

On Witness Encryption and Laconic Zero-Knowledge Arguments

Yanyi Liu*

Noam Mazor[†]

Rafael Pass[‡]

January 6, 2025

Abstract

Witness encryption (WE) (Garg et al, STOC’13) is a powerful cryptographic primitive that is closely related to the notion of *indistinguishability obfuscation* (Barak et al, JACM’12, Garg et al, FOCS’13). For a given NP-language L , WE for L enables encrypting a message m using an instance x as the public-key, while ensuring that efficient decryption is possible by anyone possessing a witness for $x \in L$, and if $x \notin L$, then the encryption is hiding.

We show that this seemingly sophisticated primitive is *equivalent* to a communication-efficient version of one of the most classic cryptographic primitives—namely that of a *zero-knowledge argument* (Goldwasser et al, SIAM’89, Brassard et al, JCSS’88): for any NP-language L , the following are equivalent:

- There exists a *witness encryption* for L ;
- There exists a *laconic* (i.e., the prover communication is bounded by $O(\log n)$) special-honest verifier zero-knowledge (SHVZK) argument for L .

Our approach is inspired by an elegant (one-sided) connection between (laconic) zero-knowledge arguments and public-key encryption established by Berman et al (CRYPTO’17) and Cramer-Shoup (EuroCrypt’02), and the equivalence between a notion of so-called “predictable arguments” and witness encryption by Faonio, Nielsen, and Venturi (PKC’17).

*Cornell Tech. E-mail: y12866@cornell.edu. Research partly supported by NSF CNS-2149305.

[†]Tel Aviv University. E-mail: noammaz@gmail.com. Research partly supported by NSF CNS-2149305, AFOSR Award FA9550-23-1-0312 and AFOSR Award FA9550-23-1-0387 and ISF Award 2338/23.

[‡]Cornell Tech, Technion and Tel Aviv University. E-mail: rafael@cs.cornell.edu. Supported in part by NSF Award CNS 2149305, AFOSR Award FA9550-23-1-0387, AFOSR Award FA9550-23-1-0312 and AFOSR Award FA9550-24-1-0267 and ISF Award 2338/23. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Government, or the AFOSR.

1 Introduction

Witness encryption (WE) is a fascinating cryptographic primitive introduced by Garg, Gentry, Sahai and Waters [GGSW13]. Roughly speaking, a WE for an NP-language L , enables encrypting a message m with an instance x as the “public key”, such that (a) if $x \in L$, then that anyone having a witness $w \in R_L(x)$ for x can recover m from the ciphertext c , yet (b) if $x \notin L$, then the ciphertext c hides m (in the standard sense of semantic security).

On top of being intriguing in its own right, WE for NP together with standard cryptographic hardness assumption yields other more advanced “trustless” encryption schemes (e.g., flexible broadcast encryption and more) [FWW23]. Furthermore, as already shown in the original work of [GGSW13], WE for NP and just the assumption of one-way functions, yields public-key encryption. The assumption of one-way functions can be further relaxed to just assuming that $\text{NP} \not\subseteq \text{P/poly}$ [KMN⁺14, HN23, LMP24]. In fact, for these results, it further suffices to assume the existence of WE even just for specific NP languages such as MCSP [KC00], or MK^tP [Kol68, Ko86] (see [LMP24] for more details).

For what languages does WE exists? If $\text{NP} = \text{P}$, then WE for every language in NP trivially exists: simply decide the instance and if it is true send m in the clear and otherwise nothing. But if $\text{NP} \neq \text{P}$, then establishing WE for NP has proven harder, and WE is generally viewed as being closely related to the powerful notion of *indistinguishability obfuscation* (iO) [BGI⁺12, GGH⁺13]. Indeed, as shown by [GGH⁺13], the existence of iO implies WE for all of NP and thus, all constructions of iO (such as, the recent constructions of Jain, Lin, and Sahai [JLS21, JLS22] which can be based on well-founded assumptions) can be used to get secure WE. Furthermore, up until recently, all known construction of WE for NP either passed through the route of iO, or were a stepping-stone/inspiration for achieving iO. More recently, however, direct WE constructions were proposed in [VWW22, Tsa22] based on a type of knowledge-based LWE assumption (referred to as evasive LWE) from which iO constructions are not known, highlighting that WE may perhaps be a weaker primitive than iO (and indeed so-called black-box separations between the primitives are known [GMM17]).

When focusing on WE for *specific* languages (as opposed to for all of NP), as noted already in [GGSW13] and further formalized in [ABP15], every language L having a so-called *smooth projective hash function* (a.k.a a *hash proof system*) [CS02] unconditionally has a WE (but the converse is not known); these include various number-theoretic languages. However, as noted in [GGSW13], all such languages are in SZK and this approach is thus unlikely to extend to all of NP (unless the Polynomial-Hierarchy collapses [AH91]).

The current state of the art thus leaves open the following basic question:

For what languages does WE exist, and under what hardness assumptions?

In this work, rather than directly providing an answer to the above question, and instead of trying to build WE based on concrete hardness assumptions, or for concrete languages, we show that, perhaps surprisingly, the above question is equivalent to a corresponding question regarding a seemingly much simpler and long-studied cryptographic primitive, namely so-called *laconic ZK arguments*. Before diving in, we highlight that we are not the first to establish a connection between WE and interactive arguments of some form; indeed, as we will discuss in more detail in Section 1.1, Faonio, Nielsen, and Venturi [FNV17] show an equivalence between a notion of a so-called “predictable argument” and WE, and several other works [BC20, Kiy24, BISW18, BLOW20] show that more restricted/non-standard forms of laconic or ZK arguments imply their notion of predictability. (In essence, as we discuss in more detail later on, all such arguments essentially require a deterministic Prover.) In contrast to those works, we here focus on just standard notions of laconicity and ZK (which most notably allow for probabilistic Provers.)

Laconic Special-HVZK *Zero-knowledge interactive proofs/arguments* [GMR89, BCC88] are one of the central cryptographic protocols; they are interactive protocols whereby a Prover, P , can convince a Verifier, V , that some instance $x \in L$, but with the additional (seemingly) paradoxical zero-knowledge (ZK) guarantee that the Verifier does not “learn anything new” except for the fact that $x \in L$. In our treatment (and as is typically the case for cryptographic applications), we will be restricting to ZK proofs/argument where the the Prover P can be implemented in polynomial time given access to any witness $w \in R_L(x)$ (a.k.a. *efficient-prover* ZK protocols).

We will be focusing on so-called *honest-verifier ZK protocols* (HVZK), where the zero-knowledge property only needs to hold with respect to the honest verifier V ; this property is formalized by requiring the existence of a polynomial-time simulator S that can simulate the view of V in a random interaction with the Prover in a computationally indistinguishable way (given just the instance x as input). In many applications of HVZK protocols, it is often convenient to consider a slight strengthening of this notion, referred to as *special-HVZK* (SHVZK) [CDS94, Dam02], where the simulator receives not only the instance x , but also a randomly-sampled random-tape r and is required to simulate the view of $V_r(x)$ (i.e., the honest verifier with random tape fixed to r) in an interaction with $P(x, w)$. (SHVZK is not only easier to use in applications (see e.g., [CDS94]), but as far as we know, all natural interactive HVZK protocol also satisfy this property.¹)

We will be consider SHVZK proofs/arguments satisfying an additional communication-efficient property: We say that a proof system is ℓ -*laconic* if the total length of the prover messages is bounded by $\ell(|x|)$, and we simply refer to the proof system as being laconic if it is $O(\log n)$ -laconic.

The study of laconic proof systems was introduced by Goldreich and Håstad [GH98] and further expanded in the work of Goldreich, Vadhan and Wigderson [GVW02]. More recently, Berman, Degwekar, Rothblum and Vasudevan [BDRV18] demonstrated an intriguing (one-sided) connection between laconic HVZK protocols (satisfying a certain quantitative form of laconicity) and *public key encryption* (PKE; in more detail, they show that the existence of an r -round statistical HVZK argument for a language L that is $O(r^{-2/3} \log^{1/3} n)$ -laconic, together with the assumption that L satisfies a certain notion of “cryptographic hardness”²), implies the existence of public-key encryption.³

1.1 Our results

The above-mentioned result of [BDRV18] will be the starting point for our work. In essence, we will show that once we consider SHVZK (instead of just HVZK), then not only we can deal with an arbitrary level of laconicity (i.e., any $O(\log n)$ prover communication complexity), but we can also construct WE instead of just public-key encryption (and also without needed to make any hardness assumption on L). Our main result is the following theorem.

Theorem 1.1 (Characterizing WE). *For any language $L \in \text{NP}$, the following are equivalent:*

- *There exists a WE (resp. a statistically-secure WE) for L .*
- *There exists an efficient-prover laconic SHVZK argument (resp. proof) for L .*

¹But it is easy to come up with examples separating the notions. In particular, NIZK constructions such as [FLS90] are trivially HVZK but the simulator is required to “program” the Common Random String (e.g., by setting the randomness to be the output of a PRG) to be able to perform the simulator. Such protocols are not SHVZK (as simulation can only be performed when the Verifier randomness is in the range of the PRG).

²Namely that there exist a PPT machine to sample YES-instances together with their witness, and a PPT machine to sample NO-instances that are computationally indistinguishable from the YES-instances.

³The also show a weak converse of this theorem, that public-key encryption also implies a relaxed type of “average-case laconic HVZK arguments of weak knowledge”.

We emphasize that our characterization considers WE protocols satisfying correctness with overwhelming probability (as opposed to probability 1)—that is, we do not characterize WE with perfect correctness.

As a consequence of Theorem 1.1 (combined with [HN23, LMP24]), we get that laconic SHVZK arguments for NP, or even just for specific languages such as MCSP, or MK^tP , together with the assumption that $NP \not\subseteq P/\text{poly}$ yields the existence WE and OWFs, which by [GGSW13] yields the existence of PKE. (Or alternatively, simply that laconic SHVZK arguments for some language satisfying cryptographic hardness implies PKE.⁴)

On WE and Predictable Arguments The elegant work of Faonio, Nielsen, and Venturi introduced the notion of *predictable arguments* [FNV17] and showed deep connections between those and witness encryption. An argument is said to be “predictable” if given the Verifier’s (private) randomness, only a single prover response can lead the Verifier to accept, and this response is required to be efficiently computable (with the Verifier randomness). They showed that for any language $L \in NP$, the existence of a predictable argument for L is equivalent to the existence of a WE for L .

Subsequently, there have been a number of works that have established sufficient conditions for constructing predictable arguments (and thus also WE) based on various proof systems with certain laconicity or zero-knowledge properties:

- Bitansky and Choudhuri [BC20] obtained a predictable argument from any *deterministic-prover* (malicious-verifier) zero-knowledge with bounded-auxiliary input (where the bounded-auxiliary input property is employed to overcome the impossibility result of [GO94]⁵.
- Kiyoshima [Kiy24] showed that any arguments with a strong (non-standard) form of a witness indistinguishability (WI) property [FS90] yields a predictable argument.⁶ This WI property is implied by a notion of “resettable statistical witness indistinguishability”, and as such, his results show that resettable *statistical* zero-knowledge implies predictable arguments.
- Boneh, Ishai, Sahai, and Wu show 1-bit laconic arguments (together with an extra average-case hardness assumptions) imply predictable arguments [BISW18]; furthermore, Barta, Ishai, Ostrovsky, and Wu extend this also to 1-group-element laconic arguments where the Verifier only performs generic group operations [BIOW20].

We note that any predictable argument directly implies a deterministic-prover SHVZK argument: the SHVZK simulator simply outputs the “predicted” prover messages (which are the same as what the deterministic prover would send). Indeed, the “warm-up” case of our proof (which deals with deterministic SHVZK) can be viewed as a generalization of the proof of [FNV17]. We highlight that the core technical contribution in this work is dealing with *probabilistic* provers, once we add the laconicity property. Indeed, as a corollary of Theorem 1.1, we thus get that any laconic SHVZK (potentially with a probabilistic prover) can be turned into a 2-round deterministic prover SHVZK, and thus a 2-round predictable argument.

We additionally note that once we add the laconicity property, the prover messages are “1/poly-predictable”, and as such any laconic argument is “1/poly-weakly predictable argument” in the

⁴Even this result is incomparable to [BDRV18] since we can handle arbitrary $O(\log n)$ laconic provers; on the other hand, we require the proof system to satisfy the stronger notion of SHVZK, as opposed to HVZK.

⁵Goldreich and Oren showed that standard deterministic-prover ZK arguments cannot exist for languages outside of BPP.

⁶Roughly speaking, this strong WI property requires that the transcript only depends on the randomness of the Prover and the Verifier (and is independent of the witness the Prover uses).

terminology of [BIOW20]. Such weak predictability was proven useful in building full-fledged predictable arguments in [BIOW20] for the case of 1-element arguments in the generic group model (when the soundness error is negligible), but it is unknown how to amplify weak predictability more generally. Indeed, we here take a different approach.

1.2 Proof Outline: Earlier Work and the Key Idea

We here provide an outline of the proof of Theorem 1.1. We provide the construction of a WE for an NP language L given any laconic SHVZK argument for L . The converse direction (i.e., that WE for L implies laconic SHVZK) directly follows from a construction in [KMN⁺14].⁷ Towards this goal, let us first explain the approach of [BDV18], which originates already in [CS02].

1.2.1 Earlier Work: PKE from Laconic ZK

Cramer and Shoup’s [CS02] notion of an Smooth Projective Hash function (a.k.a a Hash Proof System) can be viewed as a 2-message HVZK proof system with a *deterministic* prover for a cryptographically hard language L . [CS02] show that such a proof system can be turned into an PKE by letting the public key be a randomly sampled YES-instance x (that is indistinguishable from a randomly sampled NO-instance), and letting the secret key be the witness for x . To encrypt a message, generate a simulated transcript (π_1, π_2) ⁸, send π_1 and m hidden using π_2 (e.g., one-time pad encrypted using randomness extracted from π_2 through some reconstructive extractor such as the Goldreich-Levin hard-core function [GL89]). Since the prover for the HVZK is deterministic, anyone having the secret key (i.e., the witness for x) can given π_1 recover π_2 (by simply running the honest prover strategy on x, w) and thus decrypt m . Secrecy, on the other hand, follows from the cryptographic hardness of L and the soundness of the proof system: in particular: (1) any attacker to the secrecy of the PKE can recover the prover message π_2 given a randomly sampled π_1 (according to the simulator strategy) leading to an accepting execution; (2) by the cryptographic hardness of L , the attacker will still do so when switching to randomly sampled false statements, which contradicts the soundness of the HVZK proof.

[BDV18] extend this method to work also when the prover is randomized as long as the total prover communication is small—more precisely, they require it to be of length $O(\log^{1/3} n)/r^{2/3}$ where r is the number of rounds of the protocol. Let us focus on the case of just 2-round protocols. Their key insight is that if the length of the prover message π_2 is short, then we can effectively make the prover deterministic following the approach of Valiant and Vazirani [VV86]: let the encryption include a *hash*, $h(\pi_2)$ of the message π_2 sampled by the encryptor, where h is an appropriate pairwise independent hash function. The decryptor can next continuously sample the honest prover to find a prover message that matches the hash. The tricky part here consists of showing that by appropriately selecting the length of the output of h , we can simultaneously guarantee that (a) the message π_2 is uniquely determined, yet (b) the hash $h(\pi_2)$ does not computationally reveal too much about π_2 to break the soundness of the protocol. Proving this is quite nontrivial.

1.2.2 Our Approach in a Nutshell

Our first observation is that the above approaches already satisfy the functionality of witness encryption: instead of sampling a random YES instances, simply use any YES-instance, and any receiver knowing a witness can decrypt. The problem is how to argue security. Indeed, the above approach

⁷Technically, there was only argued that the construction satisfies HVZK, but the construction trivially also satisfies the SHVZK property—indeed, essentially all standard HVZK protocols also satisfy the SHVZK property.

⁸Since it is a 2-message protocol, it must be the case that π_1 is the verifier message and π_2 is the prover message.

strongly relies on the fact that we can sample a random NO-instance that is indistinguishable from YES instances which no longer is the case when considering witness encryption. We here provide a very high-level explanation of how we deal with this issue, focusing on just 2-round protocols.

Warm-up (following [GGSW13, ABP15, FNV17]): Dealing with Deterministic Provers

We first show that for the case of 2-round protocols with a deterministic prover, if the protocol satisfies the *special* HVZK property, then the above approach can be extended to work even if the protocol is just an argument!

As mentioned above, it was already observed that the approach of Cramer-Shoup extends to yield witness encryption [GGSW13, ABP15] but that approach instead relied on a different insight and in particular the unconditional soundness conditions of Hash Proof Systems). Alternatively, as discussed above, Faonio, Nielsen, and Venturi [FNV17] show that the above approach directly work if the proof systems satisfies a predictability property.

To deal with just deterministic-prover SHVZK (which as noted above is implied by predictability), we generalize the approach of [FNV17] and modify the protocol as follows: To encrypt, first run the simulator (on randomly sampled random tape for the verifier) and check that the simulator outputs views that lead to the verifier accepting. If not, simply output $\perp$, and if yes, proceed just as before.⁹ Correctness follows just as before (due to the SHVZK property, the additional extra step can only disrupt things with negligible probability). Let us now turn to secrecy. If the simulator does not output accepting views, then secrecy trivially holds. Now, the key point is that due to the SHVZK property, if the simulator does output accepting views with high probability, then the verifier message is statistically close to an honest verifier message (since we are simply running the honest verifier on a fresh random tape).¹⁰ As a consequence, it follows that if some attacker is able to predict the prover response π_2 for a verifier message π_1 sampled by the SHVZK simulator, then this response π_2 would also convince an honest verifier, which contracts the soundness of the protocol, and from which we have that secrecy holds.

Dealing with Laconic Provers The above reasoning, however, does not apply to the protocol of [BDRV18]. Instead, to deal with the case of probabilistic, but laconic, provers, we take a different (and arguably simpler) approach. Another advantage is that this approach also works when the laconicity is $O(\log n)$ (as opposed to just $O(\log^{1/3} n)$). The key insight here is that due to the SHVZK and the laconic prover property, the encryptor can estimate the distribution of the prover message π_2 given a fixed verifier message π_1 as sampled by the simulator (with a fixed random tape r for the verifier so that the message π_1 remains fixed). The decryptor, on the other hand, can estimate the distribution of the honest prover response. Due to the HVZK property (and standard Chernoff bounds), these estimates should be close. We then complete the protocol by having the encryptor and decryptor engaging in a correlated sampling procedure to jointly sample π_2 and complete the rest of the protocol as before (and again including the simulator checking stage).

More Details on the Correlated Sampling Roughly speaking, the correlated sampling consist of having both entities computing a discretized (bucketed) version of the probability distribution (where the size of the bucket for each message π_2 is proportional to its probability mass) and letting the encryptor “throw a ball into a random bucket” to sample its message π_2 by picking a value $v \in [0, 1]$ and including v as part of the ciphertext. The important part to note here is that v is chosen independently of the distribution, so including it in the ciphertext does not leak anything

⁹This modification is not needed for [FNV17] as their “predictor” always finds accepting prover messages.

¹⁰We note that this property may not hold if the protocol only satisfies the HVZK, as opposed to SHVZK, property.

that would make it easier to find an accepting prover message.¹¹ The decryptor performs the same discretization and samples a random message using the same “ball” v . We can now show that when $x \in L$, then by the SHVZK property, the encryptor and the decryptor will agree with high probability. On the other hand, when $x \notin L$, if some attacker is able to (given v) predict the message π_2 , then effectively it must be able to sample π_2 from the simulator distribution, which (as previously argued) would also convince an honest verifier, contradicting the soundness.

We warn the reader that the above outline is brushing over many important details that are nontrivial to deal with—we provide a more detailed proof overview in the next section that explain them.

1.3 Detailed Proof Overview

We proceed to a more detailed proof overview of how we build a WE for L given a laconic SHVZK argument for L .

1.3.1 The Protocol

To formalize the above approach, it will be convenient to first define and construct a weak form of WE that next can be amplified into a “full-fledged” WE.

Weak Witness Encryption We start by introducing the weak form of WE, which we simply refer to as *weak witness encryption* (*weak WE*). Roughly speaking, we think of weak WE as a WE that for random message, and the secrecy property is weakened to only require the message to be hard to compute (as opposed to it being fully hidden). In more detail, a weak WE for an NP language L consists of two PPT algorithms Enc, Dec where $\text{Enc}(x)$ with a statement $x \in \{0, 1\}^n$ samples a ciphertext c together with a “message” m . The correctness condition requires that if $x \in L$, $\text{Dec}(x, w)$ where $w \in R_L(x)$ recovers m with probability $1 - \alpha$ (for some inverse polynomial $\alpha = \alpha(n)$). The secrecy/soundness condition—referred to as *soundness security* in analogy with the standard definition of WE—requires that if $x \notin L$, no attacker can compute m given c with probability more than $1 - n \cdot \alpha$. Note, weak WE weakens the standard WE definition in two ways: (1) by only requiring m to be (weakly) hard to compute, and (2) correctness only holds w.r.t. random messages. However, as we show, any weak WE implies a full-fledged WE (for the same language); this can be shown by appealing to the hardness amplification theorem for weakly verifiable puzzles [CHS05], the Goldreich-Levin theorem [GL89], and a standard “majority” trick to enable WE correctness amplification (see Section 4 for more details). We emphasize that the final WE we thus obtain may have an exponentially small correctness error (i.e. we do not get a WE with perfect correctness).

The Weak WE Construction We proceed to show how to construct a weak WE scheme for any NP-language L given any ℓ -laconic efficient-prover SHVZK argument (P, V) with simulator S where $\ell(n) = O(\log n)$. For the ease of presentation, we assume that (P, V) has completeness $2/3$ and soundness error $1/3$.¹² The reader is referred to Section 5 for a detailed description of our weak WE scheme.

The encryption algorithm $\text{Enc}(x)$, $x \in \{0, 1\}^n$, proceeds as follows.

¹¹This is a major difference with the approach of [BDRV18] where a lot of care had to be taken to ensure that the additional hash provided does not leak too much.

¹²We remark that we are able to handle much more general parameters where the gap between completeness and soundness error is only inverse polynomial.

1. *Test* whether S generates accepting views¹³ (with probability $1/2$) on x and uniformly sampled verifier random-tape r . If not, output $c = \perp$, $m \leftarrow \{0, 1\}^n$. Fix a random r .
2. *Sample* a random round i (in which P speaks) and sample a random transcript (prefix) π' of the first $i - 1$ rounds from $S(x, r)$.
3. *Draw* a “histogram” of the (simulated) next-message distribution given π' (by repeatedly sampling from $S(x, r)$).¹⁴ Let p_σ denote the estimated probability mass for message σ .
4. *Perform correlated sampling* by thinking of each σ occupying an interval (i.e., a “bucket”) in $[0, 1]$ of length p_σ , sampling random $v \leftarrow [0, 1]$ ¹⁵ (that acts as shared randomness), and σ^* is defined to be the element whose interval covers v (i.e., whose bucket contains the “ball” v). Output $c = (x, \pi', v)$ and $m = \sigma^*$.

The decryption algorithm $\text{Dec}(c, w)$ proceeds as follows.

1. *Interpret* c as $c = (x, \pi', v)$.
2. *Draw* a “histogram” of the (real) next-message distribution given π' (by repeatedly sampling from $(P(x, w), V(x))$).
3. *Perform correlated sampling* by sampling from the histogram as in Enc with the shared randomness v . Output the outcome of the sampling, denoted by τ^* .

1.3.2 Some Useful Notations and Observations

Before analyzing our schemes, it is instructive to remark here that it is without loss of generality to assume that for *every* statement x , the simulator S will generate views in which (a) the transcript satisfies laconicity and (b) the verifier messages are honestly generated, due to the SHVZK property.¹⁶

It will also be useful to introduce some notations and makes some simple claims about them.

Good Histograms and Heavy π' We will consider “good” histograms and “heavy” partial transcript π' , where a histogram is said to be *good* if it is “ ε -close” to the actual distribution (i.e., the estimated probability mass is within ε of the real probability mass for each element in the support), and π' is said to be *heavy* (or θ -heavy) (in the underlying distribution) if it is sampled (in the distribution) with probability $\geq \theta$. We pick both $\varepsilon = \varepsilon(n)$ and $\theta = \theta(n)$ to be some sufficiently small (but a-priori bounded) inverse polynomials in n . We rely on the following three claims on good histograms and heavy π' :

Claim 1: Heavy π' lead to good histograms: We first remark that any heavy π' leads to good histograms. Since in order to generate the histogram, we will sample from the distribution for an a-priori bounded sufficiently large ($\geq \text{poly}(n)/(\theta\varepsilon)^2$) number of times, by a standard Chernoff-type argument (and a Union bound over every element in the support), we have that with overwhelming probability, we will generate a good histogram given a heavy partial transcript π' .

¹³Recall that a view (of V in an interaction with P) contains a verifier random-tape r and a transcript, where a transcript consists of all messages exchanged by P and V . A view is *accepting* if the verifier accepts in the end.

¹⁴We sample for a sufficiently large (but a-priori polynomially bounded) number of times. We will elaborate more on this later.

¹⁵For simplicity, we can let v be a real number. We can always “discretize” the interval, paying an exponentially small sampling error.

¹⁶For any S , we can consider another simulator $S'(x, r)$ which truncates the prover messages to ℓ bits, and replaces the verifier messages by the ones that are consistent with r and x (by running $V_r(x)$ over the truncated transcript).

Claim 2: Good histograms approximate the prover’s next message: We next observe that good histograms are statistically close to the distribution of the prover’s next-message. This follows from the fact that (P, V) is ℓ -laconic, so the next-message distribution has support size at most 2^ℓ . Thus, the statistical distance between a good histogram and the actual distribution is at most $\varepsilon \cdot 2^\ell = \varepsilon \cdot \text{poly}(n)$.¹⁷

Claim 3: Random prefixes π' are heavy: Finally, we mention that with very high probability, a random π' will be heavy (in $\text{Enc}(x)$). Again, due to laconicity, we can show that the support size of the distribution of π' is at most 2^ℓ . By a Union bound over every element in the support, π' is not heavy with probability at most $\theta \cdot 2^\ell = \theta \cdot \text{poly}(n)$.

We are now ready to argue correctness and security of the scheme.

1.3.3 Correctness Analysis Overview

We proceed to show that σ^* will equal τ^* with very high probability¹⁸ ($\geq 1 - \alpha$) when $x \in L$ is a YES instance and $w \in R_L(x)$ is a witness of x . For any verifier random-tape r , let $\Pi_{\text{sim}}(r)$ denote the simulated transcript distribution in $S(x, r)$, and let $\Pi_{\text{real}}(r)$ denote the real transcript distribution in an interaction between $P(x, w)$ and $V_r(x)$. The SHVZK property of (P, V) guarantees that $\{R, \Pi_{\text{sim}}(R)\}$ is computationally indistinguishable from $\{R, \Pi_{\text{real}}(R)\}$ (against non-uniform attackers), where R denotes the uniform random distribution (over verifier random-tapes). To make the correctness proof modular, it will be convenient to pass through an (a-priori) stronger notion of SHVZK, which we refer to as *almost-every-randomness SHVZK*; this notion of SHVZK will require indistinguishability to for an overwhelming fraction of random tapes r for the verifier (when fixing the verifier’s random tape to r).¹⁹ This notion can be viewed as a relaxed form of the notion of a *semi-malicious verifier* [BGJ⁺13] (which, roughly speaking, requires simulation to succeed w.r.t. *all* random tapes for the honest verifier). Our proof proceeds in the following three steps.

Step 1: Dealing with a deterministic V : Let us start by assuming that V is deterministic; alternatively, we can consider the case that V uses a single fixed random-tape $r = 0^*$. Since (P, V) is laconic, it follows that $\{r, \Pi_{\text{real}}(r)\}$ is a distribution over only polynomially many elements. In this case, (non-uniformly secure) computational indistinguishability directly implies statistical indistinguishability (since we can hardcode the transcript that distinguishes between the two distributions the best).

Now, given that $\Pi_{\text{sim}}(r)$ and $\Pi_{\text{real}}(r)$ are statistically close²⁰, we first observe that the test in step 1 will be passed with overwhelming probability (since the real transcripts are accepted with probability $\geq 2/3$). Next, we can focus our attention to the partial transcripts π' that are heavy in $\text{Enc}(x)$ (since by Claim 3 this happens with very high probability). Notice that if π' is heavy in $\text{Enc}(x)$ (where the underlying distribution is $\Pi_{\text{sim}}(r)$), it must be also heavy in Dec (where the underlying distribution is $\Pi_{\text{real}}(r)$), due to the SHVZK property. Thus, we have that

- The next-message distribution given π' in $\Pi_{\text{sim}}(r)$ is statistically close to the next-message distribution given π' in $\Pi_{\text{real}}(r)$ (due to SHVZK);

¹⁷In this proof overview, we just consider the distance to be “sufficiently small” (and often ignore it) since we can make ε as much small as we want.

¹⁸We say that something happens with very high probability if it happens with probability $1 - 1/p(n)$ and we can make the polynomial p arbitrarily large.

¹⁹We remark that almost-every-randomness SHVZK trivially implies SHVZK; in the third step of our analysis below, we show that the converse actually also holds.

²⁰In the formal proof, we instead rely on the notion of max distance, but the proof proceeds roughly in the same way.

- And the two histograms generated in **Enc** and **Dec** will be good (by Claim 1).

Combining the above two facts (and also Claim 2), we have that the two histograms will be statistically close to each other. Since our correlated sampling procedure with the same shared randomness (from the statistically close histograms) will produce the same outcome with very high probability, we conclude that $\sigma^* = \tau^*$ with very high probability.

Step 2: Randomized V and breaking almost-every-randomness SHVZK: We turn to considering the case where the verifier random-tape is uniformly distributed (rather than a fixed string). We assume for contradiction that $\sigma^* \neq \tau^*$ happens with probability $> \alpha$. By a standard averaging argument, there exists at least an $O(\alpha)$ fraction of r , such that σ^* and τ^* disagree with probability $O(\alpha)$ when the verifier random-tape is fixed to be r . As argued above, for each such r , there exists a non-uniform distinguisher D_r that distinguishes between $\{r, \Pi_{\text{sim}}(r)\}$ and $\{r, \Pi_{\text{real}}(r)\}$ with advantage $\text{poly}(\alpha, 1/n)$. The (first) caveat is that the non-uniform advice used in D_r depends on r . This can be fixed by relying on the laconicity. Recall that the advice needed in D_r is just a transcript, which given r can be fully determined by the prover messages.²¹ Since (P, V) is ℓ -laconic, the advice needed in D_r is just ℓ bits. Using probabilistic arguments, we can show that there is a (single) non-uniform attacker D that, for at least an $O(\alpha/2^\ell)$ fraction of r , distinguishes between $\{r, \Pi_{\text{sim}}(r)\}$ and $\{r, \Pi_{\text{real}}(r)\}$ with advantage $\text{poly}(\alpha, 1/n)$. This concludes that (P, V) cannot satisfy almost-every-randomness SHVZK.

Step 3: From SHVZK to almost-every-randomness SHVZK: In the final step, we show that any attacker to almost-every-randomness SHVZK can be turned (with some additional non-uniformity) into an attacker to SHVZK for any protocol (P, V) with an efficient prover P , which yields the desired contradiction (and thus $\sigma^* = \tau^*$ must hold with probability $> \alpha$).

At first, let us explain why an attacker breaking almost-every-randomness SHVZK does not necessarily break SHVZK: The reason for this is that distinguishing advantage is *not* linear (due to the absolute sign), and thus averaging arguments will not work. To overcome this issue, the key observation is that $\Pi_{\text{sim}}(r)$ is efficiently samplable given r (which is included in the view), and $\Pi_{\text{real}}(r)$ is efficiently samplable given r and the witness w (due to the efficient prover). Given an almost-every-randomness SHVZK attacker D , we can additionally add w as non-uniform advice in D . For each r , D can now compute whether the absolute sign in the distinguishing advantage shall be “flipped” and flip it when needed (by estimating the probabilities that D outputs 1 on $(r, \Pi_{\text{sim}}(r))$ and on $(r, \Pi_{\text{real}}(r))$ up to inverse polynomial precision (in our case, $\text{poly}(\alpha, 1/n)$), and flipping the output of D if, for example, D outputs 1 and the former is smaller than the later).

1.3.4 Security Analysis Overview

We prove that if $x \notin L$, over random $(c, m) \leftarrow \text{Enc}(x)$, no non-uniform polynomial-time attacker A can compute m given c with probability $\geq 1 - \beta$, where $\beta = O(\frac{1}{2^\ell \gamma})$ and γ denotes the number of rounds in (P, V) . To do so, we rely on the soundness of (P, V) with respect to the NO instance x . (As we shall see very soon, our reduction is fully black-box, and thus if (P, V) is statistically sound (i.e., a proof system), then the weak WE we obtain will have statistical security.) We start by assuming for contradiction that there exists an attacker A that computes m given c with probability $\geq 1 - \beta$. Our goal is to use A to build a cheater prover P^* that breaks the soundness of (P, V) .

²¹As we shall argue, we can also instead use the witness w as non-uniform advice (since what we need is the “best” prover messages, which can be computed from w (again, relying on laconicity)).

The Malicious Prover P^* : We will construct a “malicious” prover P^* that has access to A when interacting with V . The malicious prover $P^{*,A}$, during the interaction, proceeds as follows.

In each round $i \in [\gamma]$, if V speaks in this round, $P^{*,A}$ just receives the message. If P speaks in this round, let $\pi_{[i-1]}$ denote the transcript of the previous $i-1$ rounds. $P^{*,A}$ simply samples random $v \leftarrow [0, 1]$, and sends across to V the output of $A(c)$ where $c = (x, \pi_{[i-1]}, v)$.

Analyzing the success probability of P^* : We proceed to arguing that $P^{*,A}$ convinces V . We consider any $x \notin L$ and such that the simulator passes the “simulator test” (since otherwise, A can only success with negligible probability as m is uniformly random n -bit string.)

We may further assume w.l.o.g that the simulator outputs accepting views with probability at least $5/12$ —by a Chernoff bound, the simulator can only pass the test with negligible probability otherwise.

Note that under these restriction, any attacker for the WE is able to sample the prover’s next message for simulated views given just the partial transcript. The problem, however, is that while we can show that the *transcript* generated is correctly distributed, it does not mean that the external verifier will be accepting these transcript. Of course, the simulator is generating accepting views, but the reason for this could potentially be that its simulated prover messages are correlated with the verifier’s random tape. We show that this cannot happen if A succeeds with high probability—intuitively, this follows since A ’s view is independent of the verifier randomness r (conditioned on the partial transcript it gets), so if A succeeds in predicting simulated next messages, then these messages also cannot depend on r .

To formalize this argument, we proceed by considering any fixed verifier randomness r ; we say that r is *good* if $A(c)$ computes m with probability $\geq 1 - O(\beta)$ given that the verifier random-tape used in Enc is fixed to r (denoted by Enc_r). It follows that r is good with probability at least $1 - 0.01$. Fix some good verifier-tape r . Let Π_{adv} denote the distribution of the real transcripts between $P^{*,A}$ and V_r , and let Π_{sim} denote the transcript distribution in $S(x, r)$.

We aim to show that the statistical distance between Π_{adv} and Π_{sim} is at most 0.01 . If so, combining this with the fact that $S(x, r)$ generates accepting transcripts with probability $5/12$, and that r is good with probability $1 - 0.01$, we conclude that $P^{*,A}$ convinces V with probability

$$\frac{5}{12} - 0.01 - 0.01 > \frac{1}{3}.$$

We move on to proving that Π_{adv} and Π_{sim} is statistically close, and this is proved in a round-by-round fashion.²² For any round $i \in [\gamma]$, let $\Pi_{\text{adv},[i]}$ (resp $\Pi_{\text{sim},[i]}$) denote the distribution over the first i round messages in Π_{adv} (resp in Π_{sim}). Let d_i denote the statistical distance between $\Pi_{\text{adv},[i]}$ and $\Pi_{\text{sim},[i]}$. Consider any round $i \in [\gamma]$. Suppose that V speaks in this round, then we have that d_i will not increase (from d_{i-1}). On the other hand, suppose P speaks in round i , we can focus our attention on partial transcripts $\pi_{[i-1]}$ that is $O(2^{-\ell}/\gamma)$ -heavy in $\Pi_{\text{sim},[i-1]}$ (which happens with probability $\geq 1 - 2^\ell \cdot O(2^{-\ell}/\gamma) \geq 1 - O(1/\gamma)$ by Claim 3). Since $\pi_{[i]}$ is $O(2^{-\ell}/\gamma)$ -heavy in Π_{sim} , we have that (i) the histogram (generated in $\text{Enc}_r(x)$ when $\pi' = \pi_{[i-1]}$) is good (by Claim 1) and (ii) $A(c) = m$ holds with probability at least

$$1 - O(\beta \cdot 2^\ell \gamma) \geq 1 - O\left(\frac{1}{\gamma}\right)$$

over $(c, m) \leftarrow \text{Enc}_r$ given that $\pi' = \pi_{[i-1]}$. Notice that m is distributed according to the histogram, which approximates the distribution $\Pi_{\text{sim},i}$ given $\Pi_{\text{sim},[i-1]} = \pi_{[i-1]}$ (if the histogram is good, by

²²In the formal proof, as mentioned before, we instead rely on the notion of max distance, but the proof proceeds roughly in the same way.

Claim 2). In addition, $A(c)$ is distributed the same as $\Pi_{\text{adv},i}$ given $\Pi_{\text{adv},[i-1]} = \pi_{[i-1]}$. Taking into account the probability that $\pi_{[i-1]}$ is not heavy, it follows that the statistical distance between $\Pi_{\text{adv},[i]}$ and $\Pi_{\text{sim},[i]}$ will increase by at most

$$O(1/\gamma) + O\left(\frac{1}{\gamma}\right) \leq 0.01 \frac{1}{\gamma}$$

from d_{i-1} (by choosing the constants in $O(\cdot)$ carefully). Combining this with the fact that (i) there are γ rounds and (ii) when $i = 0$, the statistical distance is 0, we conclude that Π_{adv} and Π_{sim} is at most

$$0.01 \frac{1}{\gamma} \cdot \gamma \leq 0.01.$$

2 Preliminaries

For an array of γ variables, $\pi = (\pi_1, \dots, \pi_\gamma)$, we let $\pi_{[i]} = (\pi_1, \dots, \pi_i)$ denote the (array of the) first i variables.

2.1 Interactive Protocol and Zero-Knowledge

We start by introducing the notions of *interactive proofs* [GMR88] and *arguments* [BCC88].

Interactive Proofs and Arguments Roughly speaking, for a language $L \in \text{NP}$, an *interactive proof* for L consists of two (interactive) PPT machines, a Prover P and, a Verifier V such that if $x \in L$, given any witness w , when $P(w)$ and V interact (on the common input x), V will accept. In addition, the, so-called, soundness condition requires that if $x \notin L$, then no computationally unbounded (malicious) Prover P^* can make the Verifier V to accept. We say that (P, V) is an *interactive argument* for L if it satisfies the same condition except that the soundness condition only hold w.r.t. computationally efficient malicious provers P^* (that can be implemented in non-uniform polynomial time).

We proceed to a formal definition. Given two interactive Turing machines, P, V , let $\langle P, V \rangle(x)$ denote the output of V in an interaction between P and V on the common input x . When P receives some additional auxiliary input a , we denote it as $\langle P(a), V \rangle(x)$. Let $\text{View}_V(\langle P, V \rangle(x))$ denote the view of V in the interaction, which includes the common input x , all the messages V sends and receives, and the internal randomness of V . We refer to all the messages exchanged in the interaction as *transcript*. We let V_r denote the machine V with its internal randomness fixed to be r .

Definition 2.1 (Interactive Proofs/Arguments). *A pair of (interactive) PPT machines (P, V) is said to be an interactive proof with completeness $c(\cdot)$ and soundness error $s(\cdot)$ for a language L with witness relation R_L if the following are satisfied.*

- **Completeness:** For any $x \in L$, $n = |x|$, any witness $w \in R_L(x)$,

$$\Pr[\langle P(w), V \rangle(x) = 1] \geq c(n)$$

- **Soundness:** For any $x \notin L$, $n = |x|$, any (computationally unbounded) machines P^* ,

$$\Pr[\langle P^*, V \rangle(x) = 1] \leq s(n)$$

In addition, we say that (P, V) is an interactive argument for L if the soundness condition only holds for all non-uniform polynomial time P^* and all sufficiently long x .

We will be focusing on interactive proofs/arguments where the prover communication complexity is limited; such proof system are referred to as *laconic proofs/arguments* [GH98, GVV02].

Definition 2.2 (Laconic Proofs/Argument). *An interactive proof/argument (P, V) is referred to as being ℓ -laconic if the total length of messages from P to V given a common input $x \in \{0, 1\}^n$ is upper bounded by $\ell(n)$. We say that (P, V) is simply laconic if it is ℓ -laconic for $\ell(n) = O(\log n)$.*

Zero-Knowledge In this work, we focus on the notion of *special honest verifier zero-knowledge (SHVZK)* [GMR89, CDS94]. An interactive protocol (P, V) is said to be *special honest verifier zero-knowledge* if there exists a PPT simulator S such that for any $x \in L$, the distribution $S(x, r)$ over a uniform random r is computationally indistinguishable to the view of V in a real interaction.

Definition 2.3 (Special honest verifier zero-knowledge zero-knowledge (SHVZK)). *We say that an interactive protocol (P, V) for a language L with witness relation R_L is special honest verifier zero-knowledge (SHVZK) if there exists a PPT simulator S such that for every non-uniform polynomial-time machine D , there exists a negligible function ν such that for all $x \in L$, $n = |x|$, every $w \in R_L(x)$, it holds that*

$$|\Pr[D(\text{View}_V(P(w), V)(x)) = 1] - \Pr[r \leftarrow \mathcal{U}_{q(n)}: D(S(x, r)) = 1]| \leq \nu(n)$$

where $q(\cdot)$ is a polynomial such that V uses at most $q(|x|)$ random coins on input x for all x .

2.2 Witness Encryption

We proceed to formalizing the notion of witness encryption (WE) [GGSW13].

Definition 2.4. A witness encryption (WE) scheme for an NP language L consists of two PPT algorithms, Enc and Dec , such that the following are satisfied:

- **Correctness:** For all $x \in L$, $n = |x|$, all witnesses w of x , any $b \in \{0, 1\}$:

$$\Pr[\text{Dec}(\text{Enc}(x, b), w) = b] \geq 1 - 2^{-n}$$

- **Soundness Security:** For every non-uniform polynomial-time machine D , there exists a negligible function ν such that for all $x \notin L$, $n = |x|$,

$$|\Pr[D(\text{Enc}(x, 0)) = 1] - \Pr[D(\text{Enc}(x, 1)) = 1]| \leq \nu(n)$$

We say that (Enc, Dec) is statistically secure if the two distributions in the soundness security guarantee are statistically indistinguishable. We say that (Enc, Dec) is a witness encryption with correctness $c(\cdot)$ if the probability in the correctness condition is at least $c(n)$ all input lengths n .

2.3 Computational Indistinguishability

We recall the definition of (computational) indistinguishability [GM84].

Definition 2.5. Two ensembles $\{A_{n,x}\}_{n \in \mathbb{N}, x \in S_n}$ and $\{B_{n,x}\}_{n \in \mathbb{N}, x \in S_n}$ are said to be computational $\mu(\cdot)$ -indistinguishable, if for every non-uniform machine D (the “distinguisher”), there exists a negligible function $\mu(\cdot)$, such that for every $n \in \mathbb{N}$, $x \in S_n$,

$$|\Pr[D(1^n, x, A_{n,x}) = 1] - \Pr[D(1^n, x, B_{n,x}) = 1]| < \mu(n)$$

3 Main Theorem and Some Corollaries

Our main theorem shows the equivalence between WE for an NP-language L and the existence of a laconic SHVZK argument for L . In fact, the theorem is a bit stronger.

Theorem 3.1 (Characterizing WE). *For any language $L \in \text{NP}$, the following are equivalent:*

- *There exists a WE (resp. a statistically-secure WE) for L .*
- *There exists an efficient-prover ℓ -laconic SHVZK argument (resp. proof) for L with completeness c and soundness error s , such that $c(n) - s(n) \geq 1/\delta(n)$ for some polynomial $\delta(n)$ and $\ell(n) = O(\log n)$.*
- *For every (efficiently computable) $\ell(n) \geq 1$, there exists an (deterministic)-efficient prover ℓ -laconic two-round statistical SHVZK argument (resp. proof) for L with completeness $1 - 2^{-n}$ and soundness error $2^{-\ell(n)} + \text{negl}(n)$.*

Proof: We show in Corollary 5.2 (stated and proven in Section 5) that the second item implies the first item. In Theorem 6.1 (stated and proven in Section 6), we show that the first item implies the last item. Finally, the last item trivially implies the second item. ■

Note that as a direct consequence of Theorem 3.1, we get the following results of independent interest regarding laconic SHVZK proofs:

Round-collapse: As shown by Babai-Moran [BM88], all constant-round interactive proof for a language L can be collapsed into 2-round proofs for L . But this transformation does not preserve zero-knowledge, or the efficient-prover property (and also does not apply to arguments), and there are important barriers indicating that this may be inherent [Wee06, Vad00]. As a consequence of Theorem 3.1, we get that if the protocol is laconic, then round collapse preserving both SHVZK and the efficient-prover property is possible.

Optimal soundness error: The standard method for amplifying soundness of an interactive protocols is through parallel repetition [GMR88, BM88, BIN97, PV07, HPWP10, Hai13], but this blows up the communication complexity. As an additional consequence of Theorem 3.1, we directly get that any language having a laconic SHVZK also has one where the soundness error is $2^{-\ell} + \text{negl}(n)$. This is optimal up to the $\text{negl}(n)$ term as HVZK for non-trivial languages cannot have soundness error 0 [GO94], and thus if a cheating proof exists, it can always be guessed with probability $2^{-\ell}$.

4 From Weak to (Strong) Witness Encryption

In this section, we define a weaker notion of witness encryption (weak WE) that will be useful for our purposes, and next demonstrate how any weak WE can be turned into a (standard) WE; subsequently, we will show how to build a weak WE.

Definition 4.1. *A weak witness encryption scheme (with correctness $c(\cdot)$ and hardness $\omega(\cdot)$) for an NP language L consists of two PPT algorithms, Enc and Dec, such that the following are satisfied:*

- *(Syntax) Enc takes input x and outputs two strings (CT, y) . Dec takes input CT , and outputs a string z .*
- **Correctness:** *For all $x \in L$, $n = |x|$, all witnesses w of x :*

$$\Pr[(CT, y) \leftarrow \text{Enc}(x) : \text{Dec}(CT, w) = y] \geq c(n)$$

- **Soundness Security:** For any non-uniform polynomial-time attacker A , for all sufficiently large $n \in \mathbb{N}$, all $x \notin L$, $|x| = n$,

$$\Pr[(CT, y) \leftarrow \text{Enc}(x): A(CT) \neq y] \geq \omega(n)$$

We say that (Enc, Dec) is statistically secure if the security requirement holds with respect to any time-unbounded attacker A .

We say that (Enc, Dec) is (overwhelmingly) hard if the security condition holds for $\omega(n) = 1 - \frac{1}{p(n)}$ for every polynomial p . We say that (Enc, Dec) is correct if the correctness condition holds for $c(n) = 1 - 2^{-n}$. We say that (Enc, Dec) is just a weak witness encryption if it is both correct and overwhelmingly hard.

Proposition 4.2. Assume that there exists a weak witness encryption with correctness $1 - \frac{1}{\alpha(n)}$ and hardness $\frac{1}{\beta(n)}$ for two polynomials α, β such that $\alpha(n) \geq 4n\beta$. Then, there exists a witness encryption.

Moreover, the above statement preserves the statistical security between the schemes.

Proof: This proposition follows from Lemma 4.3, 4.5, and 4.6. ■

Lemma 4.3. Assume that there exists a weak witness encryption with correctness $1 - \frac{1}{\alpha(n)}$ and hardness $\frac{1}{\beta(n)}$ for two polynomials α, β such that $\alpha(n) \geq 4n\beta$. Then, there exists a weak witness encryption with correctness $\frac{3}{4}$ that is overwhelmingly hard.

Moreover, the above statement preserves the statistical security.

Proof: Consider any weak witness encryption (Enc, Dec) with correctness $1 - \frac{1}{\alpha(n)}$ and hardness $\frac{1}{\beta(n)}$, $\alpha(n) \geq 4n\beta$. Consider a new scheme defined as running (Enc, Dec) in parallel for $n\beta(n)$ times. The correctness of the new scheme follows from a Union bound. To argue the hardness, we rely on the notion of *weakly verifiable puzzles*, defined in [CHS05]. Informally, a weakly verifiable puzzle is a cryptographic puzzle where the verifier also gets the randomness of puzzle maker.²³ Observe that given that $(CT, y) \leftarrow \text{Enc}(x)$, we can make a weakly verifiable puzzle such that breaking the puzzle is equivalent to computing y given CT . (Simply let Enc be the puzzle maker, and the verifier can check if the attacker computes y using Enc 's internal randomness.) Also notice that the proof in [CHS05] is black box, so their result still holds if the puzzle generator and the verifier share the statement x , $x \in \{0, 1\}^n$ (as opposed just 1^n). Thus, the hardness of the new scheme follows from [CHS05]. ■

Lemma 4.4 (Goldreich-Levin [GL89, Yao82] (c.f. [MP23])). There exist an oracle probabilistic polynomial-time algorithm R such that for all $n \in \mathbb{N}$, any distribution Q over $\{0, 1\}^n \times \{0, 1\}^*$, any $\varepsilon > 0$, any (probabilistic) functions $f : \{0, 1\}^n \rightarrow \{0, 1\}$ satisfying

$$\Pr_{(x,z) \leftarrow Q, r \leftarrow \{0,1\}^n} [D(z, r, x \cdot r) = 1] - \Pr_{(x,z) \leftarrow Q, r \leftarrow \{0,1\}^n, b \leftarrow \{0,1\}} [D(z, r, b) = 1] \geq \varepsilon$$

where $x \cdot r$ denote the inner product over $\mathcal{F}_2$, then

$$\Pr[(x, z) \leftarrow Q: R^D(1^{\varepsilon^{-1}}, z) = x] \geq \frac{\varepsilon^3}{8n}$$

²³Formally, a weakly verifiable puzzle consists of a puzzle generator G and a verifier V . G will output a puzzle p and some “check information” c (where G 's internal randomness could fit in). V accepts if an answer a is correct with p and c . The hardness of the puzzle requires that no efficient attacker can find an answer a given only the puzzle p .

Lemma 4.5. *If there exists a (overwhelmingly hard) weak witness encryption with correctness $\frac{3}{4}$, then there exists a witness encryption with correctness $\frac{3}{4}$.*

Moreover, the above statement preserves the statistical security.

Proof: Let $(\text{Enc}', \text{Dec}')$ be the weak witness encryption scheme with correctness $\frac{3}{4}$. Consider the following (Enc, Dec) :

- $\text{Enc}(x, b)$ runs $(CT', y) \leftarrow \text{Enc}'(x)$, samples $r \leftarrow \{0, 1\}^{|y|}$, and simply outputs $CT = (CT', r, (y \cdot r) \oplus b)$ (where $y \cdot r$ denotes the inner product over $\mathcal{F}_2$).
- $\text{Dec}(CT, w)$ receives $CT = (CT', r, b')$, computes $z = \text{Dec}'(CT', w)$, and outputs $(r \cdot z) \oplus b'$.

The correctness of (Enc, Dec) follows from the correctness of $(\text{Enc}', \text{Dec}')$ (and the fact that when $y = z$, $(r \cdot z) \oplus b' = b$). We turn to arguing security. We will consider information-theoretic adversaries (resp non-uniform polynomial time adversaries) if $(\text{Enc}', \text{Dec}')$ is statistically hard (resp computationally hard). Since $(\text{Enc}', \text{Dec}')$ is overwhelmingly hard, it follows that for all sufficiently large n , on any $x \notin L$, it is hard to find y given CT' . By the Goldreich-Levin Theorem [GL89] (c.f. the (contra-position of) Lemma 4.4), it follows that the following two ensembles are indistinguishable:

$$\{\text{Enc}(x, 0)\}_{n \in \mathbb{N}, x \notin L, |x|=n} \text{ and } \{\text{Enc}(x, b) : b \leftarrow \{0, 1\}\}_{n \in \mathbb{N}, x \notin L, |x|=n}$$

And thus the security of Enc follows from a standard hybrid argument. ■

Lemma 4.6. *If there exists a witness encryption with correctness $\frac{3}{4}$ (and negligible security error), then there exists a witness encryption.*

Moreover, the above statement preserves the statistical security.

Proof: Given a witness encryption with correctness $3/4$, we can simply encrypt the same bit (using independent randomness) for $O(n)$ times and output ciphertexts we obtain as encryption. To decrypt, we can decrypt each ciphertext and output the majority of the decrypted bits. The security follows from a standard hybrid argument. And the correctness follows from a Chernoff-type argument. ■

5 Weak Witness Encryption from Laconic ZK

We here show how to construct a weak WE from any laconic SHVZK argument.

Theorem 5.1. *Assume that ℓ -laconic SHVZK argument with completeness c and soundness error s for an NP language L exists, such that $c(n) - s(n) \geq 1/\delta(n)$ for some polynomial $\delta(n)$ and $\ell(n) = O(\log n)$. Then, there exists a weak witness encryption for L with correctness $1 - \frac{1}{\alpha(n)}$ and hardness $\frac{1}{\beta(n)}$, where α, β are polynomials such that $\alpha(n) \geq 4n\beta(n)$.*

Moreover, if the ℓ -laconic SHVZK protocol is a proof system, then the weak witness encryption scheme satisfies statistical security.

Using Proposition 4.2, we get the following corollary.

Corollary 5.2. *Assume that ℓ -laconic SHVZK protocol with completeness c and soundness error s for an NP language L exists, such that $c(n) - s(n) \geq 1/\delta(n)$ for some polynomial $\delta(n)$ and $\ell(n) = O(\log n)$. Then, there exists a witness encryption for L .*

Moreover, if the ℓ -laconic SHVZK protocol is a proof system (resp argument system), then the weak witness encryption scheme has statistical security (resp computational security).

Proof: Immediate from Theorem 5.1 and Proposition 4.2. ■

The rest of this section is dedicate to proving Theorem 5.1.

5.1 The Protocol

We proceed to describing the protocol.

Ingredients: An ℓ -laconic special honest verifier zero-knowledge protocol (P, V, S) with completeness c and soundness error s for L , $c(n) - s(n) \geq 1/\delta(n)$ for some polynomial $\delta(n)$, $\ell(n) = O(\log n)$.

We can assume without loss of generality that S generates views where (i) the “prover” is ℓ -laconic and (ii) the verifier messages are consistent with V and the verifier’s random tape in the view. (See Remark 5.4.)

Parameters $\beta(n) = 8 \cdot 2^{2\ell} \cdot (16\delta\ell)^3$, $\alpha(n) = 16n\beta$, $T(n) = (128 \cdot 2^{3\ell}\alpha^2)^2 \cdot n\ell$ where $\delta = \delta(n)$, $\ell = \ell(n)$.

Encryption: $\text{Enc}(x)$, $x \in \{0, 1\}^n$, proceeds as the follows.

1. *Test whether S generates accepting transcript.* Test whether the simulator S generates transcripts which the verifier V will accept. (See Figure 2.) If not, simply sample $y \leftarrow \{0, 1\}^n$, output $(CT = \perp, y)$, and halt.
2. *Sample a random prefix.* Sample a random string $r \in \{0, 1\}^{q(n)}$ (that serves as V ’s random tape) and sample a view of the verifier from $S(x, r)$. Let $\pi = (\pi_1, \dots, \pi_\gamma)$ denote the transcript. Sample random $i \in [\gamma]$ such that π_i is a prover message. Let $\text{NextMsg}_S(i, \pi_{[i-1]}, r)$ denote the distribution of the i -th round message in the transcript sampled by $S(x, r)$ conditioned on the first $i - 1$ messages being $\pi_{[i-1]}$.
3. *Estimate the next-message distribution.* Repeat sampling from $S(x, r)$ (for $T(n)$ times) to estimate the probability mass for each element in the distribution $\text{NextMsg}_S(i, \pi_{[i-1]}, r)$. (See Figure 1 for the detailed algorithm.) Let D'_S denote the estimated distribution (a histogram).
4. *Use a correlated sampling strategy.* Sample $v \in \{0, 1\}^n$ (which acts as the shared randomness). Use the canonical sampling algorithm **Samp** (defined in Figure 3) to sample from D'_S with randomness v . Let $\sigma^* = \text{Samp}(D'_S, v)$.
5. Output $CT = (x, i, \pi_{[i-1]}, v)$ and $y = \sigma^*$.

Decryption: $\text{Dec}(CT, w)$, proceeds as the follows.

1. Interpret $CT = (x, i, \pi_{[i-1]}, v)$. Let $n = |x|$. Consider $\pi_{[i-1]}$ as a partial transcript.
2. *Run P on the partial transcript.* Consider the following experiment: Let $P(x, w)$ interact, for i rounds, with an imaginary verifier V' where, in each round j , V' will simply send π_j as its message. Let $\text{NextMsg}_P(i, \pi_{[i-1]}, w)$ denote the distribution of the i -th message in the experiment conditioned on the first $i - 1$ messages being $\pi_{[i-1]}$.
3. *Estimate the next-message distribution.* Repeat the above experiment (for $T(n)$ times) to estimate the probability mass for each element in the distribution $\text{NextMsg}_P(i, \pi_{[i-1]}, w)$. (See Figure 1 for the detailed algorithm.) Let D'_P denote the estimated distribution (a histogram).
4. *Sample using the shared randomness.* Use the canonical sampling algorithm **Samp** (defined in Figure 3) to sample from D'_P with randomness v . Let $\tau^* = \text{Samp}(D'_P, \tau)$.
5. Finally, output $z = \tau^*$.

Estimating a Conditional Distribution

- Given access to a (joint) distribution $\Pi = (\Pi_1, \dots, \Pi_\gamma)$, an index $i \in [\gamma]$, a prefix $\pi_{[i-1]}$, and a parameter T , the goal is to empirically estimate the distribution of Π_i conditioned on $\Pi_{[i-1]} = \pi_{[i-1]}$.
- Sample π' from Π for T times. Count the number of times where $\pi'_{[i-1]} = \pi_{[i-1]}$ (denoted by cnt). Let S be the set of σ that appears as π'_i in some π' such that $\pi'_{[i-1]} = \pi_{[i-1]}$. For each $\sigma \in S$, count the number of times where $\pi'_{[i-1]} = \pi_{[i-1]}$ and $\pi'_i = \sigma$ (denoted by cnt_σ).
- Return a distribution (or a “histogram”) over S where each $\sigma \in S$ has probability $\text{cnt}_\sigma / \text{cnt}$.

Figure 1: The distribution estimating algorithm (that generates a histogram)

Testing the Simulator S

- Given input $x \in \{0, 1\}^n$, a simulator S , a verifier V , and parameters $c(n), s(n), \delta(n)$ such that $c(n) \geq s(n) + 1/\delta(n)$, the goal is to test whether S generates transcripts which V will accept with probability close to $c(n)$.
- Sample a random string $r \in \{0, 1\}^{q(n)}$ and a transcript π from $S(x, r)$. Repeat the sampling procedure for $10n\delta(n)^2$ times. For each sampled randomness r and transcript π , simulate the verifier $V(x)$ with random tape r as follows:
For each round i , if π_i is a prover message, we feed π_i to V ; if π_i is a verifier message, we ignore it. Finally, we check whether V accepts in the end. We refer to (r, π) as being “accepted” if V accepts in the end.
- If the fraction of accepted (r, π) is more than $s(n) + 1/(2\delta(n))$, then our algorithm returns YES. Otherwise, return NO.

Figure 2: The simulated transcripts testing algorithm

The Canonical Sampling Algorithm Samp

- Given the description of a distribution D over $(\sigma_1, \dots, \sigma_m)$ with probabilities $(p_1, \dots, p_m)$, and a string $v \in \{0, 1\}^n$ (that acts as the randomness), $\text{Samp}(D, v)$ sorts $((\sigma_1, p_1), \dots, (\sigma_m, p_m))$ in lexicographic order. Then, $\text{Samp}(D, v)$ finds the largest index $j \in [m]$ such that $\sum_{i \leq j} p_i \leq v/2^n$, and returns σ_j .

Figure 3: The canonical sampling algorithm

5.2 Analyzing the Construction

Lemma 5.3. *Assume that (P, V, S) is a ℓ -laconic SHVZK protocol with completeness c and soundness error s for an NP language L , $c(n) - s(n) \geq 1/\delta(n)$ for some polynomial $\delta(n)$, $\ell(n) = O(\log n)$. Then, (Enc, Dec) is a weak witness encryption for L with correctness $1 - \frac{2}{\alpha(n)}$ and hardness $\frac{1}{2\beta(n)}$, where α, β are polynomials (defined in the protocol) such that $\alpha(n) \geq 16n\beta(n)$.*

Moreover, if (P, V, S) is a proof system (resp argument system), then (Enc, Dec) has statistical security (resp computational security).

The proof of Theorem 5.1 immediately follows by Lemma 5.3.

Proof: [of Theorem 5.1] Immediate by Lemma 5.3. ■

Remark 5.4. We will focus our attention to simulator S that only generates views in which (i) the prover is ℓ -laconic and (ii) the verifier messages are consistent with V and the random tape. (Notice that if the transcript and the random tape are fixed, the verifier messages are determined.) We can assume without loss of generality that the simulator S satisfies these: If the prover is not laconic, we can simply remove all the prover messages after $\ell(n)$ bits. If verifier messages are not consistent, we can replace all verifier messages with the correct ones. Since (P, V, S) is zero-knowledge, both actions will only affect a negligible fraction of simulated transcripts (when running the protocol on a YES instance), and thus the simulator remains valid.

We define several algorithms (or notions) that will be useful in our proof. Define Enc' to be the algorithm Enc with the first step skipped (i.e., it does not perform the test). Let Enc'_r denote the algorithm Enc' with the verifier's random tape fixed to be r (as opposed to sampling from random).

For any fixed $n \in \mathbb{N}$, any instance $x \in \{0, 1\}^n$, any verifier's random tape $r \in \{0, 1\}^{q(n)}$, let $\Pi_{\text{sim}}(r)$ denote the distribution of simulated transcripts generated by $S(x, r)$. (We will simply denote it by Π_{sim} (and omit (r)) when r is fixed and clear from the context.) Let $\gamma = \gamma(n)$ denote the number of rounds in (P, V, S) . For any $i \in [\gamma]$, and transcript π , we say that a partial transcript $\pi_{[i]}$ is θ -heavy (in $\Pi_{\text{sim},[i]}(r)$) for some $\theta > 0$ if the probability weight of $\pi_{[i]}$ in the distribution $\Pi_{\text{sim},[i]}$ is at least θ .

For any two distributions D_1, D_2 , we define the *max distance* between D_1 and D_2 , $\Delta_{\max}(D_1, D_2)$, as $\max_{\omega \in D_1 \cup D_2} \{|\Pr[D_1 = \omega] - \Pr[D_2 = \omega]|\}$.

Lemma 5.5. If a partial transcript $\pi_{[i-1]}$ is θ -heavy in $\Pi_{\text{sim},[i-1]}(r)$, and $\text{NextMsg}_S(i, \pi_{[i-1]}, r)$ has support size $\leq 2^{\ell(n)}$, then with probability at least $1 - 2^{-n}$ (over internal randomness of Enc'_r),

$$\Delta_{\max}(\text{NextMsg}_S(i, \pi_{[i-1]}, r), D'_S) \leq \frac{1}{32 \cdot 2^{2\ell(n)} \alpha(n)}$$

where θ is picked to be $\frac{1}{4 \cdot 2^{\ell(n)} \alpha(n)}$.

Proof: Let $\varepsilon = \frac{1}{32 \cdot 2^{2\ell(n)} \alpha(n)}$. Let p denote the probability of $\pi_{[i-1]}$ in $\Pi_{\text{sim},[i-1]}(r)$ (and notice that $p \geq \theta$). Consider any σ in the support of $\text{NextMsg}_S(i, \pi_{[i-1]}, r)$, and let q denote the probability of σ in $\text{NextMsg}_S(i, \pi_{[i-1]}, r)$. By a standard Chernoff-type argument, it holds that with probability $\geq 1 - 2^{-n-\ell(n)}$ in the algorithm in Figure 1, it holds that (i) $|\frac{\text{cnt}_\sigma}{T} - p| \leq \theta\varepsilon/4$; and (ii) $|\frac{\text{cnt}_\sigma}{T} - p \cdot q| \leq \theta\varepsilon/4$. It follows that

$$\begin{aligned} \left| \frac{\text{cnt}_\sigma}{\text{cnt}} - q \right| &= \left| \frac{\text{cnt}_\sigma/T}{\text{cnt}/T} - q \right| \\ &\leq \max \left\{ \frac{pq + \theta\varepsilon/4}{p - \theta\varepsilon/4} - q, q - \frac{pq - \theta\varepsilon/4}{p + \theta\varepsilon/4} \right\} \\ &= \max \left\{ \frac{\theta\varepsilon/4 + q\theta\varepsilon/4}{p - \theta\varepsilon/4}, \frac{q\theta\varepsilon/4 + \theta\varepsilon/4}{p + \theta\varepsilon/4} \right\} \\ &\leq \frac{\theta\varepsilon/2}{\theta(1 - \varepsilon/4)} \\ &\leq \varepsilon. \end{aligned}$$

where the second inequality follows from $p \geq \theta$, and the last two inequalities follow from $q \leq 1$, $\varepsilon \leq 1$. By taking a Union bound, the probability that this holds for all σ is at least

$$1 - 2^{-n-\ell(n)} \cdot 2^{\ell(n)} \geq 1 - 2^{-n}$$

■

5.2.1 Security Analysis

Lemma 5.6. *Assume that S generates ℓ -laconic-prover transcripts, $\ell(n) = O(\log n)$. Then, there exists a PPT algorithm R such that the following holds. For any attacker A , any x , $n = |x|$, such that*

$$\Pr[(CT, y) \leftarrow \text{Enc}(x) : A(CT) = y \mid CT \neq \perp] \geq 1 - \frac{1}{\beta(n)}$$

Then,

$$\Pr[\langle R^A, V \rangle(x) = 1] \geq d(n) - \frac{1}{8\delta(n)}$$

where $d(n)$ is defined as

$$d(n) \stackrel{\text{def}}{=} \Pr[r \leftarrow \{0, 1\}^{q(n)}, (x, r, \pi) \leftarrow S(x, r) : (r, \pi) \text{ is accepted}]$$

(according to the notion of being accepted defined in Figure 2)

Proof: Consider a prover R^A that proceeds as follows. On input x , $x \in \{0, 1\}^n$, it interacts with a verifier V . For each round i , let $\pi_{[i-1]} = (\pi_1, \dots, \pi_{[i-1]})$ denote the current transcript. If the prover speaks in this round, R^A samples a random $v \leftarrow \{0, 1\}^n$. Next, R^A runs A on input $(x, i, \pi_{[i-1]}, v)$, and sends the output of A to the verifier. (If the verifier speaks in this round, R^A simply receives the message.)

We move on to arguing that when R^A interacts with V on x , it will make V accept with good enough probability. Recall that Enc' is defined as the algorithm Enc with its first step skipped, and Enc'_r is defined as the algorithm Enc' with its verifier's random tape fixed to be r . Notice that the algorithm Enc' behaves exactly the same as the algorithm Enc conditioned on $CT \neq \perp$. Thus, we have that

$$\Pr[(CT, y) \leftarrow \text{Enc}'(x) : A(CT) = y] \geq 1 - \frac{1}{\beta(n)}$$

Let G be the set of V 's random tapes such that

$$G = \{r \in \{0, 1\}^{q(n)} : \Pr[(CT, y) \leftarrow \text{Enc}'_r(x) : A(CT) = y] \geq 1 - \frac{16\delta(n)}{\beta(n)}\}$$

By a standard averaging argument, it follows that, over a random $r \in \{0, 1\}^{q(n)}$, $r \in G$ with probability at least $1 - \frac{1}{16\delta(n)}$.

Consider a random run of the interaction between R^A and V , with V 's randomness fixed to be r , $r \in G$. Let $\Pi_{\text{adv}} = \Pi_{\text{adv}}(r)$ denote the distribution of “real” transcript of the interaction between R^A (which does not know r) and V_r . Recall that $\Pi_{\text{sim}} = \Pi_{\text{sim}}(r)$ is defined to be the distribution of simulated transcript generated by $S(x, r)$. We are going to show that the max distance between Π_{adv} and Π_{sim} is at most

$$\frac{1}{16 \cdot 2^{\ell(n)} \delta(n)}$$

If this holds, since the support of $S(x, r)$ is of size at most $2^{\ell(n)}$, we conclude that the statistical distance between Π_{adv} and Π_{sim} is at most $\frac{1}{16\delta(n)}$. Combining with the fact that $r \in G$ with probability $1 - \frac{1}{16\delta(n)}$ (and that $(R, \Pi_{\text{sim}}(R))$ will be accepted with probability $d(n)$ (where R denote the uniform random distribution over $\{0, 1\}^{q(n)}$)), by taking a Union bound, we conclude that $\langle R^A, V \rangle(x)$ will accept with probability

$$d(n) - \frac{1}{16\delta(n)} - \frac{1}{16\delta(n)} = d(n) - \frac{1}{8\delta(n)}$$

which will conclude the Lemma.

Finally, we show that Π_{adv} and Π_{sim} are close in maximal distance. Let γ denote the number of rounds in the interaction, and note that $\gamma \leq 3\ell(n)$. Pick

$$\varepsilon = \varepsilon(n) \stackrel{\text{def}}{=} \frac{1}{16 \cdot 2^{\ell(n)} \delta(n)} \cdot \frac{1}{\gamma}$$

We claim that in each round i , the max distance between $\Pi_{\text{adv},[i]}$ and $\Pi_{\text{sim},[i]}$ is at most $i \cdot \varepsilon$. Plugging in $i = \gamma$, we have that $\Delta_{\text{max}}(\Pi_{\text{adv}}, \Pi_{\text{sim}}) \leq \varepsilon \gamma$. We proceed to proving our claim in a round-by-round fashion. We start with that in round 0, the max distance is 0.

Next, for any $i \in [\gamma]$, assume that $\Delta_{\text{max}}(\Pi_{\text{adv},[i-1]}, \Pi_{\text{sim},[i-1]}) \leq \varepsilon \cdot (i-1)$. In round i , if the verifier speaks in this round, then the message is completely determined by the previous messages in both Π_{adv} and Π_{sim} , and therefore $\Delta_{\text{max}}(\Pi_{\text{adv},[i]}, \Pi_{\text{sim},[i]}) \leq \varepsilon \cdot (i-1)$. (See Remark 5.4.) If the prover speaks in this round, then consider any partial transcript $\pi_{[i-1]}$ of the $i-1$ messages. We only need to consider $\pi_{[i-1]}$ whose probability weights in either $\Pi_{\text{adv},[i]}$ or $\Pi_{\text{sim},[i]}$ is at least $\varepsilon \cdot i$ (since otherwise it will never influence the max distance upper bound $\varepsilon \cdot i$ for round i). Then, we can argue that such $\pi_{[i-1]}$ has probability mass at least ε in both $\Pi_{\text{adv},[i-1]}$ and $\Pi_{\text{sim},[i-1]}$, since otherwise the difference in the probability mass will be larger than $\varepsilon \cdot (i-1)$. Given any such $\pi_{[i-1]}$, consider the following four distributions:

1. D_1 is defined to be $A((x, i, \pi_{[i-1]}, v))$ where $v \leftarrow \{0, 1\}^n$. Notice that this is exactly the distribution of $\Pi_{\text{adv},i}$ conditioned on $\Pi_{\text{adv},[i-1]} = \pi_{[i-1]}$.
2. D_2 is the distribution $\text{Samp}(D'_S, v)$ where $v \leftarrow \{0, 1\}^n$, where D'_S is an empirical estimation of $\text{NextMsg}_S(i, \pi_{[i-1]}, r)$. (We will specify our choice of D'_S later in the proof.)
3. D_3 is the distribution $\text{Samp}(\text{NextMsg}_S(i, \pi_{[i-1]}, r), v)$ where $v \leftarrow \{0, 1\}^n$.
4. D_4 is defined to be just $\text{NextMsg}_S(i, \pi_{[i-1]}, r)$. Notice that this distribution is identical to $\Pi_{\text{sim},i}$ conditioned on $\Pi_{\text{sim},[i-1]} = \pi_{[i-1]}$.

We will prove that the $\Delta_{\text{max}}(D_1, D_4) \leq \varepsilon$ in the remaining of this proof. Notice that if $\Delta_{\text{max}}(D_1, D_4) \leq \varepsilon$, it implies that for any π_i , the difference in the probability mass of $(\pi_{[i-1]}, \pi_i)$ in the distribution $\Pi_{\text{adv},[i]}$ and $\Pi_{\text{sim},[i]}$ is at most

$$\begin{aligned} & |\Pr[\Pi_{\text{adv},[i]} = (\pi_{[i-1]}, \pi_i)] - \Pr[\Pi_{\text{sim},[i]} = (\pi_{[i-1]}, \pi_i)]| \\ &= |\Pr[\Pi_{\text{adv},[i-1]} = \pi_{[i-1]}] \cdot \Pr[\Pi_{\text{adv},i} = \pi_i \mid \Pi_{\text{adv},[i-1]} = \pi_{[i-1]}] \\ &\quad - \Pr[\Pi_{\text{sim},[i-1]} = \pi_{[i-1]}] \cdot \Pr[\Pi_{\text{sim},i} = \pi_i \mid \Pi_{\text{sim},[i-1]} = \pi_{[i-1]}]| \\ &\leq |1 - (1 - (i-1)\varepsilon) \cdot (1 - \varepsilon)| \\ &\leq \varepsilon \cdot i \end{aligned}$$

By considering every $\pi_{[i-1]}$ and π_i , we can conclude that $\Delta_{\text{max}}(\Pi_{\text{adv},[i]}, \Pi_{\text{sim},[i]}) \leq \varepsilon \cdot i$.

Since the probability weight of $\pi_{[i-1]}$ in the distribution $\Pi_{\text{sim},[i-1]}$ is at least ε , it follows that Enc'_r will sample $\pi_{[i-1]}$ in CT with probability at least $\frac{\varepsilon}{\gamma}$ (since it will hit $\pi_{[i-1]}$ with probability ε and i with probability $\frac{1}{\gamma}$). Recall that $r \in G$, and thus we have that A computes y with probability $\geq 1 - \frac{16\delta(n)}{\beta(n)}$ over Enc'_r . Notice that in Enc'_r , y equals to $\text{Samp}(D'_S, v)$ for a random $v \leftarrow \{0, 1\}^n$. By a standard averaging argument, it follows that with probability at least

$$1 - \frac{16\delta(n)}{\beta(n)} \cdot \frac{\gamma}{\varepsilon}$$

over random $v \leftarrow \{0, 1\}^n$ and the internal randomness (for obtaining D'_S from $\text{NextMsg}_S(i, \pi_{[i-1]}, r)$), we have that

$$A((x, i, \pi_{[i-1]}, v)) = \text{Samp}(D'_S, v)$$

Let E denote the event where $\Delta_{\max}(D'_S, \text{NextMsg}(i, \pi_{[i-1]}, r)) \leq \varepsilon/8$. Recall that the probability weight of $\pi_{[i-1]}$ in $\Pi_{\text{sim}, [i-1]}$ is at least ε , and therefore $\pi_{[i-1]}$ is ε -heavy in $\Pi_{\text{sim}, [i-1]}$. Observe that

$$\varepsilon \geq \frac{1}{4 \cdot 2^{\ell(n)} \alpha(n)}, \quad \frac{1}{32 \cdot 2^{2\ell(n)} \alpha(n)} \leq \varepsilon/8$$

By Lemma 5.5, the event E happens with probability $\geq 1 - 2^{-n}$. Relying on the standard probabilistic argument (and a Union bound), there exists $D_S^* \in E$ such that with probability at least

$$1 - \frac{16\delta(n) \cdot \gamma/\varepsilon}{\beta(n)} - 2^{-n}$$

over $v \leftarrow \{0, 1\}^n$, it holds that

$$A((x, i, \pi_{[i-1]}, v)) = \text{Samp}(D_S^*, v)$$

Fix D'_S in the definition of D_2 to be D_S^* . It follows from the definition of **Samp** (and a Union bound), we have that

$$\Delta_{\max}(D_1, D_2) \leq \frac{16\delta(n) \cdot \gamma/\varepsilon}{\beta(n)} + 2^{-n} \leq \frac{(16\delta(n)\ell(n))^2 2^{\ell(n)}}{\beta(n)} + 2^{-n} \leq \varepsilon/8 + \varepsilon/8 = \varepsilon/4$$

(since ε is inverse polynomial in n and $2^{-n} \leq \varepsilon/8$)

Since $D_S^* \in E$, it holds that $\Delta_{\max}(D_S^*, \text{NextMsg}(i, \pi_{[i-1]}, r)) \leq \varepsilon/8$. It follows from the definition of the canonical sampling algorithm **Samp** that

$$\Delta_{\max}(D_2, D_3) \leq \Delta_{\max}(D_S^*, \text{NextMsg}(i, \pi_{[i-1]}, r)) + 2 \cdot 2^{-n} \leq \varepsilon/2$$

and

$$\Delta_{\max}(D_3, D_4) \leq \Delta_{\max}(\text{NextMsg}(i, \pi_{[i-1]}, r), \text{NextMsg}(i, \pi_{[i-1]}, r)) + 2^{-n} = 2^{-n} \leq \varepsilon/8$$

Finally, we conclude that

$$\Delta_{\max}(D_1, D_4) \leq \Delta_{\max}(D_1, D_2) + \Delta_{\max}(D_2, D_3) + \Delta_{\max}(D_3, D_4) \leq \varepsilon$$

■

5.2.2 Correctness Analysis

Consider any sufficiently large $n \in \mathbb{N}$, any $x \in \{0, 1\}^n$, $x \in L$, any $w \in R_L(x)$, and any verifier's random tape $r \in \{0, 1\}^{q(n)}$.

We define $\Pi_{\text{real}} = \Pi_{\text{real}}(r, w)$ as the distribution of transcripts in a real interaction $\langle P(w), V_r \rangle(x)$ between the prover $P(w)$ and the verifier V_r on the common input x . (And, we omit (r, w) when they are fixed and clear from the context.) Recall that γ denotes the number of rounds. For each $i \in [\gamma]$, any transcript π , we can define the partial transcript $\pi_{[i]}$ as being heavy in $\Pi_{\text{real}, [i]}$ in a similar way: We say that a partial transcript $\pi_{[i]}$ is θ -heavy in $\Pi_{\text{real}, [i]}$, for some $\theta > 0$, if the probability weight of $\pi_{[i]}$ in $\Pi_{\text{real}, [i]}$ is at least θ .

Lemma 5.7. *If a partial transcript $\pi_{[i-1]}$ is θ -heavy in $\Pi_{\text{real},[i-1]}(r, w)$, and $\text{NextMsg}_P(i, \pi_{[i-1]}, w)$ has support size $\leq 2^{\ell(n)}$, then with probability at least $1 - 2^{-n}$ (over internal randomness of Dec),*

$$\Delta_{\max}(\text{NextMsg}_P(i, \pi_{[i-1]}, w), D'_P) \leq \frac{1}{32 \cdot 2^{\ell(n)} \alpha(n)}$$

where θ is picked to be $\frac{1}{4 \cdot 2^{\ell(n)} \alpha(n)}$

Proof: This Lemma essentially follows from the proof of Lemma 5.5 (by replacing Π_{sim} with Π_{real} , $\text{NextMsg}_S(i, \pi_{[i-1]}, r)$ with $\text{NextMsg}_P(i, \pi_{[i-1]}, w)$, and D'_S with D'_P). ■

We rely on the following Lemma to show the correctness of our correlated sampling strategy.

Lemma 5.8. *For any two distributions D_1, D_2 with the union of support containing m elements and max distance at most ε , it holds that*

$$\Pr[v \leftarrow \{0, 1\}^n : \text{Samp}(D_1, v) = \text{Samp}(D_2, v)] \geq 1 - m^2 \varepsilon - m 2^{-n}$$

Proof: Let $(\sigma_1, \dots, \sigma_m)$ denote the elements in the (union of the) support of the two distributions (in lexicographic order). Let $(p_1, \dots, p_m)$ denote their probabilities in D_1 , and $(q_1, \dots, q_m)$ denote their probabilities in D_2 . Since D_1 and D_2 are ε -close to each other in max distance, $|p_i - q_i| \leq \varepsilon$ for every $i \in [m]$.

For any $i \in [m]$, let I_1 denote the interval $[\sum_{j < i} p_j, \sum_{j \leq i} p_j)$ and I_2 denote the interval $[\sum_{j < i} q_j, \sum_{j \leq i} q_j)$. Notice that if $v/2^n$ is inside both intervals, $\text{Samp}(D_1, v)$ and $\text{Samp}(D_2, v)$ will simultaneously output σ_i . Also observe that I_2 is at most ε shorter than I_1 , and can only be “shifted away” from I_1 by $(i-1) \cdot \varepsilon$. Thus,

$$|I_1 \cap I_2| \geq p_i - \varepsilon - 2(i-1)\varepsilon$$

In addition, $v/2^n \in I_1 \cap I_2$ with probability at least $p_i - \varepsilon - 2(i-1)\varepsilon - 2^{-n}$.

Finally, the Lemma is concluded by taking summation over all $i \in [m]$. ■

Next, we show that the SHVZK property implies that $\Pi_{\text{sim}}(r)$ and $\Pi_{\text{real}}(r, w)$ is close in max distance for (any w and) a large fraction of r .

Lemma 5.9. *Assume that (P, V, S) is an ℓ -laconic SHVZK protocol for an NP language L (with witness set $R_L(\cdot)$), $\ell(n) = O(\log n)$. For any polynomial $p_1(\cdot)$ and $p_2(\cdot)$, for any sufficiently large $n \in \mathbb{N}$, $|x| = n$, any $w \in R_L(x)$, with probability $\geq 1 - \frac{1}{p_1(n)}$ over $r \in \{0, 1\}^{q(n)}$, we have that*

$$\Delta_{\max}(\Pi_{\text{real}}(r, w), \Pi_{\text{sim}}(r)) \leq \frac{1}{p_2(n)}$$

Proof: For any transcript π , let σ denote all prover messages in π (referred to as *prover response*). We notice that a real transcript is determined by the prover response together with the verifier random tape. We can assume this is true also for simulated transcripts (see Remark 5.4). Thus, for any verifier random tape r and prover response σ , let $f(\sigma, r)$ denote the transcript uniquely determined by σ, r .

We next claim that for any sufficiently large $n \in \mathbb{N}$, any $x \in L$, $|x| = n$, any $w \in R_L(x)$, any prover response $\sigma \in \{0, 1\}^{\ell(n)}$, it holds that with probability $\geq 1 - \frac{1}{2^{\ell(n)} p_1(n)}$ over $r \leftarrow \{0, 1\}^{q(n)}$,

$$|\Pr[\pi \leftarrow \Pi_{\text{real}}(r, w) : \pi = f(\sigma, r)] - \Pr[\pi \leftarrow \Pi_{\text{sim}}(r) : \pi = f(\sigma, r)]| \leq \frac{1}{p_2(n)}$$

If the claim holds, by a Union bound over all prover response $\sigma \in \{0, 1\}^{\ell(n)}$, the Lemma will be proved.

We proceed to proving the above claim. Suppose for contradiction that there exist infinitely many $n \in \mathbb{N}$, $x \in L$, $w \in R_L(x)$, prover response σ , such that with probability $\geq \frac{1}{2^{\ell(n)} p_1(n)}$ over $r \leftarrow \{0, 1\}^{q(n)}$,

$$|\Pr[\pi \leftarrow \Pi_{\text{real}}(r, w) : \pi = f(\sigma, r)] - \Pr[\pi \leftarrow \Pi_{\text{sim}}(r) : \pi = f(\sigma, r)]| > \frac{1}{p_2(n)}$$

We will construct a non-uniform polynomial-time attacker A that breaks the indistinguishability between $\text{View}_V(P(w), V_R)(x)$ and $S(x, R)$ (where R denotes the random variable over r). Consider any such n, x, w, σ (that the above holds), and let A hardcode all of them. In addition, notice that either (1) with probability $\geq \frac{1}{2 \cdot 2^{\ell(n)} p_1(n)}$, the above inequality holds and the value inside the absolute sign is positive; or (2) with probability $\frac{1}{2 \cdot 2^{\ell(n)} p_1(n)}$, the inequality holds and the value is negative. Let b be a bit where $b = 0$ if the former is the case (and $b = 1$ if the latter is the case). Let A also hardcode the bit b . Pick $\varepsilon = \varepsilon(n) = \frac{1}{4 \cdot 2^{\ell(n)} p_1(n) p_2(n)}$.

Our distinguisher A proceeds as follows. On input a (real or simulated) view which contains x , a verifier random tape r , and a transcript π , check if x is the statement it hardcodes. (If not, it simply outputs $\perp$.) Then, it samples from $\Pi_{\text{real}}(r, w)$ (by running $\text{View}_V(P(w), V_r)(x)$) for $L(n) = 64n\varepsilon^{-2}$ times, and computes the empirical estimation of the probability that a real interaction produces a transcript that matches the prover response σ (and denote the estimation by s_1). In addition, it samples from $\Pi_{\text{sim}}(r)$ (by running $S(x, r)$) for $L(n)$ times, and computes the empirical estimation of the probability that a simulated transcript matches σ (and denote the estimation by s_2). Finally, it checks whether π matches the prover response σ . If not, it simply outputs 0. If so, A outputs $b \oplus 1$ if $s_1 \geq s_2$ (or b otherwise).

We next present the proof for $b = 0$ (which can be easily adapted to a proof for $b = 1$). Let $p_{\text{real},r}(\pi)$ (resp $p_{\text{sim},r}(\pi)$) denote the probability of π for $\pi \leftarrow \Pi_{\text{real}}(r, w)$ (resp $\Pi_{\text{sim}}(r)$). Let E denote the set of $r \in \{0, 1\}^{q(n)}$ for which $|p_{\text{real},r}(f(\sigma, r)) - p_{\text{sim},r}(f(\sigma, r))| \geq \varepsilon$. It follows that the distinguishing advantage of A between $\text{View}_V(P(w), V_r)(x)$ and $S(x, r)$ is at least

$$\begin{aligned} & \left| \Pr[r \leftarrow \mathcal{U}_{q(n)}, \pi \leftarrow \Pi_{\text{real}}(r, w) : A(x, r, \pi) = 1] - \Pr[r \leftarrow \mathcal{U}_{q(n)}, \pi \leftarrow \Pi_{\text{sim}}(r) : A(x, r, \pi) = 1] \right| \\ &= \left| \mathbb{E}_{r \leftarrow \mathcal{U}_{q(n)}, \pi = f(\sigma, r)} [p_{\text{real},r}(\pi) \cdot \Pr[A(x, r, \pi) = 1]] - \mathbb{E}_{r \leftarrow \mathcal{U}_{q(n)}, \pi = f(\sigma, r)} [p_{\text{sim},r}(\pi) \cdot \Pr[A(x, r, \pi) = 1]] \right| \\ &\geq \left| \Pr[E] \cdot \mathbb{E}_{r \leftarrow \mathcal{U}_{q(n)} | E, \pi = f(\sigma, r)} [(p_{\text{real},r}(\pi) - p_{\text{sim},r}(\pi)) \cdot \Pr[A(x, r, \pi) = 1]] \right. \\ &\quad \left. - \Pr[\neg E] \cdot \mathbb{E}_{r \leftarrow \mathcal{U}_{q(n)} | \neg E, \pi = f(\sigma, r)} [|p_{\text{real},r}(\pi) - p_{\text{sim},r}(\pi)| \cdot \Pr[A(x, r, \pi) = 1]] \right| \\ &\geq \left| \Pr[E] \cdot \mathbb{E}_{r \leftarrow \mathcal{U}_{q(n)} | E, \pi = f(\sigma, r)} [(p_{\text{real},r}(\pi) - p_{\text{sim},r}(\pi)) \cdot \Pr[A(x, r, \pi) = 1]] \right| - \varepsilon \\ &\geq \left| \Pr[E] \cdot \mathbb{E}_{r \leftarrow \mathcal{U}_{q(n)} | E, \pi = f(\sigma, r)} [(p_{\text{real},r}(\pi) - p_{\text{sim},r}(\pi)) \cdot \mathbb{I}_{[p_{\text{real}}(\pi) > p_{\text{sim}}(\pi)]}] \right| - 2^{-n} - \varepsilon \\ &\geq \Pr_{r \leftarrow \mathcal{U}_{q(n)}, \pi = f(\sigma, r)} \left[p_{\text{real},r}(\pi) - p_{\text{sim},r}(\pi) > \frac{1}{p_2(n)} \right] \cdot \frac{1}{p_2(n)} - 2^{-n} - \varepsilon \\ &\geq \frac{1}{2 \cdot 2^{\ell(n)} p_1(n)} \cdot \frac{1}{p_2(n)} - 2^{-n} - \varepsilon \geq \varepsilon/2. \end{aligned}$$

where the first equation follows from the fact that A will only output 1 when $\pi = f(\sigma, r)$; and the first inequality follows from the triangle inequality of the absolute sign; the second inequality holds since for all $r \notin E$, $|p_{\text{real},r}(f(\sigma, r)) - p_{\text{sim},r}(f(\sigma, r))| < \varepsilon$; the third inequality follows from the fact that when $|p_{\text{real},r}(\pi) - p_{\text{sim},r}(\pi)| \geq \varepsilon$, $A(x, r, \pi) = 1$ if and only if $p_{\text{real}}(\pi) > p_{\text{sim}}(\pi)$ except with probability 2^{-2n} (which follows from a standard Chernoff-type argument) (where $\mathbb{I}$ denotes the indicator function); the fourth inequality is trivially true; the second last one follows from $b = 0$; and the last one holds due to our choice of parameters.

Finally, notice that A breaks the zero-knowledge property of (P, V, S) for infinitely many n , which is a contradiction. ■

Finally, we are ready to show the correctness of our scheme.

Lemma 5.10. *Assume that P is ℓ -laconic, $\ell(n) = O(\log n)$. Then, for all $n \in \mathbb{N}$, any $x \in \{0, 1\}^n$, $x \in L$, $w \in R_L(x)$, if with probability $\geq 1 - \frac{1}{\alpha(n)}$ over $r \leftarrow \{0, 1\}^{q(n)}$,*

$$\Delta_{\max}(\Pi_{\text{real}}(r, w), \Pi_{\text{sim}}(r)) \leq \frac{1}{128 \cdot 2^{3\ell(n)} \alpha(n)^2} \quad (1)$$

then, we have that

$$\Pr[(CT, y) \leftarrow \text{Enc}(x) : \text{Dec}(CT, w) = y \mid CT \neq \perp] \geq 1 - \frac{2}{\alpha(n)}$$

Proof: Recall that the algorithm Enc' is defined as Enc with the first step skipped, and thus the distribution of $(CT, y) \leftarrow \text{Enc}'(x)$ is identical to $(CT, y) \leftarrow \text{Enc}(x)$ given that $CT \neq \perp$. This allows us to focus our attention on Enc' and we will show that the probability that $\text{Dec}(CT, w) = y$ is at most $1 - \frac{1}{\alpha(n)}$ for $(CT, y) \leftarrow \text{Enc}'(x)$.

Consider any n , any instance $x \in L$, any $w \in R_L(x)$. We say that a verifier random tape $r \in \{0, 1\}^{q(n)}$ is *good* if Inequality 1 holds with respect to r . It follows that a random r is good with probability $\geq 1 - \frac{1}{\alpha(n)}$. Next, consider a good random tape $r \in \{0, 1\}^{q(n)}$, and a random sample $((x, i, \pi_{[i-1]}, v), y)$ from $\text{Enc}'_r(x)$. Notice that in Enc'_r , $\pi_{[i-1]}$ is sampled from $\Pi_{\text{sim}, [i]}(r)$. Let E denote the event that $\pi_{[i-1]}$ is θ -heavy in $\Pi_{\text{sim}, [i]}(r)$ where θ is picked to be

$$\theta = \frac{1}{2 \cdot 2^{\ell(n)} \alpha(n)}$$

Since S only generates transcripts in which the prover is ℓ -laconic (see Remark 5.4) and V 's random tape is fixed in $\Pi_{\text{sim}}(r)$, it follows that the support size of $\Pi_{\text{sim}, [i]}(r)$ is at most $2^{\ell(n)}$. By a Union bound, we conclude that the probability that E happens is at least

$$1 - 2^{\ell(n)} \cdot \theta \geq 1 - 2^{\ell(n)} \cdot \frac{1}{2 \cdot 2^{\ell(n)} \alpha(n)} \geq 1 - \frac{1}{2\alpha(n)}$$

We will show that conditioned on the event $\pi_{[i-1]} \in E$ and r is good, with probability at least $1 - \frac{1}{2\alpha(n)}$ over $(CT, y) \leftarrow \text{Enc}'_r(x), z \leftarrow \text{Dec}(CT, w)$, it holds that

$$y = z$$

If this is the case, combining with that E happens with high probability ($\geq 1 - \frac{1}{2\alpha(n)}$) and r is good with probability $\geq 1 - \frac{1}{\alpha(n)}$, by taking a Union bound, we have that $y = \text{Dec}(CT, w)$ with probability

$$1 - 2 \cdot \frac{1}{2\alpha(n)} - \frac{1}{\alpha(n)} \geq 1 - \frac{2}{\alpha(n)}$$

which will conclude the Lemma. We move on to proving the above claim. Consider any fixed r that is good. Fix $\Pi_{\text{sim}} = \Pi_{\text{sim}}(r)$ and $\Pi_{\text{real}} = \Pi_{\text{real}}(r, w)$. the following 4 distributions.

1. D_1 is defined to be D'_S (where D'_S is as in $\text{Enc}'_r(x)$). (D'_S is a random variable over “histograms” and we view histograms also as distributions they define.) Notice that y is identically distributed as $\text{Samp}(D'_S, v)$ for $v \leftarrow \{0, 1\}^n$.

2. D_2 is the distribution $\text{NextMsg}_S(i, \pi_{[i-1]}, r)$.
3. D_3 is the distribution $\text{NextMsg}_P(i, \pi_{[i-1]}, w)$.
4. D_4 is defined to be D'_P (where D'_P is as in Dec). (D'_P is a random variable over “histograms”.)
Notice that z is identically distributed as $\text{Samp}(D'_P, v)$ for $v \leftarrow \{0, 1\}^n$.

We rely on the following claims on the max distance between the distributions. Before stating and proving the claims, we here make a useful observation (referred to as *the support size observation*) that all of $D_1, \dots, D_4$ are of support size at most $2^{\ell(n)}$. $D_2 = \text{NextMsg}_S(i, \pi_{[i-1]}, r)$ has support size $\leq 2^{\ell(n)}$ since S only generate transcripts in which the prover is ℓ -laconic (see Remark 5.4). $D_3 = \text{NextMsg}_P(i, \pi_{[i-1]}, w)$ has support size $\leq 2^{\ell(n)}$ since P is ℓ -laconic. D_1, D_4 all have small support size since they are estimated version of D_2, D_3 . Pick

$$\varepsilon = \frac{1}{32 \cdot 2^{2\ell(n)} \alpha(n)}$$

Claim 1. *With probability at least $1 - 2^{-n}$, $\Delta_{\max}(D_1, D_2) \leq \varepsilon$.*

Proof: Since the event E holds, we know that $\pi_{[i-1]}$ is θ -heavy in $\Pi_{\text{sim}, [i-1]}$, $\theta \geq \frac{1}{4 \cdot 2^{\ell(n)} \alpha(n)}$. The claim follows from Lemma 5.5 (together with the support size observation). ■

Claim 2. *With probability at least $1 - 2^{-n}$, $\Delta_{\max}(D_3, D_4) \leq \varepsilon$.*

Proof: Since the event E holds, we know that $\pi_{[i-1]}$ is θ -heavy in $\Pi_{\text{sim}, [i-1]}$. In addition, the random tape r we are using is also good, it follows that $\Delta_{\max}(\Pi_{\text{sim}}, \Pi_{\text{real}})$ is at most

$$\frac{1}{128 \cdot 2^{3\ell(n)} \alpha(n)^2} \leq \theta/2$$

Thus, $\pi_{[i-1]}$ is also $(\theta/2)$ -heavy in $\Pi_{\text{real}, [i-1]}$.

Given that $\pi_{[i-1]}$ is also $(\theta/2)$ -heavy in $\Pi_{\text{real}, [i-1]}$, the claim follows from Lemma 5.7 (together with the support size observation). ■

Claim 3. $\Delta_{\max}(D_2, D_3) \leq \varepsilon$.

Proof: Let p denote the probability of $\pi_{[i-1]}$ in $\Pi_{\text{sim}, [i-1]}$, and let q denote the probability of $\pi_{[i-1]}$ in $\Pi_{\text{real}, [i-1]}$. Since the event E holds, $p \geq \theta$. In addition, the random tape r we are using is also good, it follows that $\Delta_{\max}(\Pi_{\text{sim}}, \Pi_{\text{real}})$ is at most

$$\frac{1}{128 \cdot 2^{3\ell(n)} \alpha(n)^2} = \theta\varepsilon/2$$

We have that $|p - q| < \theta\varepsilon/2$.

Notice that the distribution D_2 is the same as $\Pi_{\text{sim}, i}$ given $\Pi_{\text{sim}, [i-1]} = \pi_{[i-1]}$. And the distribution D_3 is identical to $\Pi_{\text{real}, i}$ given $\Pi_{\text{real}, [i-1]} = \pi_{[i-1]}$. Suppose for contradiction that there exists a message σ such that $|\Pr[D_2 = \sigma] - \Pr[D_3 = \sigma]| > \varepsilon$. Assume that $a = \Pr[D_2 = \sigma]$, $b = \Pr[D_3 = \sigma]$. Then, the probability weight of (π_{i-1}, σ) in $\Pi_{\text{sim}, [i]}$ differs from that in $\Pi_{\text{real}, [i]}$ by at least

$$|p \cdot a - q \cdot b| > (\theta - \theta\varepsilon/2) \cdot 1 - \theta \cdot (1 - \varepsilon) \geq \theta\varepsilon/2$$

(since $p \geq \theta$, $q \geq \theta - \theta\varepsilon/2$, $|a - b| > \varepsilon$), which contradicts to the fact that the max distance between Π_{sim} and Π_{real} is at most $\theta\varepsilon/2$. ■

Combining the above three claims, we have that with probability at least $1 - 2 \cdot 2^{-n}$,

$$\Delta_{\max}(D_1, D_4) \leq 3\varepsilon \leq \frac{1}{8 \cdot 2^{2\ell(n)} \alpha(n)}$$

If this holds, together with the support size observation, we apply Lemma 5.8 to conclude that

$$\Pr[v \leftarrow \{0, 1\}^n : \text{Samp}(D'_S, v) = \text{Samp}(D'_P, v)] \geq 1 - \frac{2^{2\ell(n)}}{8 \cdot 2^{2\ell(n)} \alpha(n)} - 2^{\ell(n)} \cdot 2^{-n} \geq 1 - \frac{1}{4\alpha(n)}$$

(since α is polynomial in n). Taking a Union bound (to deal with the fact that the above three claims hold with probability $1 - 2 \cdot 2^{-n}$), we conclude that conditioned on E and r being good, given $(CT, y) \leftarrow \text{Enc}'_r(x)$, $z \leftarrow \text{Dec}(CT, w)$, $y = z$ holds with probability at least

$$1 - \frac{1}{4\alpha(n)} - 2 \cdot 2^{-n} \geq 1 - \frac{1}{2\alpha(n)}$$

which completes our proof. $\blacksquare$

5.2.3 Concluding the Proof of Lemma 5.3

Proof: [of Lemma 5.3] The correctness of (Enc, Dec) directly follows from Lemma 5.9 and Lemma 5.10.²⁴ To argue (soundness) security, we rely on Lemma 5.6, and let R denotes the PPT algorithm guaranteed to exist in Lemma 5.6. Suppose for contradiction that there exists an attacker A such that for infinitely many $n \in \mathbb{N}$, $x \notin L$, $|x| = n$,

$$\Pr[(CT, y) \leftarrow \text{Enc}(x) : A(CT) = y] \geq 1 - \frac{1}{2\beta(n)}$$

Consider any n , $x \notin L$, $|x| = n$ such that the above holds, we rely on a case analysis on whether the probability that $r \leftarrow \{0, 1\}^{q(n)}$, $(x, r, \pi) \leftarrow S(x, r)$, (r, π) will be accepted by V with probability $\geq s(n) + \frac{1}{4\delta(n)}$ (where “being accepted” is defined as in Figure 2, and notice that this probability is denoted by $d(n)$ in Lemma 5.6):

- If not, by a standard Chernoff type argument, with probability $1 - 2^{-n}$, S will not pass the test in step 1 of $\text{Enc}(x)$ (see Figure 2) and Enc will output $(\perp, y)$ for a random $y \in \{0, 1\}^n$. By a Union bound, given $(CT, y) \leftarrow \text{Enc}(x)$, y is informational-theoretically hard to compute with probability $1 - 2 \cdot 2^{-n}$. Thus, there can only be finitely many n on which this happens.
- If so, notice that

$$\begin{aligned} & \Pr[(CT, y) \leftarrow \text{Enc}(x) : A(CT) = y] \\ & \leq \Pr[(CT, y) \leftarrow \text{Enc}(x) : A(CT) = y \mid CT = \perp] \\ & \quad + \Pr[(CT, y) \leftarrow \text{Enc}(x) : A(CT) = y \mid CT \neq \perp] \\ & = 2^{-n} + \Pr[(CT, y) \leftarrow \text{Enc}(x) : A(CT) = y \mid CT \neq \perp] \end{aligned}$$

It follows that

$$\Pr[(CT, y) \leftarrow \text{Enc}(x) : A(CT) = y \mid CT \neq \perp] \geq 1 - \frac{1}{2\beta(n)} - 2^{-n} \geq 1 - \frac{1}{\beta(n)}$$

²⁴Note that Lemma 5.10 only shows that the correctness requirement of (Enc, Dec) holds on sufficiently large n . We remark that such a scheme can be transformed into a scheme such that the correctness condition holds on all n (at a price of losing the security guarantee when n is small) by letting Enc always output the same string when n is small.

Thus, by Lemma 5.6, we have that

$$\Pr[\langle R^A, V \rangle(x) = 1] \geq s(n) + \frac{1}{4\delta(n)} - \frac{1}{8\delta(n)} \geq s(n) + \frac{1}{8\delta(n)}$$

which breaks the soundness condition on x .

Finally, since the first case happens only on finitely many n , we conclude that we break the soundness condition on infinitely many n , which contradicts the statistical soundness of (P, V) (if A is a time-unbounded attacker) or the computational soundness of (P, V) (if A can be computed in non-uniform polynomial time). ■

6 Laconic Zero Knowledge from Witness Encryption

In this section, we show how to construct a laconic SHVZK protocol for L given any WE for L . A closely related result, claiming only the HVZK property, was already obtained by [KM^N+14]; we simply remark that their protocol in fact already achieved the SHVZK property (as do most HVZK protocols). For the convenience of the reader, we include the proof.

Theorem 6.1. *Let L be an NP language. Assume that there exists a witness encryption for L with computational security (resp. statistical security).*

Then, for every (efficiently computable) $\ell(n) \geq 1$, there exists an (deterministic)-efficient prover ℓ -laconic two-round statistical SHVZK argument (resp. proof) for L with completeness $1 - 2^{-n}$ and soundness error $2^{-\ell(n)} + \text{negl}(n)$.

Proof: We start with the case that $\ell = 1$. Let (Enc, Dec) be the witness encryption scheme for L . Consider the following protocol (P, V) : V on input x tosses a random bit b , and sends $CT = \text{Enc}(x, b)$ to the prover. The prover, on input x , receives a witness $w \in R_L(x)$, runs $\text{Dec}(CT, w)$ to decrypt the message, and sends back the decrypted bit. Finally, V accepts if the bit prover sent to it equals b .

The completeness follows from the correctness of (Enc, Dec) , and the soundness follows from the security of (Enc, Dec) [GM84]. For zero-knowledge, notice that the simulator can simply output the verifier's random bit b as the prover message. Since the prover is able to recover b with probability $1 - 2^{-n}$, the simulator will output the correct prover message with probability $1 - 2^{-n}$.

For larger ℓ , consider the following protocol (P, V) : V on input x tosses a random bits $b_1, \dots, b_\ell$, and for each $i \in [\ell]$, sends 3 independent copies of $\text{Enc}(x, b_i)$, CT_i^1, CT_i^2, CT_i^3 to the prover. The prover, on input x , receives a witness $w \in R_L(x)$, runs $\text{Dec}(CT_i^j, w)$ to decrypt the message, and for every i sends back the majority of decrypted bits. Finally, V accepts if the bit prover sent to it equals b_i for each i .

The soundness follows by a standard hybrid argument (in which we can replace b_i with bits that are uniformly random given the verifier messages, making them information-theoretically hard to predict). The completeness follows since for every $i \in [\ell]$, the probability of an error in the recovery of b_i is at most $3 \cdot (2^{-n})^2$, and thus by the union bound the completeness error is at most $3\ell \cdot 2^{-2n} \leq 2^{-n}$. ■

Acknowledgment

We are grateful to Yuval Ishai, Nir Bitansky and Daniele Venturi for pointing out the highly relevant related work on predictable arguments and their applications.

References

- [ABP15] Michel Abdalla, Fabrice Benhamouda, and David Pointcheval. Disjunctions for hash proof systems: New constructions and applications. In *Advances in Cryptology-EUROCRYPT 2015: 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part II 34*, pages 69–100. Springer, 2015.
- [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.
- [BC20] Nir Bitansky and Arka Rai Choudhuri. Characterizing deterministic-prover zero knowledge. In *Theory of Cryptography: 18th International Conference, TCC 2020, Durham, NC, USA, November 16–19, 2020, Proceedings, Part I 18*, pages 535–566. Springer, 2020.
- [BCC88] Gilles Brassard, David Chaum, and Claude Crépeau. Minimum disclosure proofs of knowledge. *J. Comput. Syst. Sci.*, 37(2):156–189, 1988.
- [BDRV18] Itay Berman, Akshay Degwekar, Ron D Rothblum, and Prashant Nalini Vasudevan. From laconic zero-knowledge to public-key cryptography. In *Annual International Cryptology Conference*, pages 674–697. Springer, 2018.
- [BGI⁺12] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. *J. ACM*, 59(2):6:1–6:48, 2012.
- [BGJ⁺13] Elette Boyle, Sanjam Garg, Abhishek Jain, Yael Tauman Kalai, and Amit Sahai. Secure computation against adaptive auxiliary information. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 316–334. Springer, 2013.
- [BIN97] Mihir Bellare, Russell Impagliazzo, and Moni Naor. Does parallel repetition lower the error in computationally sound protocols? In *FOCS '97*, pages 374–383, 1997.
- [BIOW20] Ohad Barta, Yuval Ishai, Rafail Ostrovsky, and David J Wu. On succinct arguments and witness encryption from groups. In *Annual International Cryptology Conference*, pages 776–806. Springer, 2020.
- [BISW18] Dan Boneh, Yuval Ishai, Amit Sahai, and David J Wu. Quasi-optimal snargs via linear multi-prover interactive proofs. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 222–255. Springer, 2018.
- [BM88] László Babai and Shlomo Moran. Arthur-merlin games: A randomized proof system, and a hierarchy of complexity classes. *J. Comput. Syst. Sci.*, 36(2):254–276, 1988.
- [CDS94] Ronald Cramer, Ivan Damgård, and Berry Schoenmakers. Proofs of partial knowledge and simplified design of witness hiding protocols. In *CRYPTO*, pages 174–187, 1994.
- [CHS05] Ran Canetti, Shai Halevi, and Michael Steiner. Hardness amplification of weakly verifiable puzzles. In *Theory of Cryptography, Second Theory of Cryptography Conference, TCC 2005, Cambridge, MA, USA, February 10-12, 2005, Proceedings*, pages 17–33, 2005.

- [CS02] Ronald Cramer and Victor Shoup. Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In *International conference on the theory and applications of cryptographic techniques*, pages 45–64. Springer, 2002.
- [Dam02] Ivan Damgård. On σ -protocols. *Lecture Notes, University of Aarhus, Department for Computer Science*, 84, 2002.
- [FLS90] Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple non-interactive zero knowledge proofs based on a single random string. In *FOCS '90*, pages 308–317, 1990.
- [FNV17] Antonio Faonio, Jesper Buus Nielsen, and Daniele Venturi. Predictable arguments of knowledge. In *Public-Key Cryptography–PKC 2017: 20th IACR International Conference on Practice and Theory in Public-Key Cryptography, Amsterdam, The Netherlands, March 28–31, 2017, Proceedings, Part I 20*, pages 121–150. Springer, 2017.
- [FS90] Uriel Feige and Adi Shamir. Witness indistinguishable and witness hiding protocols. In *STOC '90*, pages 416–426, 1990.
- [FWW23] Cody Freitag, Brent Waters, and David J Wu. How to use (plain) witness encryption: Registered abe, flexible broadcast, and more. In *Annual International Cryptology Conference*, pages 498–531. Springer, 2023.
- [GGH⁺13] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *Symposium on Foundations of Computer Science (FOCS)*, pages 40–49, 2013.
- [GGSW13] Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. Witness encryption and its applications. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 467–476, 2013.
- [GH98] Oded Goldreich and Johan Håstad. On the complexity of interactive proofs with bounded communication. *Inf. Process. Lett.*, 67(4):205–214, 1998.
- [GL89] Oded Goldreich and Leonid A. Levin. A hard-core predicate for all one-way functions. In *STOC*, pages 25–32, 1989.
- [GM84] Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *J. Comput. Syst. Sci.*, 28(2):270–299, 1984.
- [GMM17] Sanjam Garg, Mohammad Mahmoody, and Ameer Mohammed. Lower bounds on obfuscation from all-or-nothing encryption primitives. In *Annual International Cryptology Conference*, pages 661–695. Springer, 2017.
- [GMR88] Shafi Goldwasser, Silvio Micali, and Ronald L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM J. Comput.*, 17(2):281–308, 1988.
- [GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989.
- [GO94] Oded Goldreich and Yair Oren. Definitions and properties of zero-knowledge proof systems. *Journal of Cryptology*, 7:1–32, 1994.

- [GVW02] Oded Goldreich, Salil Vadhan, and Avi Wigderson. On interactive proofs with a laconic prover. *Computational Complexity*, 11(1):1–53, 2002.
- [Hai13] Iftach Haitner. A parallel repetition theorem for any interactive argument. *SIAM J. Comput.*, 42(6):2487–2501, 2013.
- [HN23] Shuichi Hirahara and Mikito Nanashima. Learning in pessiland via inductive inference. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 447–457. IEEE, 2023.
- [HPWP10] Johan Håstad, Rafael Pass, Douglas Wikström, and Krzysztof Pietrzak. An efficient parallel repetition theorem. In *Theory of Cryptography, 7th Theory of Cryptography Conference, TCC 2010, Zurich, Switzerland, February 9-11, 2010. Proceedings*, pages 1–18, 2010.
- [JLS21] Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from well-founded assumptions. In *Symposium on Theory of Computing (STOC)*, pages 60–73, 2021.
- [JLS22] Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from LPN over $\mathbb{F}_p$, DLIN, and PRGs in NC^0 . In *International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, pages 670–699, 2022.
- [KC00] Valentine Kabanets and Jin-yi Cai. Circuit minimization problem. In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, May 21-23, 2000, Portland, OR, USA*, pages 73–79, 2000.
- [Kiy24] Susumu Kiyoshima. Resettable statistical zero-knowledge for np. In *Annual International Cryptology Conference*, pages 288–320. Springer, 2024.
- [KMN⁺14] Ilan Komargodski, Tal Moran, Moni Naor, Rafael Pass, Alon Rosen, and Eylon Yogev. One-way functions and (im)perfect obfuscation. *IACR Cryptology ePrint Archive*, 2014:347, 2014.
- [Ko86] Ker-I Ko. On the notion of infinite pseudorandom sequences. *Theor. Comput. Sci.*, 48(3):9–33, 1986.
- [Kol68] A. N. Kolmogorov. Three approaches to the quantitative definition of information. *International Journal of Computer Mathematics*, 2(1-4):157–168, 1968.
- [LMP24] Yanyi Liu, Noam Mazon, and Rafael Pass. A note on zero-knowledge for np and one-way functions. *Cryptology ePrint Archive*, 2024.
- [MP23] Noam Mazon and Rafael Pass. Counting unpredictable bits: A simple prg from one-way functions. *Cryptology ePrint Archive*, 2023.
- [PV07] Rafael Pass and Muthuramakrishnan Venkatasubramanian. An efficient parallel repetition theorem for arthur-merlin games. In *STOC '07*, pages 420–429, 2007.
- [Tsa22] Rotem Tsabary. Candidate witness encryption from lattice techniques. In *Annual International Cryptology Conference*, pages 535–559. Springer, 2022.

- [Vad00] Salil P. Vadhan. On transformation of interactive proofs that preserve the prover’s complexity. In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, May 21-23, 2000, Portland, OR, USA*, pages 200–207, 2000.
- [VV86] Leslie G. Valiant and Vijay V. Vazirani. NP is as easy as detecting unique solutions. *Theoretical Computer Science*, 47(3):85–93, 1986.
- [VWW22] Vinod Vaikuntanathan, Hoeteck Wee, and Daniel Wichs. Witness encryption and null-io from evasive lwe. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 195–221. Springer, 2022.
- [Wee06] Hoeteck Wee. Finding pessiland. In *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*, pages 429–442, 2006.
- [Yao82] Andrew Chi-Chih Yao. Theory and applications of trapdoor functions (extended abstract). In *23rd Annual Symposium on Foundations of Computer Science, Chicago, Illinois, USA, 3-5 November 1982*, pages 80–91, 1982.