*) open to some views (v_1^*, \dots, v_n^*) . Since $x^* \notin \mathcal{L}(\mathcal{R})$, there must exist at least one pair of views (v_i^*, v_j^*) such that $(c_i^*, c_j^*) \notin \mathcal{R}'_{i,j}$, i.e., the views are *inconsistent* with respect to the public input, which is the instance x^* . Otherwise, if all pairs of views are consistent with respect to x^* , then by perfect correctness of M it can be argued that $x^* \in \mathcal{L}(\mathcal{R})$ (see [IKOS07, Lemma 2.3] about local vs. global consistency). Since V_{WI} accepts π_{WI}^* if and only if all the underlying proofs accept, and $(i, j) = (i^*, j^*)$ with probability at least $1/n^2$, the reduction can output $((c_{i^*}^*, c_{j^*}^*), \pi_{i^*,j^*}^*)$ to break Π_{ZK} 's soundness with a $1/n^2$ loss, which is still non-negligible.¹⁵ \square

Before proving WI in Proposition 4.9, we establish some useful notation.

Notation 2 (Non-Standard String Notation).

- For a string or vector s of length n and a set $\mathcal{S} \subseteq [n]$, we use $s_{\mathcal{S}}$ to denote $(s_i)_{i \in \mathcal{S}}$.
- We let N denote $\binom{[n]}{2}$, and interpret a string $s \in \{0, 1\}^N$ as $(s_{i,j})_{(i,j) \in \binom{[n]}{2}}$.

¹⁴Note that we don't rely on correctness of decryption of Λ here and only use the fact that the encryption algorithm is a map once the random coins are fixed.

¹⁵This explains why we require Π_{ZK} to be adaptively sound (in the adaptively-sound mode) to start off with: the instance (c_i^*, c_j^*) and the associated proof $\pi_{i,j}^*$ that breaks Π_{ZK} 's soundness are determined by the output of the cheating prover P^* .

Let $n := |x|$ and $N := \binom{[n]}{2}$. For $s \in \{0,1\}^N$, the distribution $H_1 = H_{1,s}$ is defined below. For a distribution S over $\{0,1\}^N$, the hybrid distribution $\langle P_{\text{WI}}(w) \rightarrow V_{\text{WI}} \rangle_S(x)$ is defined as $H_{1,s}$ with s first sampled according to S .

$(pk, \mathbf{crs}, \mathbf{c}, \boldsymbol{\pi}) \leftarrow H_{1,s}(x, w)$

1. Sample lossy key $pk \leftarrow \text{LGen}(1^n)$
2. Generate the commitment \mathbf{c} as specified in Fig. 11, Line 1. That is:
 - (a) Run MPC as in Fig. 11, Lines 1a to 1c to generate views (v_1, \dots, v_n) .
 - (b) Compute the n -tuple of commitments \mathbf{c} , where $c_i := E(pk, v_i; q_i)$, as in Fig. 11, Line 1d.
3. Generate V_{ZK} 's views in Fig. 11, Line 2 depending on s . That is, sample $(\mathbf{crs}, \boldsymbol{\pi})$, where for $(i, j) \in \binom{[n]}{2}$

$$(crs_{i,j}, \pi_{i,j}) \leftarrow \begin{cases} \langle P_{\text{ZK}}((v_i, q_i, v_j, q_j)) \rightarrow V_{\text{ZK}} \rangle(c_i, c_j) & \text{if } s_{i,j} = 0 \\ S_{\text{ZK}}(c_i, c_j) & \text{otherwise.} \end{cases}$$

4. Output $(pk, \mathbf{crs}, \mathbf{c}, \boldsymbol{\pi})$.

Figure 12: Hybrid distribution $H_1 = H_{1,s}$.

- By $\{0,1\}_{\geq T}^N$, we denote subset of strings in $\{0,1\}^N$ with Hamming weight at least T : i.e. $\{0,1\}_{\geq T}^N := \{s \in \{0,1\}^N : \|s\|_0 \geq T\}$.

Notation 3 (Hybrid Distributions).

- Recall from Section 2 that $\langle P_{\text{WI}}(w) \rightarrow V_{\text{WI}} \rangle(x)$ denotes the random variable corresponding to V_{WI} 's views when the protocol in Fig. 11 is executed (in statistically-hiding mode) on $(x, w) \in \mathcal{R}$, i.e., $(pk, \mathbf{crs}, \mathbf{c}, \boldsymbol{\pi})$, where $(pk, \mathbf{crs}) \leftarrow \text{Gen}_{\text{WI}}(1^n)$ and $(\mathbf{c}, \boldsymbol{\pi}) \leftarrow P_{\text{WI}}((pk, \mathbf{crs}), x, w)$.
- Similarly, $\langle P_{\text{ZK}}((v_i, q_i, v_j, q_j)) \rightarrow V_{\text{ZK}} \rangle(c_i, c_j)$ denotes the random variable corresponding to V_{ZK} 's views when the protocol Π_{ZK} is executed (in statistically-hiding mode) on $((c_i, c_j), (v_i, q_i, v_j, q_j)) \in \mathcal{R}'_{i,j}$, i.e., $(crs_{i,j}, \pi_{i,j})$, where $crs_{i,j} \leftarrow \text{Gen}_{\text{ZK}}(1^n)$ and $\pi_{i,j} \leftarrow P_{\text{ZK}}(crs_{i,j}, (c_i, c_j), (v_i, q_i, v_j, q_j))$.
- For $s \in \{0,1\}^N$, we use $\langle P_{\text{WI}}(w) \rightarrow V_{\text{WI}} \rangle_s(x)$ to denote the hybrid distribution described in Fig. 12, where the views of V_{ZK} in executions (i, j) such that $s_{i,j} = 1$ are simulated using S_{ZK} (thus $\langle P_{\text{WI}}(w) \rightarrow V_{\text{WI}} \rangle_{0^N}(x)$ corresponds to the real view). For a distribution S over $\{0,1\}^N$, $\langle P_{\text{WI}}(w) \rightarrow V_{\text{WI}} \rangle_S(x)$ is defined as in Fig. 12, with s first sampled according to S .

Proposition 4.9 (Privacy Amplified). Π_{WI} is HVSWI with an error $2^{-n+1} + 2^{n+2}\delta n$.

Proof. The proof proceeds in two steps. We first prove in Claim 4.10 that Π_{WI} (in statistically-hiding mode) is a combiner; that is, if a large enough fraction of the NISZKA proofs are perfect ZK, then the resulting protocol is WI (with a negligible WI error). Then, taking a common approach in the literature, we prove in Claim 4.11 that any such combiner is also a good amplifier; that is, provided that every NISZK has a small enough ZK error ε , the resulting protocol is WI (with a negligible WI error related to that of the corresponding combiner). Specifically, the proof of the latter claim follows the ideas developed in [LM20].

Claim 4.10 (Π_{WI} is a Threshold Combiner). For $T := N - n/4 + 1$ and any $s \in \{0,1\}_{\geq T}^N$, $(x, w), (x, w') \in \mathcal{R}$,

$$\text{SD}(\langle P_{\text{WI}}(w) \rightarrow V_{\text{WI}} \rangle_s(x), \langle P_{\text{WI}}(w') \rightarrow V_{\text{WI}} \rangle_s(x)) \leq 2\delta n.$$

Let $n := |x|$ and $N := \binom{[n]}{2}$. For $s \in \{0, 1\}^N$, the distribution $H_2 = H_{2,s}$ is defined below.

$(pk, crs, c, \pi) \leftarrow H_{2,s}(x, w)$

1. Sample lossy key $pk \leftarrow \text{LGen}(1^n)$
2. Generate the commitment c with dummy messages for the parties in \mathcal{H} . That is:
 - (a) Run MPC as in Fig. 11, Lines 1a to 1c to generate views (v_1, \dots, v_n) .
 - (b) Compute the n -tuple of commitments c depending on h , where $h \in \{0, 1\}^n$ denotes the indicator string for \mathcal{H} determined by s :

$$c_i \leftarrow \begin{cases} \mathbb{E}(pk, v_i; q_i) \text{ for } q_i \leftarrow \{0, 1\}^{\text{poly}(|x|)} & \text{if } h_i = 0 \\ \mathbb{E}(pk, 0^{|v_i|}) & \text{otherwise} \end{cases}$$

3. Generate V_{ZK} 's views in Fig. 11, Line 2 depending on s . That is, sample (crs, π) , where for $(i, j) \in \binom{[n]}{2}$

$$(crs_{i,j}, \pi_{i,j}) \leftarrow \begin{cases} \langle \text{P}_{\text{ZK}}((v_i, q_i, v_j, q_j)) \rightarrow \text{V}_{\text{ZK}} \rangle(c_i, c_j) & \text{if } s_{i,j} = 0 \\ \text{S}_{\text{ZK}}(c_i, c_j) & \text{otherwise} \end{cases}$$

4. Output (pk, crs, c, π)

Figure 13: Hybrid distribution $H_2 = H_{2,s}$.

Claim 4.11 (Amplification from Threshold Combiners). *For $T := N - n/4 + 1$ and any $s \in \{0, 1\}_{\geq T}^N$, $(x, w), (x, w') \in \mathcal{R}$,*

$$\begin{aligned} & \text{SD}(\langle \text{P}_{\text{WI}}(w) \rightarrow \text{V}_{\text{WI}} \rangle(x), \langle \text{P}_{\text{WI}}(w') \rightarrow \text{V}_{\text{WI}} \rangle(x)) \leq \\ & 2^{-n+1} + 2^{n+1} \cdot \max_{s \in \{0, 1\}_{\geq T}^N} \text{SD}(\langle \text{P}_{\text{WI}}(w) \rightarrow \text{V}_{\text{WI}} \rangle_s(x), \langle \text{P}_{\text{WI}}(w') \rightarrow \text{V}_{\text{WI}} \rangle_s(x)) . \end{aligned}$$

Proof of Claim 4.10. We proceed via a hybrid argument, and let $H_1 = H_{1,s}$ denote the distribution $\langle \text{P}_{\text{WI}}(w) \rightarrow \text{V}_{\text{WI}} \rangle_s(x)$ from Fig. 12. Since $\|s\|_0 \geq T = N - n/4 + 1$ and as each proof depends on at most two parties, there exists a set $\mathcal{H} \subset [n]$ determined by s of size at least $n/2$ such that $s_{i,j} = 1$ holds for every $i \in \mathcal{H}$ and $j \in [n] \setminus \{i\}$. We think of these as the *honest* parties of the MPC protocol.

- In hybrid $H_2 = H_{2,s}$, we switch the messages underlying the ciphertexts $c_{\mathcal{H}}$ from honestly-generated views $v_{\mathcal{H}}$ to a dummy message independent of the witness w : see Fig. 13. To see why $\text{SD}(H_1, H_2) \leq \delta n$ (for any s), fix any $i \in \mathcal{H}$. Since the view $\langle \text{P}_{\text{ZK}}((v_i, q_i, v_j, q_j)) \rightarrow \text{V}_{\text{ZK}} \rangle(c_i, c_j)$ in every execution (i, j) , $j \in [n] \setminus \{i\}$, is simulated, it follows that the random coins q_i of the ciphertext c_i (which serve as part of witness for Π_{ZK}) are no longer required for generating proofs. Therefore, it is possible to use δ -statistical-hiding of Λ switch all ciphertexts in \mathcal{H} (of which there are at most n of).¹⁶
- In the next hybrid $H_3 = H_{3,s}$, we simulate the joint views $v_{\overline{\mathcal{H}}}$ of the remaining parties using the MPC simulator S_{MPC} : see Fig. 14. The ciphertexts and proofs that depend on $v_{\overline{\mathcal{H}}}$ are generated accordingly. Note that H_3 is distributed identical to H_2 thanks to $n/2$ -privacy of M .

We get that $\text{SD}(H_1, H_3) \leq \delta n$. By a symmetric argument to above it is possible to show that $\text{SD}(H_3, \langle \text{P}_{\text{WI}}(w') \rightarrow \text{V}_{\text{WI}} \rangle_s(x)) \leq \delta n$. The claim now follows by another application of the triangle inequality. \square

¹⁶Note that even though the random coins used to sample the keys are revealed, we have hiding as long as the keys are sampled correctly.

Let $n := |x|$ and $N := \binom{[n]}{2}$. For $s \in \{0, 1\}^N$, the distribution $H_2 = H_{2,s}$ is defined below.

$(pk, crs, c, \pi) \leftarrow H_{3,s}(x, w)$

1. Sample lossy key $pk \leftarrow \text{LGen}(1^n)$
2. Generate the commitment c with dummy messages for the parties in \mathcal{H} and simulated views for parties in $\overline{\mathcal{H}}$. That is:
 - (a) Simulate the MPC (joint) views $v_{\overline{\mathcal{H}}} := \text{S}_{\text{MPC}}(\overline{\mathcal{H}}, x, (w_i)_{i \in \overline{\mathcal{H}}}, (1, \dots, 1))$
 - (b) Compute the n -tuple of commitments c depending on h , where $h \in \{0, 1\}^n$ denotes the indicator string for \mathcal{H} determined by s :

$$c_i \leftarrow \begin{cases} \text{E}(pk, v_i; q_i) \text{ for } q_i \leftarrow \{0, 1\}^{\text{poly}(|x|)} & \text{if } h_i = 0 \\ \text{E}(pk, 0^{|v_i|}) & \text{otherwise} \end{cases}$$

3. Generate V_{ZK} 's views in Fig. 11, Line 2 depending on s . That is, sample (crs, π) , where for $(i, j) \in \binom{[n]}{2}$

$$(crs_{i,j}, \pi_{i,j}) \leftarrow \begin{cases} \langle \text{P}_{\text{ZK}}((v_i, q_i, v_j, q_j)) \rightarrow \text{V}_{\text{ZK}} \rangle(c_i, c_j) & \text{if } s_{i,j} = 0 \\ \text{S}_{\text{ZK}}(c_i, c_j) & \text{otherwise} \end{cases}$$

4. Output (pk, crs, c, π)

Figure 14: Hybrid distribution $H_3 = H_{3,s}$.

Proof of Claim 4.11. For $s \in \{0, 1\}^N$ and distribution S over $\{0, 1\}^N$, recall the distributions $H_{1,s} = \langle \text{P}_{\text{WI}}(w) \rightarrow \text{V}_{\text{WI}} \rangle_s(x)$ and $H_{1,S} = \langle \text{P}_{\text{WI}}(w) \rightarrow \text{V}_{\text{WI}} \rangle_S(x)$ defined in Fig. 12. Similarly, let $H'_{1,s}$ and $H'_{1,S}$ denote $\langle \text{P}_{\text{WI}}(w') \rightarrow \text{V}_{\text{WI}} \rangle_s(x)$ and $\langle \text{P}_{\text{WI}}(w') \rightarrow \text{V}_{\text{WI}} \rangle_S(x)$, respectively. Recall that our goal is to show that

$$\text{SD}(H_{1,0^N}, H'_{1,0^N}) \leq 2^{-n+1} + 2^{n+1} \cdot \max_{s \in \{0,1\}_{\geq T}^N} \text{SD}(H_{1,s}, H'_{1,s}).$$

First, for any distributions Z over $\{0, 1\}_{\geq T}^N \cup \{0^N\}$, we have

$$\begin{aligned} \text{SD}(H_{1,Z}, H'_{1,Z}) &= \frac{1}{2} \sum_h |H_{1,Z}(h) - H'_{1,Z}(h)| \\ &= \frac{1}{2} \sum_h \left| \sum_{z \in \{0,1\}_{\geq T}^N \cup \{0^N\}} Z(z) (H_{1,z}(h) - H'_{1,z}(h)) \right| \\ &\geq \frac{1}{2} \sum_h \left(Z(0^N) |H_{1,0^N}(h) - H'_{1,0^N}(h)| - \sum_{z \in \{0,1\}_{\geq T}^N} Z(z) |H_{1,z}(h) - H'_{1,z}(h)| \right) \\ &= Z(0^N) \cdot \text{SD}(H_{1,0^N}, H'_{1,0^N}) - \sum_{z \in \{0,1\}_{\geq T}^N} Z(z) \cdot \text{SD}(H_{1,z}, H'_{1,z}) \\ &\geq Z(0^N) \cdot \text{SD}(H_{1,0^N}, H'_{1,0^N}) - (1 - Z(0^N)) \cdot \max_{z \in \{0,1\}_{\geq T}^N} \text{SD}(H_{1,z}, H'_{1,z}). \end{aligned}$$

Thus, for any distribution Z over $\{0, 1\}_{\geq T}^N \cup \{0^N\}$ with $Z(0^N) > 0$, and distribution S over $\{0, 1\}_{\geq T}^N$,

$$\begin{aligned} \text{SD}(H_{1,0^N}, H'_{1,0^N}) &\leq Z(0^N)^{-1} \cdot \text{SD}(H_{1,Z}, H'_{1,Z}) + \max_{z \in \{0,1\}_{\geq T}^N} \text{SD}(H_{1,z}, H'_{1,z}) \\ &\leq Z(0^N)^{-1} \cdot (\text{SD}(H_{1,Z}, H_{1,S}) + \text{SD}(H_{1,S}, H'_{1,S}) + \text{SD}(H'_{1,S}, H'_{1,Z})) + \max_{z \in \{0,1\}_{\geq T}^N} \text{SD}(H_{1,z}, H'_{1,z}) \\ &\leq Z(0^N)^{-1} \cdot (\text{SD}(H_{1,Z}, H_{1,S}) + \text{SD}(H'_{1,Z}, H'_{1,S}) + 2Z(0^N)^{-1} \cdot \max_{z \in \{0,1\}_{\geq T}^N} \text{SD}(H_{1,z}, H'_{1,z}) . \end{aligned}$$

To complete the proof, we use the following Lemma.

Lemma 4.12. *There exist two distributions Z and S , where Z is over $\{0, 1\}_{\geq T}^N \cup \{0^N\}$, with $Z(0^N) > 2^{-n}$, and S is over $\{0, 1\}_{\geq T}^N$, such that*

$$Z(0^N)^{-1} \cdot \max \{ \text{SD}(H_{1,Z}, H_{1,S}), \text{SD}(H'_{1,Z}, H'_{1,S}) \} \leq \left(\frac{4eN}{n} \varepsilon \right)^{n/4} . \quad (8)$$

Indeed, for our setting of parameters, i.e., $N = \binom{n}{2}$ and $\varepsilon = 1/100n$, the value of Eq. (8) is at most $1/2^n$.

The proof of Lemma 4.12 is based on a coupling argument and roughly follows [LM20]. The proof can be found in Appendix A. This concludes the proof of Claim 4.11. \square

This concludes the proof of Proposition 4.9 \square

This concludes the proof of Theorem 4.6. \square

Acknowledgements

We thank Oded Goldreich for valuable comments and discussion. We also thank Chaya Ganesh for pointing us to the results in [CLP15].

Nir Bitansky is a member of the Checkpoint Institute of Information Security and is supported by the European Research Council (ERC) under the European Union's Horizon Europe research and innovation programme (grant agreement No. 101042417, acronym SPP).

Chethan Kamath is supported by Azrieli International Postdoctoral Fellowship, by the European Research Council (ERC) under the European Union's Horizon Europe research and innovation programme (grant agreement No. 101042417, acronym SPP), and by ISF grant 1789/19.

Omer Paneth is a member of the Checkpoint Institute of Information Security and is supported by an Azrieli Faculty Fellowship, and ISF grant 1789/19.

Ron Rothblum is funded by the European Union (ERC, FASTPROOF, 101041208). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

Prashant Nalini Vasudevan is supported by the National Research Foundation, Singapore, under its NRF Fellowship programme, award no. NRF-NRFF14-2022-0010.

References

- [AH91] William Aiello and Johan Håstad. Statistical zero-knowledge languages can be recognized in two rounds. *J. Comput. Syst. Sci.*, 42(3):327–345, 1991. 5
- [AL17] Gilad Asharov and Yehuda Lindell. A full proof of the BGW protocol for perfectly secure multiparty computation. *Journal of Cryptology*, 30(1):58–151, January 2017. 38

- [Ald83] David Aldous. Random walks on finite groups and rapidly mixing Markov chains. In *Séminaire de Probabilités XVII 1981/82: Proceedings*, pages 243–297. Springer, 1983. [38](#), [50](#)
- [AR21] Gal Arnon and Guy N. Rothblum. On prover-efficient public-coin emulation of interactive proofs. In Stefano Tessaro, editor, *2nd Conference on Information-Theoretic Cryptography, ITC 2021, July 23-26, 2021, Virtual Conference*, volume 199 of *LIPIcs*, pages 3:1–3:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. [4](#)
- [AR23] Noga Amit and Guy N. Rothblum. Constant-round arguments from one-way functions. In Barna Saha and Rocco A. Servedio, editors, *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, pages 1537–1544. ACM, 2023. [5](#)
- [BBD⁺20] Marshall Ball, Elette Boyle, Akshay Degwekar, Apoorvaa Deshpande, Alon Rosen, Vinod Vaikuntanathan, and Prashant Nalini Vasudevan. Cryptography from information loss. In Thomas Vidick, editor, *ITCS 2020*, volume 151, pages 81:1–81:27. LIPIcs, January 2020. [7](#)
- [BDRV18] Itay Berman, Akshay Degwekar, Ron D. Rothblum, and Prashant Nalini Vasudevan. Multi-collision resistant hash functions and their applications. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part II*, volume 10821 of *Lecture Notes in Computer Science*, pages 133–161. Springer, 2018. [5](#), [25](#)
- [BGW88] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In *20th ACM STOC*, pages 1–10. ACM Press, May 1988. [10](#), [38](#)
- [BHK17] Zvika Brakerski, Justin Holmgren, and Yael Tauman Kalai. Non-interactive delegation and batch NP verification from standard computational assumptions. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 474–482. ACM, 2017. [4](#)
- [BHKY19] Nir Bitansky, Iftach Haitner, Ilan Komargodski, and Eylon Yogev. Distributional collision resistance beyond one-way functions. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2019 - 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19-23, 2019, Proceedings, Part III*, volume 11478 of *Lecture Notes in Computer Science*, pages 667–695. Springer, 2019. [25](#)
- [BKP18] Nir Bitansky, Yael Tauman Kalai, and Omer Paneth. Multi-collision resistance: a paradigm for keyless hash functions. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 671–684. ACM, 2018. [3](#), [5](#)
- [Blu81] Manuel Blum. Coin flipping by telephone. In Allen Gersho, editor, *CRYPTO’81*, volume ECE Report 82-04, pages 11–15. U.C. Santa Barbara, Dept. of Elec. and Computer Eng., 1981. [28](#)
- [BMO90] Mihir Bellare, Silvio Micali, and Rafail Ostrovsky. The (true) complexity of statistical zero knowledge. In Harriet Ortiz, editor, *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing, May 13-17, 1990, Baltimore, Maryland, USA*, pages 494–502. ACM, 1990. [10](#)
- [CJJ21] Arka Rai Choudhuri, Abhishek Jain, and Zhengzhong Jin. Non-interactive batch arguments for NP from standard assumptions. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part IV*, volume 12828 of *LNCS*, pages 394–423, Virtual Event, August 2021. Springer, Heidelberg. [4](#)

- [CJJ22] Arka Rai Choudhuri, Abhishek Jain, and Zhengzhong Jin. SNARGs for \mathcal{P} from LWE. In *62nd FOCS*, pages 68–79. IEEE Computer Society Press, February 2022. 4, 5, 15
- [CKGS98] Benny Chor, Eyal Kushilevitz, Oded Goldreich, and Madhu Sudan. Private information retrieval. *J. ACM*, 45(6):965–981, 1998. 53
- [CLP15] Kai-Min Chung, Edward Lui, and Rafael Pass. From weak to strong zero-knowledge and applications. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015, Part I*, volume 9014 of *LNCS*, pages 66–92. Springer, Heidelberg, March 2015. 7, 12, 44
- [CW23] Jeffrey Champion and David J. Wu. Non-interactive zero-knowledge from non-interactive batch arguments. Cryptology ePrint Archive, Paper 2023/695, 2023. <https://eprint.iacr.org/2023/695>, to appear at Crypto’23. 7
- [Del16] Holger Dell. And-compression of NP-complete problems: Streamlined proof and minor observations. *Algorithmica*, 75(2):403–423, 2016. 7, 8, 18
- [DMO00] Giovanni Di Crescenzo, Tal Malkin, and Rafail Ostrovsky. Single database private information retrieval implies oblivious transfer. In Bart Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 122–138. Springer, Heidelberg, May 2000. 7, 8, 18, 53
- [DNR04] Cynthia Dwork, Moni Naor, and Omer Reingold. Immunizing encryption schemes from decryption errors. In Christian Cachin and Jan Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 342–360. Springer, Heidelberg, May 2004. 32
- [DNRS99] Cynthia Dwork, Moni Naor, Omer Reingold, and Larry J. Stockmeyer. Magic functions. In *40th FOCS*, pages 523–534. IEEE Computer Society Press, October 1999. 12
- [DPP97] Ivan Damgård, Torben P. Pedersen, and Birgit Pfitzmann. On the existence of statistically hiding bit commitment schemes and fail-stop signatures. *J. Cryptol.*, 10(3):163–194, 1997. 25
- [Dru15] Andrew Drucker. New limits to classical and quantum instance compression. *SIAM J. Comput.*, 44(5):1443–1479, 2015. 7, 8, 18
- [FLS90] Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In *31st Annual Symposium on Foundations of Computer Science, St. Louis, Missouri, USA, October 22-24, 1990, Volume I*, pages 308–317. IEEE Computer Society, 1990. 10, 28, 29
- [For89] Lance Fortnow. The complexity of perfect zero-knowledge. *Adv. Comput. Res.*, 5:327–343, 1989. 5
- [FS90] Uriel Feige and Adi Shamir. Witness indistinguishable and witness hiding protocols. In Harriet Ortiz, editor, *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing, May 13-17, 1990, Baltimore, Maryland, USA*, pages 416–426. ACM, 1990. 10, 25, 28
- [FS08] Lance Fortnow and Rahul Santhanam. Infeasibility of instance compression and succinct PCPs for NP. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 133–142. ACM Press, May 2008. 6
- [GH98] Oded Goldreich and Johan Håstad. On the complexity of interactive proofs with bounded communication. *Inf. Process. Lett.*, 67(4):205–214, 1998. 6
- [GJS19] Vipul Goyal, Aayush Jain, and Amit Sahai. Simultaneous amplification: The case of non-interactive zero-knowledge. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 608–637. Springer, Heidelberg, August 2019. 10, 38

- [GK96] Oded Goldreich and Ariel Kahan. How to construct constant-round zero-knowledge proof systems for NP. *J. Cryptol.*, 9(3):167–190, 1996. [28](#)
- [GMW86] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity and a methodology of cryptographic protocol design (extended abstract). In *27th Annual Symposium on Foundations of Computer Science, Toronto, Canada, 27-29 October 1986*, pages 174–187. IEEE Computer Society, 1986. [10](#), [25](#), [28](#)
- [Gol93] Oded Goldreich. A uniform-complexity treatment of encryption and zero-knowledge. *Journal of Cryptology*, 6(1):21–53, March 1993. [12](#)
- [Gol18] Oded Goldreich. On doubly-efficient interactive proof systems. *Found. Trends Theor. Comput. Sci.*, 13(3):158–246, 2018. [3](#), [11](#)
- [GOS06] Jens Groth, Rafail Ostrovsky, and Amit Sahai. Non-interactive zaps and new techniques for NIZK. In Cynthia Dwork, editor, *CRYPTO 2006*, volume 4117 of *LNCS*, pages 97–111. Springer, Heidelberg, August 2006. [10](#)
- [GOS12] Jens Groth, Rafail Ostrovsky, and Amit Sahai. New techniques for noninteractive zero-knowledge. *J. ACM*, 59(3):11:1–11:35, 2012. [14](#), [33](#)
- [GS89] Shafi Goldwasser and Michael Sipser. Private coins versus public coins in interactive proof systems. *Adv. Comput. Res.*, 5:73–90, 1989. [4](#)
- [GSV98] Oded Goldreich, Amit Sahai, and Salil P. Vadhan. Honest-verifier statistical zero-knowledge equals general statistical zero-knowledge. In Jeffrey Scott Vitter, editor, *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*, pages 399–408. ACM, 1998. [9](#)
- [GVW02] Oded Goldreich, Salil P. Vadhan, and Avi Wigderson. On interactive proofs with a laconic prover. *Comput. Complex.*, 11(1-2):1–53, 2002. [6](#), [7](#)
- [HHRS15] Iftach Haitner, Jonathan J. Hoch, Omer Reingold, and Gil Segev. Finding collisions in interactive protocols - tight lower bounds on the round and communication complexities of statistically hiding commitments. *SIAM J. Comput.*, 44(1):193–242, 2015. [5](#)
- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999. [28](#)
- [HJKS22] James Hulett, Ruta Jawale, Dakshita Khurana, and Akshayaram Srinivasan. SNARGs for P from sub-exponential DDH and QR. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 520–549. Springer, Heidelberg, May / June 2022. [4](#)
- [HM96] Shai Halevi and Silvio Micali. Practical and provably-secure commitment schemes from collision-free hashing. In Neal Koblitz, editor, *CRYPTO’96*, volume 1109 of *LNCS*, pages 201–215. Springer, Heidelberg, August 1996. [25](#)
- [HN10] Danny Harnik and Moni Naor. On the compressibility of NP instances and cryptographic applications. *SIAM J. Comput.*, 39(5):1667–1713, 2010. [6](#)
- [HNO⁺09] Iftach Haitner, Minh-Huyen Nguyen, Shien Jin Ong, Omer Reingold, and Salil P. Vadhan. Statistically hiding commitments and statistical zero-knowledge arguments from any one-way function. *SIAM J. Comput.*, 39(3):1153–1218, 2009. [25](#)
- [Hol05] Thomas Holenstein. Key agreement from weak bit agreement. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 664–673. ACM Press, May 2005. [38](#)

- [HRV18] Pavel Hubáček, Alon Rosen, and Margarita Vald. An efficiency-preserving transformation from honest-verifier statistical zero-knowledge to statistical zero-knowledge. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part III*, volume 10822 of *Lecture Notes in Computer Science*, pages 66–87. Springer, 2018. [9](#)
- [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. ACM Press, June 2007. [37](#), [38](#), [40](#)
- [Imp95] Russell Impagliazzo. Hard-core distributions for somewhat hard problems. In *36th FOCS*, pages 538–545. IEEE Computer Society Press, October 1995. [38](#)
- [Kil92] Joe Kilian. A note on efficient zero-knowledge proofs and arguments (extended abstract). In S. Rao Kosaraju, Mike Fellows, Avi Wigderson, and John A. Ellis, editors, *Proceedings of the 24th Annual ACM Symposium on Theory of Computing, May 4-6, 1992, Victoria, British Columbia, Canada*, pages 723–732. ACM, 1992. [3](#)
- [KLVW22] Yael Tauman Kalai, Alex Lombardi, Vinod Vaikuntanathan, and Daniel Wichs. Boosting batch arguments and RAM delegation. *IACR Cryptol. ePrint Arch.*, page 1320, 2022. [4](#), [6](#)
- [KMY20] Fuyuki Kitagawa, Takahiro Matsuda, and Takashi Yamakawa. NIZK from SNARG. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part I*, volume 12550 of *LNCS*, pages 567–595. Springer, Heidelberg, November 2020. [7](#)
- [KNY18] Ilan Komargodski, Moni Naor, and Eylon Yogev. Collision resistant hashing for paranoids: Dealing with multiple collisions. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part II*, volume 10821 of *Lecture Notes in Computer Science*, pages 162–194. Springer, 2018. [3](#), [25](#)
- [KRR⁺20] Inbar Kaslasi, Guy N. Rothblum, Ron D. Rothblum, Adam Sealfon, and Prashant Nalini Vasudevan. Batch verification for statistical zero knowledge proofs. In Rafael Pass and Krzysztof Pietrzak, editors, *Theory of Cryptography - 18th International Conference, TCC 2020, Durham, NC, USA, November 16-19, 2020, Proceedings, Part II*, volume 12551 of *Lecture Notes in Computer Science*, pages 139–167. Springer, 2020. [6](#)
- [KRV21] Inbar Kaslasi, Ron D. Rothblum, and Prashant Nalini Vasudevan. Public-coin statistical zero-knowledge batch verification against malicious verifiers. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology - EUROCRYPT 2021 - 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, October 17-21, 2021, Proceedings, Part III*, volume 12698 of *Lecture Notes in Computer Science*, pages 219–246. Springer, 2021. [6](#)
- [LFKN92] Carsten Lund, Lance Fortnow, Howard J. Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *J. ACM*, 39(4):859–868, 1992. [3](#)
- [LM20] David Lanzenberger and Ueli Maurer. Coupling of random systems. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part III*, volume 12552 of *LNCS*, pages 207–240. Springer, Heidelberg, November 2020. [10](#), [32](#), [38](#), [41](#), [44](#), [50](#), [51](#), [53](#)
- [LS19] Alex Lombardi and Luke Schaeffer. A note on key agreement and non-interactive commitments. Cryptology ePrint Archive, Report 2019/279, 2019. <https://eprint.iacr.org/2019/279>. [32](#)

- [LY94] Richard J. Lipton and Neal E. Young. Simple strategies for large zero-sum games with applications to complexity theory. In *26th ACM STOC*, pages 734–740. ACM Press, May 1994. [9](#), [19](#)
- [Mic94] Silvio Micali. CS proofs (extended abstracts). In *35th FOCS*, pages 436–453. IEEE Computer Society Press, November 1994. [4](#)
- [Nao91] Moni Naor. Bit commitment using pseudorandomness. *J. Cryptol.*, 4(2):151–158, 1991. [25](#)
- [Oka96] Tatsuaki Okamoto. On relationships between statistical zero-knowledge proofs. In Gary L. Miller, editor, *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996*, pages 649–658. ACM, 1996. [10](#)
- [OVY93] Rafail Ostrovsky, Ramarathnam Venkatesan, and Moti Yung. Interactive hashing simplifies zero-knowledge protocol design. In Tor Hellesest, editor, *Advances in Cryptology - EUROCRYPT '93, Workshop on the Theory and Application of Cryptographic Techniques, Lofthus, Norway, May 23-27, 1993, Proceedings*, volume 765 of *Lecture Notes in Computer Science*, pages 267–273. Springer, 1993. [10](#)
- [PP22] Omer Paneth and Rafael Pass. Incrementally verifiable computation via rate-1 batch arguments. In *63rd FOCS*, pages 1045–1056. IEEE Computer Society Press, October / November 2022. [4](#)
- [PVW08] Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A framework for efficient and composable oblivious transfer. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 554–571. Springer, Heidelberg, August 2008. [53](#)
- [PW08] Chris Peikert and Brent Waters. Lossy trapdoor functions and their applications. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 187–196. ACM Press, May 2008. [6](#)
- [RR20] Guy N. Rothblum and Ron D. Rothblum. Batch verification and proofs of proximity with polylog overhead. In Rafael Pass and Krzysztof Pietrzak, editors, *Theory of Cryptography - 18th International Conference, TCC 2020, Durham, NC, USA, November 16-19, 2020, Proceedings, Part II*, volume 12551 of *Lecture Notes in Computer Science*, pages 108–138. Springer, 2020. [3](#)
- [RRR18] Omer Reingold, Guy N. Rothblum, and Ron D. Rothblum. Efficient batch verification for UP. In Rocco A. Servedio, editor, *33rd Computational Complexity Conference, CCC 2018, June 22-24, 2018, San Diego, CA, USA*, volume 102 of *LIPICs*, pages 22:1–22:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. [3](#)
- [RRR21] Omer Reingold, Guy N. Rothblum, and Ron D. Rothblum. Constant-round interactive proofs for delegating computation. *SIAM J. Comput.*, 50(3), 2021. [3](#), [6](#)
- [Sha92] Adi Shamir. $IP = PSPACE$. *J. ACM*, 39(4):869–877, 1992. [3](#)
- [Vad99] Salil Pravin Vadhan. *A study of statistical zero-knowledge proofs*. PhD thesis, Massachusetts Institute of Technology, 1999. [5](#), [25](#)
- [vN28] J v. Neumann. Zur theorie der gesellschaftsspiele. *Mathematische annalen*, 100(1):295–320, 1928. [19](#)
- [VZ13] Salil P. Vadhan and Colin Jia Zheng. A uniform min-max theorem with applications in cryptography. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 93–110. Springer, Heidelberg, August 2013. [7](#), [9](#), [21](#)

- [Wee05] Hoeteck Wee. On round-efficient argument systems. In Luís Caires, Giuseppe F. Italiano, Luís Monteiro, Catuscia Palamidessi, and Moti Yung, editors, *Automata, Languages and Programming, 32nd International Colloquium, ICALP 2005, Lisbon, Portugal, July 11-15, 2005, Proceedings*, volume 3580 of *Lecture Notes in Computer Science*, pages 140–152. Springer, 2005. [3](#)
- [WW22] Brent Waters and David J. Wu. Batch arguments for NP and more from standard bilinear group assumptions. In *CRYPTO (2)*, volume 13508 of *Lecture Notes in Computer Science*, pages 433–463. Springer, 2022. [4](#)

A Proof of Lemma 4.12

Before stating and (re)proving Lemma 4.12, which completes the proof of Claim 4.11 and therefore Theorem 4.6, we define statistical coupling (Definition A.1) and recall a lemma about statistical coupling (Lemma A.2) that will be key to the proof. We also introduce some notation (Notation 4) that will help reduce clutter.

Definition A.1 (Statistical Coupling). *Let X and Y be two probability distributions defined on a finite set Ω . A joint probability distribution XY on Ω^2 is a statistical coupling of X and Y if its marginal distributions are X and Y , respectively, i.e., for every $x \in \Omega$:*

$$X(x) = \sum_{y \in \Omega} XY(x, y),$$

and for every $y \in \Omega$:

$$Y(y) = \sum_{x \in \Omega} XY(x, y).$$

Lemma A.2 (Coupling Lemma [Ald83]). *Let X and Y be probability distributions over the same set Ω . Then*

1. For every coupling XY of X and Y ,

$$\text{SD}(X, Y) \leq \Pr_{(x, y) \leftarrow XY} [x \neq y]$$

2. There exists an “optimal” coupling XY^* such that

$$\text{SD}(X, Y) = \Pr_{(x, y) \leftarrow XY^*} [x \neq y]$$

Notation 4 ([LM20]). *For two vectors of objects (e.g., distributions) $\mathbf{a} = (a_1, \dots, a_n)$ and $\mathbf{b} = (b_1, \dots, b_n)$, and a string $s \in \{0, 1\}^n$, we use $\langle \mathbf{a}/\mathbf{b} \rangle_s = \langle a_1/b_1, a_2/b_2, \dots, a_n/b_n \rangle_s$ to denote the vector $\mathbf{c} = (c_1, \dots, c_n)$ where*

$$c_i := \begin{cases} a_i & \text{if } s_i = 0 \\ b_i & \text{if } s_i = 1. \end{cases}$$

The following lemma is a restatement of Lemma 4.12 and is based on Lemma 7, Theorem 3 and Corollary 1 from [LM20]. The notation and presentation has been altered for the sake of compatibility with this paper.

Lemma A.3 ([LM20]). *There exists two distributions Z and S , where Z is over $\{0, 1\}_{\geq T}^N \cup \{0^N\}$, with $Z(0^N) > 2^{-n}$, and S is over $\{0, 1\}_{\geq T}^N$, such that*

$$Z(0^N)^{-1} \cdot \max \{ \text{SD}(H_{1,Z}, H_{1,S}), \text{SD}(H'_{1,Z}, H'_{1,S}) \} \leq \left(\frac{4eN}{n} \varepsilon \right)^{n/4}, \quad (9)$$

where $N, n, T, \varepsilon, H_{1,Z}, H_{1,S}, H'_{1,Z}$ and $H'_{1,S}$ are as defined in Section 4.2.

Proof. We proceed in two steps. First, we define what it means for a pair of distributions Z and S , as in the statement of the lemma, to be “good”. Then, we show that if Z and S are good then the lemma follows (this roughly corresponds to [LM20, Lemma 7]).

Step I. For two strings $s, e \in \{0, 1\}^N$, let $\text{blind}(s, e)$ be the function that returns the substring of s at indices i such that $e_i = 1$.

Definition A.4 (Good pair of distributions). Let $E = (E_{i,j})_{(i,j) \in \binom{[n]}{2}}$ denote a tuple of N independent Bernoulli distributions with bias at most ε , i.e., for every $(i, j) \in \binom{[n]}{2}$, it holds that

$$\Pr_{e_{i,j} \leftarrow E_{i,j}} [e_{i,j} = 1] \leq \varepsilon.$$

A pair of distributions Z and S is good if

- Z is supported on $\{0, 1\}_{\geq T}^N \cup \{0^N\}$,
- $Z(0^N) > 2^{-n}$,
- S is supported on $\{0, 1\}_{\geq T}^N$,
- $\mathbb{E}_{e \leftarrow E} [\text{SD}(\text{blind}(Z, e), \text{blind}(S, e))] \leq Z(0^N) \left(\frac{4eN}{n}\varepsilon\right)^{n/4}$.

In [LM20], a pair of such good distributions is constructed explicitly. For the sake of completeness, we include a description of the [LM20] distributions in Appendix A.1.

Step II. We now prove the lemma given that the pair of distributions Z and S described in Stage I is good. To be specific, we show

$$\text{SD}(H_{1,Z}, H_{1,S}) \leq \mathbb{E}_{e \leftarrow E} [\text{SD}(\text{blind}(Z, e), \text{blind}(S, e))]. \quad (10)$$

The proof of the corresponding claim for H' is similar and is hence omitted.

Recall the distributions $H_{1,Z}$ and $H_{1,S}$ from Fig. 12. Next, consider $H_{1,Z}$ and $H_{1,S}$ with the execution of M and Λ fixed, i.e., the lossy key pk , views \mathbf{v} and ciphertext-random coin pair (\mathbf{c}, \mathbf{q}) in distribution $H_{1,Z}$ and $H_{1,S}$ are fixed. To prove the lemma, it suffices to show Eq. (10) holds for every $(pk, \mathbf{v}, \mathbf{c}, \mathbf{q})$. Hence, from here on, let's consider $H_{1,Z}$ and $H_{1,S}$ with $(pk, \mathbf{v}, \mathbf{c}, \mathbf{q})$ fixed.

For $(i, j) \in \binom{[n]}{2}$, let's denote by $R_{i,j}$ and $I_{i,j}$ the random variables corresponding to the real and simulated execution of Π_{ZK} , respectively (see Fig. 12, Line 3). Following [LM20], we denote $H_{1,Z}$ by $H_1(\langle \mathbf{R}, \mathbf{I} \rangle_Z)$ (see Notation 4). Since Π_{ZK} is ZK with error at most $\varepsilon = 1/100n$, we have $\text{SD}(R_{i,j}, I_{i,j}) \leq \varepsilon$. As a result, by Lemma A.2, there exists an optimal coupling RI_{ij}^* of $R_{i,j}$ and $I_{i,j}$ such that

$$\Pr_{((crs_{R,i,j}, \pi_{R,i,j}), (crs_{I,i,j}, \pi_{I,i,j})) \leftarrow RI_{ij}^*} [(crs_{R,i,j}, \pi_{R,i,j}) \neq (crs_{I,i,j}, \pi_{I,i,j})] \leq \varepsilon. \quad (11)$$

Let's use $R_{i,j}^*$ and $I_{i,j}^*$ to denote the first and second argument of RI_{ij}^* , respectively. Then, we have

$$\begin{aligned} \text{SD}(H_{1,Z}, H_{1,S}) &= \text{SD}(H_1(\langle \mathbf{R}, \mathbf{I} \rangle_Z), H_1(\langle \mathbf{R}, \mathbf{I} \rangle_S)) \\ &= \text{SD}(H_1(\langle \mathbf{R}^*, \mathbf{I}^* \rangle_Z), H_1(\langle \mathbf{R}^*, \mathbf{I}^* \rangle_S)) \\ &\leq \text{SD}(\langle \mathbf{R}^*, \mathbf{I}^* \rangle_Z, \langle \mathbf{R}^*, \mathbf{I}^* \rangle_S). \end{aligned} \quad (12)$$

Here, the second equality follows by the definition of coupling (which requires the marginals to match) and the inequality is a consequence of data processing inequality. To upper bound Eq. (12), we set up a coupling experiment G , described in Fig. 15, involving $\mathbf{R}\mathbf{I}^*$. To see why G is a valid coupling, we claim that the marginal distributions of z and s sampled as part of G are Z and S respectively, z is independent of \mathbf{R} , and s is independent of \mathbf{I} . As a result, the marginal distributions of ζ and σ are the same as $H_{1,Z}$ and $H_{1,S}$ respectively. To see why the (marginal) distribution of z sampled as part of G is Z (the argument for s and S is analogous) and independent of \mathbf{R} , note that for any e sampled in the first step, z' is distributed as $\text{blind}(Z, e)$, and z is sampled from Z conditioned on z' .

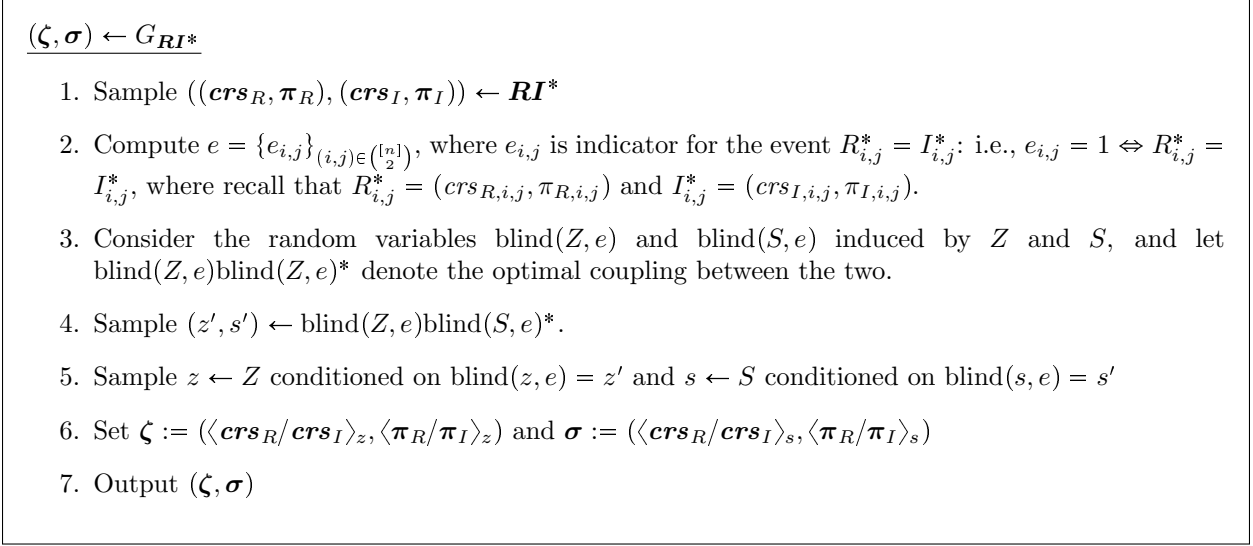


Figure 15: Coupling experiment G .

Therefore, by Lemma A.2 (“for every” claim), we have from Eq. (12) that

$$\begin{aligned}
\text{SD}(\langle \mathbf{R}^*, \mathbf{I}^* \rangle_Z, \langle \mathbf{R}^*, \mathbf{I}^* \rangle_S) &\leq \Pr_G[\zeta \neq \sigma] \\
&= \sum_{e' \in \binom{[n]}{2}} \Pr_G[\zeta \neq \sigma, e = e'] \\
&= \sum_{e' \in \binom{[n]}{2}} \Pr_G[\text{blind}(z, e) \neq \text{blind}(s, e), e = e'] \\
&= \sum_{e' \in \binom{[n]}{2}} \Pr_{G|e=e'}[\text{blind}(z, e') \neq \text{blind}(s, e')] \cdot \Pr_G[e = e'] \\
&= \sum_{e' \in \binom{[n]}{2}} \text{SD}(\text{blind}(Z, e'), \text{blind}(S, e')) \cdot \Pr_G[e = e'] \tag{13}
\end{aligned}$$

where ζ, σ, e and (z, s) above are sampled as part of G , and Eq. (13) follows by optimality of $\text{blind}(Z, e)\text{blind}(S, e)^*$. Since each $e_{i,j}$ in G is distributed as required (because Π_{zk} is zero-knowledge with error ε and the executions are independent), we get from Eqs. (12) and (13) that

$$Z(0^N)^{-1} \cdot \text{SD}(H_{1,Z}, H_{1,S}) \leq Z(0^N)^{-1} \cdot \mathbb{E}_{e' \leftarrow E} [\text{SD}(\text{blind}(Z, e'), \text{blind}(S, e))] \leq \left(\frac{4eN}{n} \varepsilon \right)^{n/4},$$

where the final inequality follows from the fact that Z and S is a pair of good distributions. □

A.1 Good Distributions

We recall the good pair of distributions Z and S defined in [LM20], described using multisets, \mathcal{Z} and \mathcal{S} , respectively:¹⁷

$$\begin{aligned} \mathcal{S} &:= \bigcup_{j \in \{T, T+2, \dots, N\}} \left\{ \left(b, \binom{j-1}{T-1} \right) : b \in \{0, 1\}^N, \|b\|_0 = j \right\}, \text{ and} \\ \mathcal{Z} &:= \{(0^N, 1)\} \cup \bigcup_{j \in \{T+1, T+3, \dots, N\}} \left\{ \left(b, \binom{j-1}{T-1} \right) : b \in \{0, 1\}^N, \|b\|_0 = j \right\}. \end{aligned}$$

The proof that Z and S constitute a pair of good distributions can be found in [LM20, Theorem 3 and Corollary 1].

B LPKE via Somewhere Extractable NIBARGs

In this section, we recall the definition of *somewhere extractable* BARGs from the literature, and also define a variant thereof, which we call *honestly somewhere extractable*. We prove that somewhere-extractable BARGs imply (single databased) *private information retrieval* (PIR), which in turn is known to imply *statistically sender-private oblivious transfer* and lossy public-key (LPKE) encryption [DMO00, PVW08].

The resulting LPKE suffers from a negligible decryption error, which makes it insufficient for the NISZKA amplification theorem in Section 4.2. We observe that if the BARGs satisfy honest-somewhere extraction then the resulting LPKE has a stronger correctness guarantee, which is also sufficient for our amplification theorem.

B.1 PIR from Somewhere Extractability

Definition B.1 (Somewhere Extractability). *A batch protocol $(\text{Gen}, \text{TGen}, \text{P}, \text{V})$ for a relation \mathcal{R} is somewhere extractable if it satisfies CRS indistinguishability, and if there is a PPT extractor E such that, for every polynomial t and polynomial-size circuit family of provers $\text{P}^* = (\text{P}_\lambda^*)_{\lambda \in \mathbb{N}}$, there is a negligible function μ such that for every $\lambda \in \mathbb{N}$, $t = t(\lambda)$, and $i^* \in [t]$:*

$$\Pr_{\text{crs}^*, \text{td}, \text{E}} [\text{V}(\text{crs}^*, (x_1, \dots, x_t), \pi) \text{ accepts} \wedge (x_{i^*}, w) \notin \mathcal{R}] \leq \mu(\lambda),$$

where $(\text{crs}^*, \text{td}) \leftarrow \text{TGen}(1^\lambda, 1^t, i^*)$, $((x_1, \dots, x_t), \pi) \leftarrow \text{P}_\lambda^*(\text{crs}^*, i^*)$, and $w \leftarrow \text{E}(\text{td}, i^*, \text{crs}^*, (x_1, \dots, x_t), \pi)$.

Definition B.2 (PIR [CKGS98]). *A one-round, single-database PIR is a tuple of polynomial-time algorithms $(\text{Q}, \text{D}, \text{R})$ with the following syntax:*

- $(k, Q) \leftarrow \text{Q}(1^\lambda, \ell, i)$. *The randomized user query algorithm takes as input a security parameter $\lambda \in \mathbb{N}$, a parameter $\ell \in \mathbb{N}$ that represents the length of the database, and a target index $i \in [\ell]$. It outputs a key k and a query Q .*
- $a := \text{D}(D, Q)$. *The deterministic database answer algorithm takes as input a database $D := (D_1, \dots, D_\ell) \in \{0, 1\}^\ell$ and a query Q and outputs an answer a .*
- $d := \text{R}(k, a)$. *The deterministic user reconstruct algorithm takes as input the key k and answer a and outputs a data bit d .*

We require the following properties:

¹⁷A multiset \mathcal{M} over a domain Ω is represented as $\{(x, m_x) : x \in \Omega\}$ where $m_x \in \mathbb{N}$ is the multiplicity of the element x . The cardinality of \mathcal{M} is then defined as $|\mathcal{M}| := \sum_{x \in \Omega} m_x$. The probability distribution M induced by \mathcal{M} is defined naturally: the probability of an element $x \in \Omega$ is $m_x/|\mathcal{M}|$.

PIR scheme (Q, D, R) , built using a somewhere-extractable NIBARG $(\text{Gen}, \text{TGen}, P, V)$ and hard sampler for the relation \mathcal{R}_f from Eq. (14).

$(k, Q) \leftarrow Q(1^\lambda, \ell, i)$

1. Use the hard sampler for \mathcal{R}_f to generate ℓ instance-witness pairs

$$\mathbf{q} := (((y_{1,0}, y_{1,1}), x_{1,0}, x_{1,1}), \dots, ((y_{\ell,0}, y_{\ell,1}), x_{\ell,0}, x_{\ell,1})).$$

2. Use TGen to sample a CRS with trapdoor set up at index i : $(crs^*, td) \leftarrow \text{TGen}(1^\lambda, 1^\ell, i)$
3. Output $(td, (crs^*, \mathbf{q}))$

$a := D(D, Q)$

1. Run the batch prover on witnesses determined by D :

$$\pi \leftarrow P(crs^*, ((y_{1,0}, y_{1,1}), \dots, (y_{\ell,0}, y_{\ell,1})), (x_{1,D_1}, \dots, x_{\ell,D_\ell})).$$

2. Output π

$d := R(k, a)$

1. Use BARG extractor to extract witness at i : $w \leftarrow E(td, i, crs^*, ((y_{1,0}, y_{1,1}), \dots, (y_{\ell,0}, y_{\ell,1})), \pi)$
2. Halt without output if the BARG verifier rejects, i.e., $V(crs^*, ((y_{1,0}, y_{1,1}), \dots, (y_{\ell,0}, y_{\ell,1})), \pi) = 0$
3. Otherwise,

$$d := \begin{cases} 0 & \text{if } f(w) = y_{i,0} \\ 1 & \text{otherwise.} \end{cases}$$

4. Output d

Figure 16: PIR scheme (Q, D, R) .

1. Correctness of reconstruction. *There exists a negligible function μ such that for every $\lambda \in \mathbb{N}$, $\ell \in \text{poly}(\lambda)$, database $D \in \{0, 1\}^\ell$ and query $i \in [\ell]$:*

$$\Pr_{(k, Q) \leftarrow Q(1^\lambda, \ell, i)} [R(k, D(D, Q)) = D_i] \geq 1 - \mu(\lambda).$$

2. Succinctness. *We say that the PIR is succinct if $|a| \leq \ell^\epsilon$ for some $\epsilon < 1$. We say that the PIR is fully succinct if there exists poly such that $|a| \leq \text{poly}(\lambda)$.*
3. Computational user privacy. *No efficient adversary can distinguish between user queries on two target indices. That is, for every polynomial-size circuit family of distinguishers $\mathbf{A} = \{A_\lambda\}_{\lambda \in \mathbb{N}}$, there is a negligible function μ , such that for all $\lambda \in \mathbb{N}$, $\ell \in \text{poly}(\lambda)$ and $i, j \in [\ell]$*

$$\left| \Pr_{(k, Q) \leftarrow Q(1^\lambda, \ell, i)} [1 \leftarrow A_\lambda(Q)] - \Pr_{(k, Q) \leftarrow Q(1^\lambda, \ell, j)} [1 \leftarrow A_\lambda(Q)] \right| \leq \mu(\lambda).$$

The PIR scheme constructed from somewhere-extractable NIBARG is described in Fig. 16. It relies on

the fact that somewhere-extractable NIBARG implies one-way functions (OWFs), and given a OWF f , we can define an NP relation

$$\mathcal{R}_f := \{((y_0, y_1), x) : f(x) = y_0 \vee f(x) = y_1\} \quad (14)$$

that allows sampling an instance along with *two* witnesses. To be precise, the hard sampler for \mathcal{R}_f invokes the OWF on two random preimages x_0 and x_1 , and then outputs the instance $(y_0 := f(x_0), y_1 := f(x_1))$.

Theorem B.3 (Somewhere-Extractable NIBARG Implies PIR). *If there exists a somewhere-extractable NIBARG $(\text{Gen}, \text{TGen}, \text{P}, \text{V})$, then the scheme in Fig. 16 is a one-round, single database PIR. If the size of BARGs is independent of the number of instances, then the PIR is fully succinct.*

Proof Sketch. User privacy follows from CRS indistinguishability (Definitions 2.21 and B.1). Correctness, follows from somewhere extractability and one-wayness. By somewhere extractability, it is guaranteed that with overwhelming probability the extractor returns *some* witness of $(y_{i,0}, y_{i,1})$, i.e., *some* pre-image of $y_{i,0}$ or $y_{i,1}$ under f . One-wayness of f ensures that it returns a witness corresponding to y_{i,D_i} and not y_{i,\bar{D}_i} . Indeed, since Q generates the NIBARG proof based only on the witness/pre-image x_{i,D_i} , it is oblivious of the other witnesses x_{i,\bar{D}_i} . As a result, the extractor outputting x_{i,\bar{D}_i} is tantamount to breaking f 's one-wayness. □ Theorem B.3

B.2 Honest Somewhere Extractability

Definition B.4 (Honest Somewhere Extractability). *A batch protocol $(\text{Gen}, \text{TGen}, \text{P}, \text{V})$ for a relation \mathcal{R} is honestly somewhere extractable if it satisfies CRS indistinguishability, and if there is a PPT extractor E such that, for every $\lambda \in \mathbb{N}$, $t = t(\lambda)$, $(x_1, w_1), \dots, (x_t, w_t) \in \mathcal{R}$ and $i^* \in [t]$:*

$$\Pr_{crs^*, td, \text{E}} [w_{i^*} \neq w] = 0,$$

where $(crs^*, td) \leftarrow \text{TGen}(1^\lambda, 1^t, i^*)$, $\pi \leftarrow \text{P}(crs^*, (x_1, w_1), \dots, (x_t, w_t))$, and $w \leftarrow \text{E}(td, i^*, crs^*, (x_1, \dots, x_t), \pi)$.

Remark 13. *We can in fact further weaken the above requirement, asking for perfect correctness for almost any CRS. Namely, that with overwhelming probability over the choice of CRS, extraction is perfect.*

Going back to the construction in Fig. 16, in case the BARG satisfies *honest* somewhere extractability (Definition B.4), then the construction satisfies perfect correctness of reconstruction. Indeed, the NIBARG proof generated by Q is honest, the extractor is guaranteed to return the actual witness used at position i , which is x_{i,D_i} .