Certified-Everlasting Quantum NIZK Proofs

Nikhil Pappu
nikpappu@pdx.edu

Portland State University

Abstract

We study non-interactive zero-knowledge proofs (NIZKs) for NP satisfying: 1) statistical soundness, 2) computational zero-knowledge (ZK) and 3) certified-everlasting zero-knowledge (CE-ZK). The CE-ZK property allows a verifier of a quantum proof to revoke the proof in a way that can be checked (certified) by the prover. Conditioned on successful certification, the verifier's state can be efficiently simulated with only the statement, in a *statistically indistinguishable* way. Our contributions regarding these certified-everlasting NIZKs (CE-NIZKs) are as follows:

- We identify a barrier to obtaining CE-NIZKs in the CRS model via generalizations of known interactive ZK proofs that satisfy CE-ZK.
- We circumvent this by constructing CE-NIZK using non-black-box use of NIZK for NP satisfying certain properties, along with OWFs. As a result, we obtain CE-NIZKs for NP in the CRS model, based on polynomial hardness of the learning with errors (LWE) assumption.
- In addition, we observe that the aforementioned barrier does not apply to the shared EPR model. We leverage this to construct CE-NIZK for NP in this model based on any statistical binding hidden-bits generator, which is known from LWE. The only quantum computation here involves single-qubit measurements of the EPR pairs.

Contents

1 Introduction **3**
1.1 Contributions 4
1.2 Related Work 5

2 Technical Overview **6**
2.1 Defining Certified-Everlasting NIZK 6
2.2 Non-Existence of Deletion-Resistant Protocols 7
2.3 On Generalizations of Prior Work 8
2.4 Feasibility in the CRS Model 9
2.5 Quantum-Efficient Protocol in the Shared EPR Model 12
2.6 Relation with Other Quantum-Enabled NIZKs 14
2.7 Open Questions 15

3 Preliminaries **16**
3.1 Notation 16
3.2 Quantum Information 17
3.3 Cryptography Basics 17
3.4 Superposition-Secure Non-Interactive Zero-Knowledge 18
3.5 Certified Deletion of BB84 States 19

4 Certified-Everlasting NIZK Proofs **20**

5 CE-NIZK in the CRS Model **22**
5.1 Impossibility of Deletion-Resistant CE-NIZKs 22
5.2 Construction from LWE 23

6 CE-NIZK in the Shared EPR Model **31**
6.1 The Hidden-Bits Paradigm 31
6.2 Construction from any Hidden Bits Generator 33

1 Introduction

A non-interactive zero-knowledge proof (NIZK) [BFM88] allows a *prover* to demonstrate the validity of a statement to a *verifier* with a *single* message, without revealing any other information. As an example, if Alice knows a solution to a Sudoku puzzle, she can convince Bob that the puzzle is indeed solvable, without leaking the solution itself. Crucially, if the puzzle does not actually have a solution, then a cheating Alice should not be able to fool Bob, a guarantee formally known as *soundness*. On the other hand, the inability of a cheating Bob to learn any additional information is referred to as *zero-knowledge*. This latter property is formalized by the existence of an efficient algorithm called a *simulator*, that can recreate Bob's view using the statement alone. Intuitively, it captures that whatever Bob learns from a proof, he can compute efficiently without it.

NIZKs have been a cornerstone of cryptography with widespread theoretical and practical impact. For example, they play an important role in constructions of CCA-secure PKE [NY90], digital signatures [BG90], ZAPs [DN00], round-efficient MPC [AJL+12], and functional encryption [GGH+13] to name a few. On the practical front, they have been instrumental in designing anonymous credentials [BCC+09], group signatures [BCC+16], cryptocurrencies [SCG+14] and verifiable computation [GGP10].

Setup Models. It is well-known that NIZKs suffer a drawback compared to their multi-round (interactive) counterparts in that they are impossible to obtain in the plain model [GO94]. Fortunately, this can be circumvented in a realistic setup model called the common reference string (CRS) model. Here, a trusted party samples a random (possibly non-uniform) public string before the protocol, and security is required to hold with overwhelming probability over the choice of CRS. Other commonly studied setup models include the designated verifier [QRW19] and random oracle models [BR93]. In the quantum world, recent works [MY22a, ABKK23] have considered two-party protocols where the parties share different registers of an entangled quantum state. These are sometimes known to provide stronger guarantees than the CRS model [BKS23]. Recently, more complex yet localized quantum setups have also garnered attention in two-party protocols [AGL24, MNY24].

Unbounded Adversaries. Due to their widespread use cases, NIZKs satisfying various security properties have been explored. Of these, perhaps the most important characterization is in regards to whether security holds against an unbounded adversary, or only a computationally-bounded one. For instance, *soundness* typically only requires that an *efficient* malicious prover cannot convince the verifier of a *false* statement. Likewise, the standard notion of *zero-knowledge* requires that an *efficient* verifier learns nothing from the transcript. However, one can consider upgrades of these in the form of *statistical soundness* and *statistical zero-knowledge*, which provide security against an unbounded prover and an unbounded distinguisher (equivalently, the verifier in the NIZK case) respectively. Hence, many NIZKs in the literature satisfy these stronger properties [BFM88, FLS90, FLS99, GO94, GOS06, SW14, BP15, CCH+19, PS19, WWW25].

Statistical Soundness or Zero-Knowledge?. Although statistical variants of soundness and zero-knowledge are achievable from standard assumptions, it is unlikely to obtain both simultaneously, except for a small subclass of NP languages [PS05]. This is the case even for interactive-protocols [For87], with evidence that quantum protocols do not help either [MW18]. Consequently, one has to choose which of these two properties to give up in favor of the other, depending on the application. Typically, statistical zero-knowledge is preferred as soundness only needs to hold during the online phase, while zero-knowledge can be compromised anytime in the future to leak

sensitive information. Still, this is not always ideal. Consider a large corporation (prover) that might have various means to cut down the time to break certain computational assumptions, with the verifier being an average user. In this case, a computational bounded assumption on the verifier could be more reasonable. A statistically sound NIZK can also remove the need to refresh the CRS periodically in such a setting. However, this introduces other issues, such as the verifier being able to sell the data to a powerful rival corporation.

A Quantum Compromise. Such a tradeoff seems inevitable in the classical setting, but previous works have proposed a novel compromise using quantum resources, in the context of interactive zero-knowledge [HMNY22, BK23]. In particular, the prover and verifier engage in a statistically-sound quantum protocol, after which the verifier is asked to return a *certificate of deletion*. Then, an everlasting guarantee akin to statistical zero-knowledge holds, given that the certificate is verified to be valid. Note that this only makes sense in a quantum world due to the no-cloning theorem, as the verifier can retain a classical proof indefinitely.

Observe also that the requirement to delete the proof does not interfere with verification. This is because verification can be performed first without disturbing the state (by gentle measurement), and the proof can be deleted having served its purpose. Another interesting aspect is that the verifier can immediately verify and revoke the proof, which can then be checked. As a result, the parties need not maintain the state for long unlike other primitives with certified-deletion. This can be quite useful due to the difficulty of maintaining quantum states.

The aforementioned works then showed that such proofs can be obtained for all languages in QMA (the quantum generalization of NP) assuming only OWFs, or even weaker quantum-assumptions such as the existence of pseudo-random states [JLS18]. Naturally, one can consider the non-interactive analogue where the prover sends a single message to the verifier, and the verifier sends back a deletion certificate. While one might argue this is no better than a two-round protocol, the deployed applications wouldn't be slowed down by the second message, which could be batched and sent at a later point. Hence, at the expense of using quantum resources and relying on the ability to penalize a non-confirming verifier, there exists the possibility of obtaining statistically-sound NIZKs with an everlasting zero-knowledge guarantee. Despite the hope, and the construction of many advanced primitives with certified-everlasting forms of security (See Section 1.2), CE-NIZKs are yet to be realized. We explore this possibility in detail.

1.1 Contributions

1. *Definitions and Impossibility.* We present definitions of certified-everlasting NIZK proofs (CE-NIZKs) in the CRS and shared-EPR models to begin with. Then, we examine natural approaches to constructing CE-NIZKs in the CRS model. Specifically, we identify a class of protocols we call *deletion-resistant* CE-NIZKs. At a high level, these admit a deletion algorithm that splits the proof state into two registers, one that passes verification of deletion, and another that passes proof verification. We then rule out the existence of such protocols by showing that they imply NIZKs that are both statistically-sound and statistically zero-knowledge. We also argue that non-interactive counterparts of known interactive protocols fall into this category.

2. *Feasibility in the CRS Model.* As a result of this barrier, we provide a new approach to constructing CE-NIZKs in the CRS model, where we employ two "layers" of NIZK proofs that are generated in quantum-superposition. In order to prove security, we employ the two-slot technique [NY90] and the OR proof strategy [FLS90, FLS99], and reduce to a certified-deletion

theorem of BB84 states [BK23]. This allows us to construct a CE-NIZK for NP based on (non-black-box use of) a superposition-secure NIZK proof for NP[1] (Definition 3.6), along with just post-quantum OWFs. We observe that superposition-secure NIZKs (a stronger requirement than post-quantum NIZKs) are implied by LWE due to the construction of [WWW25]. Hence, we obtain our result based on LWE, matching the state-of-the-art assumption for post-quantum NIZK.

3. *Bypassing the Impossibility with Shared Entanglement.* We then explore a natural quantum setup where the prover and verifier share entanglement in the form of halves of EPR pairs. Surprisingly, we find that the impossibility regarding *deletion-resistant* protocols no longer applies to this setting. As a consequence, we show that a generalization of the hidden-bits compiler of [QRW19] suffices to obtain CE-NIZKs for NP in this setting. In particular, we obtain CE-NIZKs based on any statistically binding hidden-bits generator [QRW19], which can be based on LWE [WWW25]. Although the assumption is the same as our CRS protocol, the protocol is highly efficient in regards to the quantum operations involved. For comparison, our CRS protocol generates and verifies proofs in superposition of a highly entangled state. On the other hand, this protocol only requires single-qubit measurements of the EPR halves in one of two basis, and uses classical computation and communication otherwise.

1.2 Related Work

Certified-Everlasting Cryptography. The pioneering work of Broadbent and Islam [BI20] constructed encryption with certified-deletion, where ciphertexts can be deleted in a way that the plaintext becomes unrecoverable (by an unbounded adversary) even if the secret-key is leaked later. This was followed by a series of works on certified-everlasting security, which refers to a bounded adversary that cannot break some guarantee in the future using unbounded computation, assuming it provides a valid certificate during the computationally bounded stage.

The work of Hiroka, Morimae, Nishimaki and Yamakawa [HMNY22] constructed a certified-everlasting statistical-binding commitment scheme and used it to obtain a certified-everlasting zero-knowledge proof system for QMA. Later, the work of Bartusek and Khurana [BK23] introduced an elegant compiler for obtaining primitives with certified-deletion. Through it, they obtained PKE, ABE, FHE, commitments (and ZK via [HMNY22]), witness encryption, timed-release encryption, MPC protocols and secret-sharing schemes with certified-everlasting security. The work of Hiroka et al. [HKM$^+$24] constructed public-key functional encryption (and other variants), compute-and-compare obfuscation and garbled circuits with such security guarantees. Advanced secret-sharing schemes with certified-everlasting security were studied in the work of Bartusek and Raizes [BR24]. The work of Champion, Kitagawa, Nishimaki and Yamakawa [CKNY25] constructed a certified-everlasting variant of untelegraphable encryption, a relaxation of unclonable encryption. A notion of differing inputs obfuscation (dIO) with certified-everlasting security was introduced and constructed in the work of Bartusek et al. [BGK$^+$24]. Their construction provides the publicly-verifiable deletion property [KNY23, BKP23, BKM$^+$23, KS25].

Quantum-Enabled NIZKs. In recent years, several works have studied NIZKs with interesting properties that are only feasible in a quantum world. Some of these works construct schemes that prevent copying attacks, i.e., they ensure that the recipient of a proof cannot convince several

[1]We require two special notions from such a NIZK: superposition-secure adaptive zero-knowledge and post-quantum non-adaptive zero-knowledge with statistical CRS indistinguishability which are satisfied by the LWE-based NIZKs of [Wat24, WWW25].

other verifiers, without itself knowing a witness. These "unclonable" NIZKs and variants were studied in the concurrent works of Jawale and Khurana [JK24] and Goyal, Malavolta and Raizes [GMR24]. A related notion of NIZK with certified-deletion (NIZK-CD) was explored by the work of Abbaszadeh and Katz [AK25] where the recipient must generate a valid proof of deletion, and cannot convince another verifier thereafter. While the unclonable variants were shown to imply public-key quantum money, NIZK-CD was constructed in [AK25] from standard assumptions. In Section 2, we informally argue that CE-NIZK implies a form of NIZK-CD (without the extractability notion considered in [AK25]). We also argue NIZK-CD does not imply CE-NIZK.

The work of Çakan, Goyal and Raizes [CGR26] studied NIZK with a stronger form of certified-deletion called certified-deniability. Intuitively, this is a simulation-based notion that guarantees that the state after deletion of the proof could be computed without the proof. The work presents a feasibility for this in the random oracle model (ROM), with evidence against its existence in the plain model. Our simulation definition shares some similarity to theirs, while additionally requiring statistical indistinguishability. Crucially, their impossibility does not contradict our CRS feasibility, as our simulator has the ability to manipulate the CRS. In the certified-deniability context, this ability would be restricted, similar to the work's restriction on programmability in the ROM.

Finally, the work of [GLR+25] studies a notion of one-time security for NIZK and presents a construction from iO and LWE. This notions allows an authority to distribute quantum proving tokens to a prover. Then, the prover can only convince a verifier of as many statements as the number of tokens received. This notion is shown to imply public-key quantum money which justifies the use of the strong assumption of iO. We believe the guarantees of all these NIZK notions (apart from NIZK-CD) to be orthogonal to that of CE-ZK. For a detailed comparison of our work with some of these works, see Section 2.6.

Revocable Cryptography. Revocable cryptography is closely related to the notions discussed above in that it involves a deletion (revocation) phase, after which certain information/capability is lost. This was pioneered by the work of Unruh [Unr14] which introduced revocable time-released encryption. Revocable notions for encryption/signatures were also studied in [AMP25, MPY24], and are similar to the notions of certified-deletion in the sense that secret data is being deleted. On the other hand, a long line of works on secure software leasing [AL21, KNY21, BJL+21, CMP24] and secure key leasing [KN22, APV23, AKN+23, PWYZ24, AHH24, CGJL25, KMY25, KNP25, KNP26, KLYY26] consider the following high-level premise. A user is provided a quantum state that can be used to evaluate some software or cryptographic functionality. Later, the user can provide a deletion-certificate, after which it is guaranteed that they can no longer evaluate the functionality. We remark that almost all these schemes only consider a computationally-bounded second-stage (post-deletion) adversary. This restriction is necessary in several contexts like the leasing of decryption keys of a PKE scheme.

2 Technical Overview

2.1 Defining Certified-Everlasting NIZK

We will begin by defining the syntax of a certified-everlasting non-interactive zero-knowledge proof (CE-NIZK) (Definition 4.1). A CE-NIZK for a language $\mathcal{L} \in \mathsf{NP}$ is a tuple of four algorithms $(\mathsf{Setup}, \mathsf{P}, \mathsf{V}, \mathsf{Cert})$ described as follows. The classical algorithm $\mathsf{Setup}(1^\lambda)$ outputs a common reference string crs as usual. The quantum prover's algorithm P on input the CRS crs, a statement $x \in \mathcal{L}$ and corresponding witness $\omega \in R_{\mathcal{L}}(x)$, outputs a quantum proof state σ, along with a (possibly classical) quantum state ρ_{P}, which will later be used to validate returned proofs. The

quantum verifier's algorithm V on input crs, x and proof σ outputs a decision bit b along with a quantum state ρ_{V}. The quantum certification algorithm Cert takes as input quantum states $\rho_{\mathsf{P}}, \rho_{\mathsf{V}}$ and outputs $\top$ (accept) or $\bot$ (reject). The context is that the verifier would check the validity of the proof σ to obtain ρ_{V} (these states would be close to one another in trace distance due to completeness and gentle-measurement) and send it to the prover who checks that the proof is revoked by running $\mathsf{Cert}(\rho_{\mathsf{P}}, \rho_{\mathsf{V}})$. In general, Cert may end up destroying ρ_{P} and ρ_{V}. Note that when we refer to a CE-NIZK for NP, we mean one that is a CE-NIZK for an NP-complete language.

We require a CE-NIZK for NP to satisfy completeness, statistical soundness and computational zero-knowledge, which must also hold for a standard NIZK proof for NP. In addition, it must also satisfy certified-everlasting zero-knowledge (CE-ZK). Intuitively, the notion guarantees that the view of any QPT malicious verifier V^* can be simulated in a statistically close manner with just the statement x, conditioned on certification being successful. Formally, for every QPT verifier V^*, there must exist a QPT simulator $\mathsf{Sim}_{\mathsf{V}^*}$ such that the following guarantee holds for every true statement x and corresponding witness ω:

$$\left[\begin{array}{ll} & \mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ & (\sigma, \rho_{\mathsf{P}}) \leftarrow \mathsf{P}(\mathsf{crs}, x, \omega) \\ \widetilde{\rho} \approx_s \rho' \ : & \rho_{\mathrm{A}\otimes\mathrm{B}} \leftarrow \mathsf{V}^*(\mathsf{crs}, x, \sigma) \\ & \text{If } \mathsf{Cert}(\rho_{\mathsf{P}}, \rho_{\mathrm{A}}) = \top \text{ then } \widetilde{\rho} := \rho_{\mathrm{B}}. \\ & \text{Else } \widetilde{\rho} := \bot. \\ & \rho' \leftarrow \mathsf{Sim}_{\mathsf{V}^*}(x) \end{array}\right]$$

Let us explain this formalism. First, a crs is sampled using Setup and a real proof σ and prover state ρ_{P} are computed. Next, the proof σ is provided to the malicious verifier, who outputs a bipartite state $\rho_{\mathrm{A}\otimes\mathrm{B}}$. Then, the notion sets a state $\widetilde{\rho}$ to be ρ_{B} (the residual state on register B after Cert is applied to register A) only if the state ρ_{A} on register A passes certification by $\mathsf{Cert}(\rho_{\mathsf{P}}, \rho_{\mathrm{A}})$. If not, $\widetilde{\rho}$ is simply set to $\bot$, thereby giving a distinguisher no additional information when the certification is unsuccessful. Then, we require that the simulator $\mathsf{Sim}_{\mathsf{V}^*}$ on input x produces a simulated state ρ' that is close in trace distance to $\widetilde{\rho}$ from the real execution. Note that our protocols will come with a simulator that makes black-box use of V^*, but we define a non-black-box one here for generality.

2.2 Non-Existence of Deletion-Resistant Protocols

Before we introduce our approach to realizing CE-NIZKs, we will discuss some challenges in achieving them. Specifically, we identify a natural class of CE-NIZKs that are impossible to construct. The intuition is that these come with a deletion algorithm that splits the proof state into two, the first which passes prover certification, and a second which passes NP verification. Consequently, this operation can be performed at the prover's end before sending the second part to the verifier. We observe that these imply NIZKs which are both statistically sound and statistically zero-knowledge which are unlikely to exist. We define this subclass of CE-NIZKs formally as follows.

Definition 2.1 (Deletion-Resistant CE-NIZKs)**.** A non-interactive protocol with the syntax of CE-NIZK (Definition 4.1) is said to be a deletion-resistant CE-NIZK if it satisfies the CE-ZK property, along with the following properties. There exist QPT algorithms $\mathsf{V}^*, \widetilde{\mathsf{V}}$ such that the following conditions hold for every $x \in \mathcal{L}$, $\omega \in R_{\mathcal{L}}(x)$ and $\widetilde{x} \notin \mathcal{L}$, where $(\mathsf{crs}, \sigma, \rho_{\mathsf{P}}, \rho_{\mathrm{A}\otimes\mathrm{B}})$ are sampled as

$\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda), (\sigma, \rho_{\mathsf{P}}) \leftarrow \mathsf{P}(\mathsf{crs}, x, \omega)$ and $\rho_{\mathrm{A}\otimes\mathrm{B}} \leftarrow \mathsf{V}^*(\sigma)$:

$$\Pr\left[\mathsf{Cert}(\rho_{\mathsf{P}}, \rho_{\mathrm{A}}) \to \top\right] \geq 1/\mathsf{poly}(\lambda) \tag{1}$$

$$\Pr\left[\widetilde{\mathsf{V}}(\mathsf{crs}, x, \rho_{\mathrm{B}}) \to 1 \;\middle|\; \mathsf{Cert}(\rho_{\mathsf{P}}, \rho_{\mathrm{A}}) \to \top\right] \geq 1 - \mathsf{negl}(\lambda) \tag{2}$$

$$\Pr\left[\exists \widetilde{\sigma} : \widetilde{\mathsf{V}}(\mathsf{crs}, \widetilde{x}, \widetilde{\sigma}) \to 1\right] \leq \mathsf{negl}(\lambda) \tag{3}$$

Theorem 2.2. *If there exists a deletion-resistant CE-NIZK for a language $\mathcal{L}$, there is a quantum NIZK for $\mathcal{L}$ with statistical soundness and statistical ZK.*

The proof is simple, so we show it in Section 5. Also, even if (2) holds with only inverse polynomial probability, it results in a NIZK proof but with inverse polynomial notions of completeness and statistical ZK. However, we believe an extension of the argument of Pass and Shelat [PS05] works even in this case (See the third paragraph in Section 2.6), which would show that the language $\mathcal{L}$ belongs to $\mathsf{QAM} \cap \mathsf{coQAM}$ (and hence unlikely to be NP-complete).

2.3 On Generalizations of Prior Work

We now recall previous works on CE-ZK proofs in the interactive setting by Hiroka et al. [HMNY22] and Bartusek and Khurana [BK23]. The work of Hiroka et al. constructed CE-ZK proofs for languages in QMA based on a new primitive they introduced called commitments with statistical binding and certified-everlasting hiding. The certified-everlasting hiding property refers to the fact that after deletion, even an unbounded receiver is unable to learn anything about the contents of the commitment. While such commitments were constructed in the random oracle model in their work, Bartusek and Khurana constructed them from any statistical binding commitment scheme in the plain model. Note that such statistically binding commitments can be based on OWFs [Nao91] and from potentially weaker quantum-cryptographic assumptions [MY22b, AQY22].

We now argue that these protocols do not generalize easily for constructing CE-NIZKs. At a high level, the CE-ZK proofs established by these works rely on the following template:

- First, the prover sends some $n = \mathsf{poly}(\lambda, |x|, |w|)$ commitments $\{|\mathsf{com}^i\rangle\}_{i\in[n]}$ to the verifier, generated using the certified-everlasting commitment scheme.
- Then, the verifier sends a randomly chosen opening set $S \subset [n]$ to the prover, along with deletion certificates $\{\mathsf{cert}^i\}_{i\in[n]\setminus S}$ corresponding to the commitments that are not to be opened.
- Finally, the prover sends the openings of the commitments $\{\mathsf{open}^i\}_{i\in S}$ for the set S. These openings are then sufficient for the verifier to determine the validity of the statement.

Note that while these works use classical certificates $\{\mathsf{cert}^i\}_{i\in[n]\setminus S}$, the same ideas also apply if these certificates were quantum. The main idea here is that since the unopened commitments are deleted by the verifier, even an unbounded distinguisher cannot tell apart the real and simulated views. Naturally, one might expect that a Fiat-Shamir based approach [FS87, CCH+19, PS19] in the CRS model could be employed to turn these interactive protocols into non-interactive ones. However, we argue that such protocols would be deletion-resistant, and hence run into Theorem 2.2. To see this, consider a CE-NIZK $\Pi = (\mathsf{Setup}, \mathsf{P}, \mathsf{V}, \mathsf{Cert})$ with a similar template, where the prover P determines the opening set S by itself, by applying a hash function to some classical part of the generated proof. Since V would only need to check the commitments in S as part of the protocol, we can consider $\widetilde{\mathsf{V}} := \mathsf{V}$ and a malicious QPT verifier V^* that splits the commitments in S and the commitments in $[n] \setminus S$ into two different registers. Consequently, the state $\rho_{\mathrm{A}\otimes\mathrm{B}}$

produced by V^* is such that ρ_{A} would pass certification while ρ_{B} would be accepted by V. Notice that both verification and certification could check the validity of S, but the classical information that is hashed to produce S can simply be cloned. Furthermore, designing $\mathsf{V}, \mathsf{Cert}$ to artificially deviate from the template doesn't help. For e.g., V cannot check all the commitments, as Π should satisfy computational zero-knowledge. Moreover, even if Cert requires all the commitments to be revoked, V^* can clone the contents of the opened commitments. We note that natural attempts at generalizing the hidden-bits paradigm [FLS90, FLS99] to the CE-NIZK setting do not work, due to a similar argument.

2.4 Feasibility in the CRS Model

Despite the aforementioned barrier, we show that CE-NIZKs exist in the CRS model based on standard assumptions. To this end, our first idea is to utilize a superposition of NIZK proofs entangled with some unclonable quantum state. The hope is that this at least prevents the splitting attacks mentioned in the previous subsections. We explain the structure of this approach next.

Two-Layers of NIZK Proofs. Consider a BB84 state $|y\rangle^{\theta} = H^{\theta[1]}|y[1]\rangle \otimes \ldots \otimes H^{\theta[\lambda]}|y[\lambda]\rangle$, where $y, \theta \leftarrow \{0,1\}^{\lambda}$ and H is the Hadamard transform. In the work of Bartusek and Khurana [BK23], it was shown that such states satisfy a certified deletion property. Specifically, if an adversary $\mathcal{A}$ is given $|y\rangle^{\theta}$, and $\mathcal{A}$ produces a deletion certificate $\mathsf{cert} \in \{0,1\}^{\lambda}$ such that $\mathsf{cert}[i] = y[i]$ for all $i : \theta[i] = 1$ (i.e., the values of the Hadamard basis positions are consistent), then the computational basis parity $p = \bigoplus_{i:\theta[i]=0} y[i]$ remains statistically close to uniform in the view of $\mathcal{A}$, even if the "secret-key" θ is revealed to $\mathcal{A}$ after cert is produced. In our construction, a long BB84 state of $\ell \cdot \lambda$ qubits is used, i.e., $|y\rangle^{\theta} := |y^1\rangle^{\theta^1} \otimes \ldots \otimes |y^{\ell}\rangle^{\theta^{\ell}}$, where each $y^j, \theta^j \leftarrow \{0,1\}^{\lambda}$. In the scheme, the prover P uses a classical NIZK for NP $\Pi_{\mathsf{nizk}} := (\mathsf{Setup}_{\mathsf{nizk}}, \mathsf{P}_{\mathsf{nizk}}, \mathsf{V}_{\mathsf{nizk}})$ and first generates an "inner" NIZK proof $\pi_{\mathsf{in}} \leftarrow \mathsf{P}_{\mathsf{nizk}}(\mathsf{crs}_{\mathsf{in}}, x, \omega)$ wrt a CRS $\mathsf{crs}_{\mathsf{in}} \leftarrow \mathsf{Setup}_{\mathsf{nizk}}(1^{\lambda})$. Then, π_{in} is masked as $\mathsf{ct} := \pi_{\mathsf{in}} \oplus \mathsf{pad} \oplus \mathsf{k}$, where k is a uniform string[2] and $\mathsf{pad} := p_1 \| \ldots \| p_{\ell}$, where each p_i is the computational-basis parity bit wrt the i-th BB84 state. The intuition is that if $|y\rangle^{\theta}$ is "deleted", then π_{in} cannot be recovered. However, the state $|y\rangle^{\theta}$ and "ciphertext" ct are not useful by themselves, as ct now just appears uniform. Consequently, we employ a second "outer layer" of proofs entangled with $|y\rangle^{\theta} = \sum_{\mathsf{z}} \alpha_{\mathsf{z}} |\mathsf{z}\rangle$. Specifically, we generate $\pi^{\mathsf{z}}_{\mathsf{out}} \leftarrow \mathsf{P}_{\mathsf{nizk}}(\mathsf{crs}_{\mathsf{out}}, x^{\mathsf{z}}, \omega_{\mathsf{out}} := (\theta, \mathsf{k}))$ for each superposition term z wrt the same CRS $\mathsf{crs}_{\mathsf{out}} \leftarrow \mathsf{Setup}(1^{\lambda})$, for the following "meta-statement" x^{z}:

$$\underline{\text{Meta-Statement } x^{\mathsf{z}}}: \quad \exists(\theta, \mathsf{k}) \ : \ \mathsf{V}_{\mathsf{nizk}}\Big(\mathsf{crs}_{\mathsf{in}}, x, \mathsf{ct} \oplus f(\theta, \mathsf{z}) \oplus \mathsf{k}\Big) = 1$$

with $f(\theta, \mathsf{z}) := \bigoplus_{i:\theta^1[i]=0} \mathsf{z}^1[i] \| \ldots \| \bigoplus_{i:\theta^{\ell}[i]=0} \mathsf{z}^{\ell}[i]$ where $\theta, \mathsf{z} \in \{0,1\}^{\ell\lambda}$ and $\theta = \theta^1 \| \ldots \| \theta^{\ell}$ and $\mathsf{z} = \mathsf{z}^1 \| \ldots \| \mathsf{z}^{\ell}$. Observe now that $f(\theta, \mathsf{z}) = \mathsf{pad}$ for each z in the superposition of $|y\rangle^{\theta} = \sum_{\mathsf{z}} \alpha_{\mathsf{z}} |\mathsf{z}\rangle$. Now, using the proof state $\sigma := \big(\sum_{\mathsf{z}} \alpha_{\mathsf{z}} |\mathsf{z}\rangle |\pi^{\mathsf{z}}_{\mathsf{out}}\rangle, \mathsf{ct} = \pi_{\mathsf{in}} \oplus \mathsf{pad} \oplus \mathsf{k}\big)$, the verifier can check in superposition that the proofs $\pi^{\mathsf{z}}_{\mathsf{out}}$ are valid wrt the statements x^{z}, by executing $\mathsf{V}_{\mathsf{nizk}}(\mathsf{crs}_{\mathsf{out}}, x^{\mathsf{z}}, \pi^{\mathsf{z}}_{\mathsf{out}})$. Completeness follows easily, because $\mathsf{ct} \oplus f(\theta, \mathsf{z}) \oplus \mathsf{k} = \pi_{\mathsf{in}}$ which is accepted by $\mathsf{V}_{\mathsf{nizk}}(\mathsf{crs}_{\mathsf{in}}, x, \pi_{\mathsf{in}})$ by completeness of Π_{nizk}. Next, consider the case when x is false, which means there is no accepting proof $\widetilde{\pi}_{\mathsf{in}}$ that makes $\mathsf{V}_{\mathsf{nizk}}(\mathsf{crs}_{\mathsf{in}}, x, \cdot)$ accept (whp over choice of $\mathsf{crs}_{\mathsf{in}}$). As a result, all of the meta-statements x^{z} would also be false and hence would be rejected whp by $\mathsf{V}_{\mathsf{nizk}}(\mathsf{crs}_{\mathsf{out}}, x^{\mathsf{z}}, \cdot)$. Note that due to adaptive soundness of Π_{nizk}, it is not a problem even if x^{z} are computed adversarially based on $\mathsf{crs}_{\mathsf{out}}$. As for computational ZK, it follows since $(\mathsf{crs}_{\mathsf{in}}, \pi_{\mathsf{in}})$ can be replaced with $(\mathsf{crs}_{\mathsf{Sim}}, \pi_{\mathsf{Sim}})$ which is guaranteed by the computational ZK of Π_{nizk}. We also remark that since the scheme proves

[2] This uniform value is used due to a technicality in reducing to the theorem of [BK23].

statements involving the description of $\mathsf{V}_{\mathsf{nizk}}$, it makes non-black-box use of Π_{nizk}. Hereafter, we will focus on achieving the certified-everlasting zero-knowledge (CE-ZK) property.

Reduction to Certified-Deletion Theorem. In order to enforce CE-ZK, we require the verifier to return (revoke) the entire proof state $\sigma = (\sum_{\mathsf{z}} \alpha_{\mathsf{z}} |\mathsf{z}\rangle_{\mathrm{R}} |\pi^{\mathsf{z}}_{\mathsf{out}}\rangle_{\mathrm{P}}, \mathsf{ct})$. Then, the certification algorithm Cert uncomputes the proofs $\pi^{\mathsf{z}}_{\mathsf{out}}$ in superposition, followed by measuring register R in the Hadamard basis to get a classical outcome cert. Finally, Cert checks if cert is consistent with y at the Hadamard basis positions specified by θ. Now, our goal is to prove CE-ZK by reducing to the aforementioned certified-deletion theorem of BB84 states [BK23]. This theorem is more general than previously discussed in that a QPT adversary $\mathcal{A}$ can be given an arbitrary state $|\psi\rangle$ that is generated by some distribution $\mathcal{Z}(\theta, b \oplus p, |y\rangle^{\theta})$ where $y, \theta \leftarrow \{0,1\}^{\lambda}$, $p := \bigoplus_{i:\theta[i]=0} y[i]$ and $b \in \{0,1\}$. Then, conditioned on QPT $\mathcal{A}$ outputting a valid certificate cert (consistent with the Hadamard positions), the advantage in distinguishing $b = 0/1$ using $\mathcal{A}$'s leftover state is negligible, even for unbounded distinguishers. However, this guarantee only holds for distributions $\mathcal{Z}$ which computationally hide the basis information θ, i.e., the pre-condition $\mathcal{Z}(\theta, b', |\phi\rangle) \approx_c \mathcal{Z}(0^{\lambda}, b', |\phi\rangle)$ must hold for all $\theta, b', |\phi\rangle$. Next, we describe how this distribution $\mathcal{Z}$ is defined in our context:

$\underline{\mathcal{Z}(\theta, b', |\phi\rangle):}$

- The state $|\phi\rangle$ is parsed as $|\phi\rangle = \sum_{\mathsf{z}} \alpha_{\mathsf{z}} |\mathsf{z}\rangle_{\mathrm{R}}$ on register R.
- Then, the proof state $\sigma := (\sum_{\mathsf{z}} \alpha_{\mathsf{z}} |\mathsf{z}\rangle_{\mathrm{R}} |\pi^{\mathsf{z}}_{\mathsf{out}}\rangle_{\mathrm{P}}, \mathsf{ct})$ is computed using (θ, b') (mostly similar to the construction, but we expand on it shortly).
- Next, the QPT V^* (attacking CE-ZK) is run on σ, producing a state $\rho_{\mathrm{A}\otimes\mathrm{B}}$ on registers A, B, with A denoting the "revocation" register.
- Let ρ_{A} denote the state on register A at this point, which is parsed as a state on registers $\mathrm{R} \otimes \mathrm{P}$ (recall that honest revocation returns the proof).
- Finally, the proofs $\pi^{\mathsf{z}}_{\mathsf{out}}$ on register P are uncomputed, resulting in the final state $\rho'_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{B}}$ on the three registers. This state ρ' is then output.

The context here is as follows. Consider a QPT malicious verifier V^* (which defines $\mathcal{Z}$) and an unbounded distinguisher D breaking CE-ZK. Then, in the BB84 deletion game wrt $\mathcal{Z}$, the reduction $\mathcal{R}$ receives the state $\rho'_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{B}} \leftarrow \mathcal{Z}(\theta, b \oplus p, |y\rangle^{\theta})$, after which it simply measures register R in the Hadamard basis to produce cert and outputs the state on register B. If V^* passes certification in the CE-ZK game, then cert will be a valid certificate in the BB84 game. We will discuss shortly that this allows to statistically switch a bit in the CE-ZK game, as otherwise, D would distinguish $b = 0/1$ in the BB84 game using the output of $\mathcal{R}$. Before that, we discuss how the pre-condition is satisfied.

The key idea we use to establish the pre-condition is that even though $\mathcal{Z}$ depends on θ as the witness ω_{out} for generating the proofs $\pi^{\mathsf{z}}_{\mathsf{out}}$, the computational ZK of Π_{nizk} can be used to simulate these proofs without access to θ. Note however, that there are exponentially many proofs in superposition, generated wrt the same CRS $\mathsf{crs}_{\mathsf{out}}$. Still, we can rely on the superposition-secure adaptive ZK property (Definition 3.6) for this simulation. Consequently, we define $\mathcal{Z}$ (See Lemma 5.4) so that it generates simulated proofs in superposition (instead of the real proofs $\pi^{\mathsf{z}}_{\mathsf{out}}$), when its first input is 0^{λ}. Note that this simulation can also handle the uncomputation step above, allowing to show $\mathcal{Z}(\theta, b', |\phi\rangle) \approx_c \mathcal{Z}(0^{\lambda}, b', |\phi\rangle)$. This in-turn allows to switch a bit $b = 1$ to $b = 0$ in the CE-ZK security game, in a *statistically indistinguishable* way (conditioned on successful certification). Next, we elaborate on what this bit flip actually does in our context.

Utilizing Two-Slots. The purpose of the statistically indistinguishable bit flip is to "zero-out" the inner proof π_{in} (i.e., replacing it with 0^λ) in a bit-by-bit manner, thereby removing the dependence on the witness ω. The problem with this bit flip is that even though the proofs $\pi^{\mathsf{z}}_{\mathsf{out}}$ satisfy the NP relation wrt the statements $x^{\mathsf{z}}_{\mathsf{out}}$ and witness $\omega_{\mathsf{out}} = (\theta, \mathsf{k})$ when considering the distribution $\mathcal{Z}(\theta, b' = 1, |\phi\rangle)$, this no longer holds for $\mathcal{Z}(\theta, b' = 0, |\phi\rangle)$. This is because the unmasked inner proof π'_{in} in the latter case may no longer make $\mathsf{V}_{\mathsf{nizk}}(\mathsf{crs}_{\mathsf{in}}, x, \pi'_{\mathsf{in}})$ accept. For this reason, we use a $2\ell\lambda$-qubit BB84 state instead and modify the meta-statement with a dummy slot, described as follows:

$$\underline{\text{Meta-Statement } x^{\mathsf{z}}}: \quad \exists\left(\theta, \mathsf{k}^0, \mathsf{k}^1\right) \ : \ \mathsf{V}_{\mathsf{nizk}}\Big(\mathsf{crs}_{\mathsf{in}}, x, \mathsf{ct}^0 \oplus f^0(\theta, \mathsf{z}) \oplus \mathsf{k}^0\Big) = 1$$
$$\text{OR}$$
$$\mathsf{V}_{\mathsf{nizk}}\Big(\mathsf{crs}_{\mathsf{in}}, x, \mathsf{ct}^1 \oplus f^1(\theta, \mathsf{z}) \oplus \mathsf{k}^1\Big) = 1$$

where f^0 is defined similarly to f for the last ℓ-qubits (which encode pad^0) and f^1 corresponds to the first ℓ-qubits (which encode pad^1), with $\mathsf{ct}^0 := \pi_{\mathsf{in}} \oplus \mathsf{pad}^0 \oplus \mathsf{k}^0$ and $\mathsf{ct}^1 := 0^\lambda \oplus \mathsf{pad}^1 \oplus \mathsf{k}^1$.

Note that the reason for performing a check using $\mathsf{V}_{\mathsf{nizk}}$ in the dummy slot is because we cannot simply prove a trivial statement in the simulation. For example, we cannot have the real CRS be a commitment to 1 and use a commitment to 0 in the simulation. This is because the simulated CRS distribution must be statistically close to the real CRS distribution (because the classical CRS cannot be "deleted"), on top of the proof being statistically sound. Consequently, we will rely on the non-adaptive computational ZK with statistical CRS indistinguishability property (See Definition 3.6) of Π_{nizk}. Intuitively, this means that the marginal distribution of the simulated CRS $\mathsf{crs}_{\mathsf{Sim}}$ is statistically close to that of the real CRS $\mathsf{crs}_{\mathsf{in}}$, while the overall ZK guarantee is still computational. We observe in Section 3.4 that the LWE-based NIZK proofs of [Wat24, WWW25] obtained via the hidden-bits approach satisfy this property. For current LWE-based NIZK proofs obtained via the Fiat-Shamir approach [CCH+19, PS19], this is not satisfied as the CRS includes a "fake" hash key that is only computationally close to a real key.

Using this property, we first switch $\mathsf{crs}_{\mathsf{in}}$ into $\mathsf{crs}_{\mathsf{Sim}}$ statistically. Then, due to the nature of the OR meta-statement, we are able to swap the 0^ℓ plaintext in ct^1 into a simulated proof π_{Sim} (generated along with $\mathsf{crs}_{\mathsf{Sim}}$) one bit at a time. This is because we no longer have the aforementioned issue as π_{in} inside ct^0 is sufficient to satisfy the OR condition, and hence invoke superposition-secure zero-knowledge of the outer proofs. Then, once π_{Sim} is placed in ct^1, we can replace π_{in} bit-by-bit into 0^ℓ since π_{Sim} also satisfies $\mathsf{V}_{\mathsf{nizk}}(\mathsf{crs}_{\mathsf{Sim}}, x, \pi_{\mathsf{Sim}}) = 1$ by the computational zero-knowledge of Π_{nizk}. We remark that this is the only place where computational ZK *for the inner layer proof* is used in the proof of CE-ZK. Finally, observe that the resulting distribution can be produced with the statement x alone (without ω), ensuring that the CE-ZK property is satisfied.

Some Subtleties. We remark that instead of the inner proof π_{in}, we could encrypt the witness ω directly in the construction and check the NP relation in the first slot of the meta-statement. However, we chose the current variant as we anyway use the $\mathsf{V}_{\mathsf{nizk}}$ check for the second slot and the proof of ZK is simpler in this case.

Additionally, apart from a NIZK for NP, we rely on a quantum-secure PRF (implied by OWFs) to simulate uniform randomness in superposition, used to compute and uncompute the outer NIZK proofs.

We also face the following issue in the reduction to the deletion theorem. The quantum state returned by the adversary may also have superposition terms $\widetilde{\mathsf{z}}$ such that $\omega_{\mathsf{out}} = (\theta, \mathsf{k}^0, \mathsf{k}^1)$ is not a valid witness for statement $x^{\widetilde{\mathsf{z}}}$. Consequently, we cannot argue that the state after uncomputation of real proofs is computationally close to the state after uncomputation of simulated proofs, based on the superposition-secure ZK property. To circumvent this, we employ a simpler (information-theoretic) variant of the key-testability technique of [KNY23, KNP25]. At a high level, this involves producing an information-theoretic MAC tag for each term of the BB84 state, which can be verified after revocation. Intuitively, this ensures that the adversary cannot produce a state with new superposition terms (that have non-negligible amplitude) without breaking unforgeability of the MAC. In summary, we obtain the following theorem:

Theorem 2.3 (CE-NIZK in the CRS Model)**.** *There exist CE-NIZKs for* NP *in the CRS model satisfying 1) statistical soundness, 2) computational zero-knowledge, and 3) certified-everlasting zero-knowledge, assuming the existence of OWFs and NIZKs for* NP *in the CRS model satisfying 1) statistical (adaptive) soundness, 2) computational (non-adaptive) zero-knowledge with statistical CRS indistinguishability, and 3) superposition-secure (adaptive) zero-knowledge.*

Since such NIZKs are implied by LWE (See Section 3.4) due to [WWW25], we have:

Corollary 2.4 (CE-NIZK from LWE)**.** *There exist CE-NIZKs for* NP *in the CRS model based on the polynomial hardness of the LWE problem.*

2.5 Quantum-Efficient Protocol in the Shared EPR Model

Previously, we argued that generalizations of the Fiat-Shamir and hidden-bits paradigms to the CE-NIZK setting do not work, as they run into Theorem 2.2. We observe that this impossibility does not apply to the shared EPR model where the parties are initialized with halves of polynomially-many EPR pairs. Such a model has recently garnered attention in the context of non-interactive MPC [ABKK23, BKS23] and NIZK for QMA [BKS23, MY22a]. Recall that Theorem 2.2 relies on the following idea: a prover could split the quantum proof the verifier is supposed to receive into two parts, one which passes certification while the other suffices for verification. Hence, the prover need only send the part which is required for verification. This then implies a statistically secure NIZK. In the EPR model, the prover cannot perform this "split" as it does not control the entire quantum proof state, since the verifier's EPR halves do not originate from the prover.

Our observation has some similarities to the work of [BKS23], which separated the CRS and shared EPR models in the context of non-interactive randomized oblivious transfer. In our case, the separation between the models only applies to the aforementioned deletion-resistant protocols, as evident from our CRS feasibility result. However, we demonstrate that this difference is useful as well, by constructing a protocol in the shared EPR model that is much more efficient in its use of quantum resources. Specifically, our protocol only involves performing single-qubit computational or Hadamard basis measurements on the EPR halves, and classical computation otherwise. We proceed to describe our approach.

We make use of the well-known hidden-bits paradigm [FLS90, FLS99] that we adapt to the shared EPR model for the context of obtaining certified-everlasting ZK. We note that Fiat-Shamir variants of previous interactive protocols would also bypass the impossibility in the EPR model. However, we find the hidden-bits approach to be cleaner and also helps to avoid the statistically far simulated CRS issue with Fiat-Shamir approaches mentioned earlier. We also emphasize that we assume a common reference string in addition to the shared EPRs. This is because it is not clear how to simulate a structured reference string using EPRs unlike a common random string.

Moreover, statistically sound NIZKs from LWE [CCH+19, PS19, Wat24, WWW25] are all in the common reference string model. While we could consider a more complex entangled state that can simulated arbitrary reference strings, we stick with EPR setups for their simplicity and practicality. We believe that the NIZK proof of Bitansky and Paneth [BP15] based on iO and OWFs would also work for our purpose, allowing to use EPRs alone as it only requires a uniform CRS.

The hidden bits paradigm for constructing NIZK utilizes two components: (1) a hidden bit generator (HBG) and (2) a statistically sound and statistically zero-knowledge NIZK in the hidden-bits model. In the (idealized) hidden-bits model, there is a uniformly random string that is sampled and made available to the prover, but is hidden from the verifier. Then, the prover has the ability to specify some indices of its choice to a trusted third party, which then reveals the corresponding hidden bits to the verifier (and nothing about the other bits). Since NIZKs in the hidden-bits model are known unconditionally [FLS90, FLS99], most works focus on designing an HBG. Intuitively, the purpose of the HBG is to realize the idealized hidden bits in the CRS model. In particular, it allows the prover P to sample a succinct commitment com, which along with the CRS, fixes a hidden bit string. The HBG then allows the prover to provably open certain positions of this fixed hidden bit string to the verifier. Crucially, the *statistical-binding* security of the HBG ensures that the prover is bound to this string, and cannot reveal anything else to the verifier. Moreover, *computational hiding* ensures that all unopened hidden bits are pseudorandom in the verifier's view.

Our idea is that instead of determining the actual hidden bits string via the commitment com and the CRS crs, these are utilized to fix a common basis string θ for both parties. Then, the prover would reveal parts of θ to the verifier, and both parties would obtain the actual hidden bits by measuring their halves of the EPR pairs in the θ basis. Specifically, the EPRs are divided into polynomially many "blocks". Then, the i-th block is used to derive the i-th hidden bit, set to be the XOR of computational basis measurements of the positions j where $\theta[j] = 0$. Note that the nature of EPR pairs ensures that both parties would derive the same measurement outcomes on measuring in the same basis. Also, statistical soundness is ensured due to the observation that we can consider an equivalent verifier that measures all its EPRs in the computational basis before the prover is even initialized. Consequently, the prover can only manipulate the hidden bits via the basis string θ which is restricted by the HBG commitment com. Note that the inability of this hypothetical verifier to pass certification is not a problem, as it is only used as a proof tool to argue soundness. While computational zero-knowledge follows easily by arguments similar to prior work [QRW19], certified-everlasting zero-knowledge is argued based on the following fact. The verifier is asked to delete all the unopened hidden bits by measuring the corresponding EPRs in the Hadamard basis. Consequently, even if computational hiding of the HBG eventually leaks θ to an unbounded distinguisher, the certified-deletion theorem of [BK23] can be applied to ensure that these hidden bits are information-theoretically hidden. This is because shared EPR pairs exhibit properties similar to BB84 states. Consequently, we have the following theorem:

Theorem 2.5. *CE-NIZKs for* NP *exist in the shared EPR Model (Definition 4.3), assuming statistical binding hidden bit generators (Definition 6.1) exist.*

From the fact that statistical binding HBGs are known from LWE [WWW25], and due to the structure of our protocol, we have the following corollary:

Corollary 2.6. *Assuming the polynomial hardness of LWE, there exist CE-NIZKs for* NP *in the shared EPR Model, where the prover and verifier only perform single-qubit measurements apart from arbitrary classical computation.*

2.6 Relation with Other Quantum-Enabled NIZKs

Unclonable NIZKs. The works of [JK24] and [GMR24] introduced unclonable NIZK which is a primitive that at a high level, guarantees the following property. Consider an adversary that is provided with a quantum proof of some NP statement x. Then, if it can produce two separate quantum proofs that both pass verification, one can efficiently extract out a witness w for the statement x from the adversary. Intuitively, this ensures that if an adversary only has the proof and not the witness, it cannot produce more than one valid proof. Clearly, this is only feasible in the quantum setting due to the no-cloning principle. In [JK24], it was also informally mentioned that such proofs are a prerequisite for obtaining NIZKs with some form of certified deletion, while also mentioning certified everlasting security. The intuition is that if the adversary could clone proof states, then it could delete one of them successfully while retaining the other. Moreover, these works show that unclonable NIZK implies public-key quantum money (PKQM). On the other hand, we establish CE-NIZK in this paper from LWE, which is not known to imply PKQM. The reasoning is that CE-NIZKs need not be unclonable in the sense of [JK24], which actually doesn't require the adversary to generate two copies of the state. To see this, observe the following attack on our CE-NIZK in the CRS model that clearly breaks unclonable security. The adversary on obtaining the proof state $\sigma = \sum_{\mathsf{t}} \alpha_{\mathsf{t}} |\mathsf{t}\rangle$, simply applies CNOTs to get the state $\sum_{\mathsf{t}} \alpha_{\mathsf{t}} |\mathsf{t}\rangle |\mathsf{t}\rangle$ and gives the two entangled registers to two verifiers. Since verification is done entirely in the computational basis, both verifiers would accept the proof. Natural "unentangled" attacks also work, which involving collapsing σ and sending the outcome to both verifiers. We also mention that by utilizing coset states along with iO (which implies PKQM assuming OWFs) instead of BB84 states, one might be able to obtain both unclonability and CE-ZK. Note that unclonability also doesn't imply CE-ZK in general, because the unclonability adversary never receives unbounded computational power.

NIZKs with Certified-Deletion. In the related work of [AK25], a NIZK with certified-deletion (NIZK-CD) was obtained from standard assumptions, specifically from a classical NIZK (with certain properties), and a statistically binding commitment. The security intuitively ensures that if an adversary is able to pass the deletion check and simultaneously produce an accepting proof, then one can efficiently extract a witness from the adversary. We will discuss a natural weakening of this definition, where the adversary should simply be unable to output a valid certificate and a valid proof, for statement-witness pairs drawn from some hard distribution. We refer to this weaker variant as wNIZK-CD. We consider this variant for comparison because CE-NIZK is not defined with an extractability guarantee in our work. We now argue informally CE-NIZK implies wNIZK-CD.

Remark 2.7. *In all our discussions regarding wNIZK-CD and NIZK-CD, we assume the schemes are statistically sound, as it is the focus of our work.*

CE-NIZK $\Rightarrow$ wNIZK-CD. Consider some language $L \in \mathsf{NP}$ and some sufficiently hard distribution D over statement-witness pairs. Let $\Pi := (\mathsf{Setup}, \mathsf{P}, \mathsf{V}, \mathsf{Cert})$ be a CE-NIZK scheme for L. Consider a QPT $\mathcal{A}$ that is given $\sigma \leftarrow \mathsf{P}(\mathsf{crs}, x, \omega)$ for $(x, \omega) \leftarrow D$, after which it produces $\rho_{\mathrm{A}\otimes\mathrm{B}}$ such that $\mathsf{Cert}(\rho_{\mathsf{P}}, \rho_{\mathrm{A}}) = 1$ and $\mathsf{V}(\mathsf{crs}, x, \rho_{\mathrm{B}}) = 1$. Assume for now that this condition holds with probability 1. We can now construct a QPT prover P^* that on input $(x, \omega) \leftarrow D$, computes $\sigma \leftarrow \mathsf{P}(\mathsf{crs}, x, \omega)$ and $\rho_{\mathrm{A}\otimes\mathrm{B}} \leftarrow \mathcal{A}(\sigma)$, followed by applying $\mathsf{Cert}(\rho_{\mathsf{P}}, \rho_{\mathrm{A}})$ and finally outputting ρ_{B} as the proof. Observe that the NIZK $(\mathsf{Setup}, \mathsf{P}^*, \mathsf{V})$ satisfies completeness (on average wrt D), statistical soundness (a requirement of CE-NIZK), and statistical ZK (on average, and from the CE-ZK of Π). We believe that such a proof system leads to contradiction. To see why, recall that Pass and

Shelat [PS05] show that if a language L has a NIZK with both statistical soundness and statistical ZK, then $L \in \mathsf{AM} \cap \mathsf{coAM}$. Even if the event $\mathsf{Cert}(\rho_{\mathsf{P}}, \rho_{\mathrm{A}}) = 1 \wedge \mathsf{V}(\mathsf{crs}, x, \rho_{\mathrm{B}}) = 1$ occurs with only inverse polynomial probability, thereby implying inverse polynomial variants of completeness and statistical ZK, the proof of [PS05] goes through. This is because the completeness of their coAM protocol follows from $\mathsf{negl}(\lambda)$-soundness of the NIZK, while inverse polynomial soundness follows from the completeness and statistical ZK of the NIZK. We note that a similar result would also hold for analogous (average case) quantum complexity classes, which are unlikely to contain sufficiently hard (L, D).

Next, we argue the reverse direction does not hold.

wNIZK-CD $\not\Rightarrow$ CE-NIZK. Consider a QPT verifier V^* and unbounded distinguisher D that break the CE-ZK property of Π. The problem is that D is unbounded, but a wNIZK-CD adversary must efficiently generate an accepting proof. Another issue is that D need not recover the entire witness. Hence, it is unclear how a reduction would work. We believe explicit counterexamples can also be shown, by considering statistically sound extensions of existing NIZK-CD schemes [AK25]. In other words, these NIZK-CD schemes would not satisfy CE-ZK. This is because the classical NIZK proofs attached to the quantum proof state in these schemes can be compromised by unbounded attackers.

2.7 Open Questions

1. *Publicly-Verifiable Deletion:* While our construction only satisfies privately-verifiable deletion, obtaining the publicly-verifiable variant is more desirable. We believe this should be possible by employing coset states [CLLZ21] instead of BB84 states, along with iO (by employing techniques from [BGK+24]). Note that coset states with iO and OWFs are sufficient to obtain public-key quantum money (PKQM). However, it is not clear that CE-NIZK with publicly-verifiable deletion implies PKQM because an efficient PKQM verifier may be unable to exploit the fact that an adversarial state contains some information about the witness. Still, it is not obvious how to obtain publicly-verifiable deletion without iO, and we leave it as an interesting open question.

2. *Classical Certificates*: Our construction in the CRS model requires quantum revocation, while it is ideal if the verifier could only send back a classical certificate. One approach towards achieving this would be to utilize a hybrid quantum fully homomorphic encryption (QFHE) scheme that is known from LWE [Mah20, Bra18], as used in the works of [Shm22a, Shm22b]. In particular, the verifier could be provided with a QFHE ciphertext that encrypts $(\theta, \mathsf{k}^0, \mathsf{k}^1)$ so that it could perform the uncomputation of the superposition proofs using QFHE evaluation. This way, it would be able to obtain a certificate that is a QFHE encryption of Hadamard basis measurements, which could be decrypted and checked by the prover. The problem here is that we cannot simply replace the QFHE plaintext with a dummy, followed by using a similar certified-deletion argument as [BK23]. This is because deletion verification requires the QFHE secret-key, making it unclear how to reduce to the semantic security of QFHE. Despite this, there might be new techniques that would allow to prove the security of such a construction, and we leave this for future work.

3. *Certified-Everlasting NIWIs in the Plain Model*: It is also interesting to consider other non-interactive proof systems such as NIWIs, which are possible to obtain even in the plain model [DN00, BOV03]. In the post-quantum setting, these can be based on post-quantum iO [BP15].

In particular, one can try to obtain a NIWI with statistical soundness along with computational witness indistinguishability (WI) and certified-everlasting WI in the plain model. Our construction seems to be general enough to work for this case, except that the outer proofs generated in superposition can no longer be simulated by the computational zero-knowledge property. Recall that this is an important part of our proof, as it is needed to invoke the certified deletion theorem. It might still be possible just to rely on the witness indistinguishability property. However, we think this approach would require a stronger deletion theorem, and possibly techniques from leakage-resilient cryptography [KR19, ÇGLZR24].

4. *CE-NIZKs for QMA*: In this work, we focus on obtaining CE-NIZK for NP, due to the challenges in obtaining NIZK for QMA in the CRS model (See [Shm21]). On the other hand, interactive CE-ZK is known for all of QMA [HMNY22]. In this regard, it would be interesting to explore CE-NIZK for QMA in other setup models. For inspiration, the works of [Shm21] and [MY22a] show NIZKs for QMA in the designated-verifier and shared-EPR models respectively, by utilizing variants of the Fiat-Shamir paradigm. On the other hand, one could consider CRS model constructions that make heuristic use of iO [BJRT26].

3 Preliminaries

3.1 Notation

We denote a (possibly quantum) algorithm P with input 1^λ and outputs m, ct as $\mathsf{P}(1^\lambda) \to (\mathsf{m}, \mathsf{ct})$ or $(\mathsf{m}, \mathsf{ct}) \leftarrow \mathsf{P}(1^\lambda)$. Here, λ denotes a security parameter and $1^\lambda := 1\|\dots(\lambda \text{ times})\|1$, where $\|$ denotes concatenation and $:=$ denotes defining (also used for assigning/substituting). We use $\mathsf{P}(1^\lambda)[1]$ and $\mathsf{P}(1^\lambda)[2]$ to denote the first and second indices of the output tuple respectively. For a set S, by $s \leftarrow S$, we mean s to be sampled uniformly from S. For a distribution D, by $s \leftarrow D$, we mean sampling s according to D. The functions $\mathsf{poly}(\lambda)$ and $\mathsf{negl}(\lambda)$ denote a polynomial and a negligible function in λ respectively. We use the notation $\mathcal{A}^{|\mathsf{P}(\mathsf{crs},\cdot,\cdot)\rangle}$ to denote that the quantum algorithm $\mathcal{A}$ gets quantum access to a quantum algorithm $\mathsf{P}(\mathsf{crs},\cdot,\cdot)$ (with classical inputs and outputs) on the second and third inputs. For a language $\mathcal{L} \in \mathsf{NP}$, we consider the relation $R_\mathcal{L}$ consisting of statement-witness pairs. Hence, $R_\mathcal{L}(x)$ denotes the set of witnesses of a statement x. We denote quantum registers as $\mathrm{A}, \mathrm{B}, \mathrm{R}$ etc. We use the notation y_i to denote the i-th bit of a string y (unless it is explicitly defined to be something else), and R_i to denote the i-th qubit of R. Sometimes, we also use $y[i]$ to denote the i-th bit. Additionally, we use the notation y_I to denote $\{y_i\}_{i\in I}$ for some set I. Note that y^i doesn't mean the i-th bit of y. However, the notation y^I denotes $\{y^i\}_{i\in I}$. Sometimes, we use $\overline{I}$ to denote the complement of a set I, when the universe is clear from context. We also employ the notation $s_I := t_I \oplus r_I$ to denote $s_i := t_i \oplus r_i$ for each $i \in I$. For a classical algorithm C, we use $\mathsf{C}(s; r)$ to denote running it on input s and random tape r. We denote pure quantum states as $|\psi\rangle, |\mathsf{z}\rangle$ etc while the quantum states ρ, σ will either refer to pure or mixed states based on the context. We use the notation $|\psi\rangle_\mathrm{R}$ to denote a state on register R and $\rho_{\mathrm{A}\otimes\mathrm{B}}$ to denote a bi-partite state on register $\mathrm{A}\otimes\mathrm{B}$. We denote the state on B obtained by tracing out A as $\rho_\mathrm{B} := \mathsf{Tr}_\mathrm{A}(\rho_{\mathrm{A}\otimes\mathrm{B}})$. For simplicity, we often use ρ_B to denote the state on register B at different points when it is clear from context. For e.g., we do not use a different ρ'_B to refer to the state obtained after A is collapsed to some specific value. An EPR pair is the following state on a pair of registers P, V: $\frac{1}{\sqrt{2}}\left(|0\rangle_\mathrm{P}|0\rangle_\mathrm{V} + |1\rangle_\mathrm{P}|1\rangle_\mathrm{V}\right)$, where P, V are often referred to as EPR halves. For two distributions D_0, D_1, by the notation $D_0 \approx_c D_1$, we mean the two are computationally indistinguishable by quantum polynomial time (QPT) adversaries with non-uniform quantum advice. Often, we simply use $d_0 \approx_c d_1$ for samples $d^0 \leftarrow D_0, d^1 \leftarrow D_1$ drawn from these distributions to denote the same. We

use $D_0 \equiv D_1$ to denote the distributions are identical and $D_0 \approx_s D_1$ to denote they are statistically close. Likewise, for quantum states ρ and σ, $\rho \approx_s \sigma$ denotes that the states are negligibly close in trace distance.

3.2 Quantum Information

A pure state $|\psi\rangle$ is a vector in a complex Hilbert space $\mathcal{H}$ with unit norm $\||\psi\rangle\| = 1$. An operator P which satisfies $P^\dagger = P$ is said to be a Hermitian operator. Let $S(\mathcal{H})$ be the set of Hermitian operators on $\mathcal{H}$. A density matrix $\rho \in S(\mathcal{H})$ is a positive semi-definite operator with unit trace $\mathsf{Tr}(\rho) = 1$. A density matrix represents a mixture over pure states, and the density matrix of a pure state $|\psi\rangle$ is $|\psi\rangle\langle\psi|$. A Hilbert space can be split into registers $\mathcal{H} := \mathcal{H}_{\mathrm{R}_1} \otimes \mathcal{H}_{\mathrm{R}_2}$. A unitary matrix U is one that satisfies $UU^\dagger = U^\dagger U = \mathbb{I}$. Applying U to a state $|\psi\rangle$ transforms it into the state $U|\psi\rangle$. For a mixed state ρ, applying U transforms it into $U\rho U^\dagger$. A projector Π is a Hermitian operator additionally satisfying $\Pi^2 = \Pi$. The trace distance between two mixed states ρ, σ is given by $\mathsf{TD}(\rho, \sigma) := \frac{1}{2}\mathsf{Tr}(\sqrt{(\rho-\sigma)^\dagger(\rho-\sigma)})$. It provides an upper bound on the distinguishing advantage of ρ, σ by any quantum algorithm. A quantum polynomial-time (QPT) algorithm refers to a polynomial-size quantum circuit composed of some universal gate set, along with some non-uniform advice state $|\psi\rangle$. We denote the Hadamard transform/operator by H, which is defined as $H = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$. For $y, \theta \in \{0,1\}^n$ for some polynomial n in λ, we consider BB84 (Wiesner) states of the following form, where y_i, θ_i denote the i-th bits of y, θ respectively:

$$|y\rangle^\theta := H^{\theta_1}|y_1\rangle \otimes \ldots \otimes H^{\theta_n}|y_n\rangle$$

3.3 Cryptography Basics

Definition 3.1 (Learning with Errors (LWE) [Reg05])**.** Consider integer functions $n, m, q \in \mathbb{Z}$ of the security parameter λ. Let $\chi(\lambda)$ be an error distribution over $\mathbb{Z}$. The LWE assumption is said to hold if for every QPT adversary $\mathcal{A}$, there exists a negligible function $n(\lambda)$ such that:

$$\left| \Pr\left[\mathcal{A}(A, y) \to 1 \;\middle|\; \begin{array}{l} A \leftarrow \mathbb{Z}_q^{n\times m} \\ s \leftarrow \mathbb{Z}_q^n \\ e \leftarrow \chi^m \\ y := s^\intercal A + e \end{array} \right] - \Pr\left[\mathcal{A}(A, u) \to 1 \;\middle|\; \begin{array}{l} A \leftarrow \mathbb{Z}_q^{n\times m} \\ u \leftarrow \mathbb{Z}_q^m \end{array} \right] \right| \leq n(\lambda)$$

Definition 3.2 (Quantum-Secure Pseudo-Random Function)**.** Let $\{\mathsf{F}_{\mathsf{prfk}} : \{0,1\}^{\ell_{\mathsf{in}}} \to \{0,1\}^{\ell_{\mathsf{out}}}\}_{\mathsf{prfk}\in\{0,1\}^\lambda}$ be an efficiently computable family of functions, where $\ell_{\mathsf{in}}, \ell_{\mathsf{out}}$ are some polynomials in λ. Then, $\{\mathsf{F}_{\mathsf{prfk}}\}_{\mathsf{prfk}}$ is said to be a quantum-secure PRF if for every QPT adversary $\mathcal{A}$, there exists a negligible function $n(\lambda)$ such that the following holds, where $\mathcal{R}$ is the set of all functions with input ℓ_{in} and output ℓ_{out}:

$$\left| \Pr\left[1 \leftarrow \mathcal{A}^{|f(\cdot)\rangle}(1^\lambda) \mid f \leftarrow \mathcal{R}\right] - \Pr\left[1 \leftarrow \mathcal{A}^{|\mathsf{F}_{\mathsf{prfk}}(\cdot)\rangle}(1^\lambda) \mid \mathsf{prfk} \leftarrow \{0,1\}^\lambda\right] \right| \leq n(\lambda)$$

Theorem 3.3 ([Zha12])**.** *Quantum-secure PRFs exist, assuming post-quantum secure one-way functions exist.*

Lemma 3.4 (One-Way to Hiding (O2H) [AHU19])**.** *Consider any functions $G, H : \mathcal{X} \to \mathcal{Y}$, an arbitrary value z, and some set $S \subseteq \mathcal{X}$ such that for every $x \notin S$, $G(x) = H(x)$. The values*

G, H, S, z may be arbitrarily distributed. Consider now an oracle algorithm $\mathcal{A}$ that makes at most q queries to its oracle. Let $\mathcal{B}$ be an oracle algorithm such that $\mathcal{B}^H$ executes as follows on input z:

$\underline{\mathcal{B}^H(z):}$

1. *Choose $i \leftarrow [q]$.*
2. *Execute $\mathcal{A}^H(z)$ until it makes its i-th query.*
3. *Measure the i-th query and output the outcome.*

Then, the following holds:

$$\left| \Pr\left[\mathcal{A}^H(z) = 1\right] - \Pr\left[\mathcal{A}^G(z) = 1\right] \right| \leq 2q \cdot \sqrt{\Pr\left[\mathcal{B}^H(z) \in S\right]}$$

Remark 3.5. *It is easy to see from the proof of [AHU19] that the lemma holds even when the auxiliary input z is some quantum state ρ_z. A similar observation was made in [HXY19], where it was mentioned that z could even be exponentially large, and hence allows to describe quantum states as well.*

3.4 Superposition-Secure Non-Interactive Zero-Knowledge

Definition 3.6. A triple of PPT algorithms $(\mathsf{Setup}, \mathsf{P}, \mathsf{V})$ is a superposition-secure non-interactive zero-knowledge proof (NIZK) in the common reference string (CRS) model for a language $\mathcal{L}$, if it satisfies the following properties:

- *Completeness:* The following holds for every $x \in \mathcal{L}$ and $\omega \in R_{\mathcal{L}}(x)$:

$$\Pr\left[\mathsf{V}(\mathsf{crs}, x, \pi) \to 0 \ : \ \begin{array}{c} \mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ \pi \leftarrow \mathsf{P}(\mathsf{crs}, x, \omega) \end{array}\right] \leq \mathsf{negl}(\lambda)$$

- *Statistical (Non-Adaptive) Soundness:* For every $\widetilde{x} \notin \mathcal{L}$, the following holds:

$$\Pr_{\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda)}[\exists \widetilde{\pi} : \mathsf{V}(\mathsf{crs}, \widetilde{x}, \widetilde{\pi}) \to 1] \leq \mathsf{negl}(\lambda).$$

- *Statistical Adaptive Soundness:*

$$\Pr_{\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda)}[\exists (\widetilde{x}, \widetilde{\pi}) : \widetilde{x} \notin \mathcal{L} \wedge \mathsf{V}(\mathsf{crs}, \widetilde{x}, \widetilde{\pi}) \to 1] \leq \mathsf{negl}(\lambda).$$

- *Computational (Non-Adaptive) Zero-Knowledge with Statistical CRS Indistinguishability:* There exists a QPT simulator Sim such that for all $x \in \mathcal{L}$ and $\omega \in R_{\mathcal{L}}(x)$, the following holds wrt non-uniform QPT distinguishers:

$$\left[\begin{array}{c} (x, \mathsf{crs}, \pi) \approx_c (x, \mathsf{crs}', \pi') \\ \wedge \\ (x, \mathsf{crs}) \approx_s (x, \mathsf{crs}') \end{array} \ : \ \begin{array}{c} \mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ \pi \leftarrow \mathsf{P}(\mathsf{crs}, x, \omega) \\ (\mathsf{crs}', \pi') \leftarrow \mathsf{Sim}(x) \end{array}\right]$$

Remark 3.7. *Statistical CRS indistinguishability cannot hold for the CRS output by S_1 of an adaptive ZK simulator $(\mathsf{S}_1, \mathsf{S}_2)$ of a statistically-sound NIZK. This was mentioned in [Gro06], and can be seen from the proof of [Pas13].*

- *Superposition-Secure (Adaptive) Zero-Knowledge:* There exists a PPT two-part simulator $(\mathsf{S}_1, \mathsf{S}_2)$ such that the following holds for every QPT adversary $\mathcal{A}$ that only queries *valid* statement-witness pairs in superposition, i.e., terms of the form $|x, \omega\rangle$ where $x \in \mathcal{L} \wedge \omega \in R_{\mathcal{L}}(x)$:

$$\left[\begin{array}{c} \mathcal{A}^{|\mathsf{S}_2(\mathsf{td},\cdot)\rangle}(\mathsf{crs}') \\ \approx_c \\ \mathcal{A}^{|\mathsf{P}(\mathsf{crs},\cdot,\cdot)\rangle}(\mathsf{crs}) \end{array} : \begin{array}{l} \mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ (\mathsf{crs}', \mathsf{td}) \leftarrow \mathsf{S}_1(1^\lambda) \end{array}\right]$$

Theorem 3.8 ([QRW19, WWW25, ACE+24]). *Assuming the polynomial hardness of the LWE problem, there exist NIZKs for* NP *in the CRS model with the following properties:*

- *Statistical (Adaptive) Soundness*
- *Superposition-Secure (Adaptive) Zero-Knowledge*
- *Computational (Non-Adaptive) ZK with Statistical CRS Indistinguishability*

In more detail, the work of [QRW19] showed that any statistical-binding hidden bit generator (HBG) (Definition 6.1) implies a statistically-sound NIZK. The works of [Wat24, WWW25] constructed such HBGs from LWE with sub-exponential and polynomial modulus respectively. In the work of [ACE+24], it was argued that the NIZK of [BP15] from iO and OWFs satisfies the stronger notion of superposition-secure (adaptive) zero-knowledge. We observe that their reasoning also applies to the NIZK of [WWW25] and the compiler of [QRW19] due to the following reason:

The high level intuition is that these constructions simulate the idealized hidden bits model (HBM) [FLS90] in the CRS model. Note that the idealized HBM provides *perfect* zero-knowledge [FLS99, HU19] (along with statistical soundness), and hence also superposition-secure zero-knowledge. The observation made by [ACE+24] is that the HBM to CRS model transformation also plays well with superposition queries. This is because the proof of the zero-knowledge property only requires switching the CRS distribution computationally, before appealing to the *perfect* zero-knowledge guarantee of NIZK in the HBM. Since the CRS is classical and independent of the statement-witness queries, the quantum-security implication follows. For more details, see Appendix D of [ACE+24].

Finally, we observe that in the HBG to NIZK compiler of [QRW19], the CRS in the construction is computed as $(\mathsf{crs}_{\mathsf{bg}}, s)$ where the former is the CRS of the HBG and s is a truly random string. In the simulation $\mathsf{Sim}(x)$, $\mathsf{crs}_{\mathsf{bg}}$ is computed the same way, but s is computed as $s := r_{\mathsf{bg}} \oplus r_{\mathsf{hb}}$ where r_{bg} is a pseudo-random string generated by the HBG, while r_{hb} is a uniform hidden-bit string output by the hidden bits simulator $\mathsf{Sim}_{\mathsf{hb}}(x)$. Hence, statistical CRS indistinguishability follows by the perfect zero-knowledge guarantee of the hidden bits NIZK proof.

3.5 Certified Deletion of BB84 States

Theorem 3.9 (Certified Deletion of BB84 States [BK23]). *Consider a distribution $\mathcal{Z}(\cdot,\cdot,\cdot)$ with three arguments: 1) a λ-bit string θ, 2) a bit b, and 3) a λ-qubit register* D*. For all $\theta \in \{0,1\}^\lambda$, $b' \in \{0,1\}$ and quantum states $|\psi\rangle_{\mathrm{D}\otimes\mathrm{C}}$ on a λ-qubit register* D *and an arbitrary size register* C*, let the following hold for all QPT distinguishers $\mathcal{D}$:*

$$\left|\Pr\left[\mathcal{D}\left(1^\lambda, \mathcal{Z}(\theta, b', \mathrm{D}), \mathrm{C}\right) = 1\right] - \Pr\left[\mathcal{D}\left(1^\lambda, \mathcal{Z}(0^\lambda, b', \mathrm{D}), \mathrm{C}\right) = 1\right]\right| \leq \mathsf{negl}(\lambda)$$

In other words, $\mathcal{Z}$ is semantically-secure wrt θ. Consider now the following experiment wrt an adversary $\mathcal{A}$:

$\underline{\mathsf{Exp}_{\mathcal{Z},\mathcal{A}}(b)}$:

1. *The experiment samples $y, \theta \leftarrow \{0,1\}^\lambda$ and initializes $\mathcal{A}$ with input 1^λ along with the following input:*
$$\mathcal{Z}\left(\theta, b \oplus \bigoplus_{i:\theta_i=0} y_i, |y\rangle^\theta\right)$$
2. *$\mathcal{A}$ sends a string $\mathsf{cert} \in \{0,1\}^\lambda$ and a quantum state ρ_{B} on register B.*
3. *If $\forall i \in [\lambda]$ such that $\theta_i = 1$, $\mathsf{cert}_i = y_i$, then the experiment outputs ρ_{B}. Else, it outputs $\bot$.*

Then, the following guarantee holds for all QPT adversaries $\mathcal{A}$:

$$\mathsf{TD}\left(\mathsf{Exp}_{\mathcal{Z},\mathcal{A}}(0), \mathsf{Exp}_{\mathcal{Z},\mathcal{A}}(1)\right) = \mathsf{negl}(\lambda)$$

4 Certified-Everlasting NIZK Proofs

Definition 4.1 (CE-NIZK in the CRS Model)**.** A certified-everlasting non-interactive zero-knowledge proof (CE-NIZK) in the CRS model for a language $\mathcal{L}$ is a tuple of four algorithms $(\mathsf{Setup}, \mathsf{P}, \mathsf{V}, \mathsf{Cert})$ with the following syntax:

- $\mathsf{Setup}(1^\lambda) \to \mathsf{crs}$: The classical setup algorithm takes as input a security parameter, and outputs a common reference string (CRS) crs.
- $\mathsf{P}(\mathsf{crs}, x, \omega) \to (\sigma, \rho_{\mathsf{P}})$: The quantum prover's algorithm takes as input a CRS crs, a statement $x \in \mathcal{L}$, and a witness $\omega \in R_{\mathcal{L}}(x)$. It outputs a quantum proof state σ and a residual state ρ_{P}.
- $\mathsf{V}(\mathsf{crs}, x, \sigma) \to (b, \rho_{\mathsf{V}})$: The quantum verifier's algorithm takes as input the CRS crs, a statement $x \in \mathcal{L}$, and a quantum proof σ. It outputs a bit b along with a residual state ρ_{V}.
- $\mathsf{Cert}(\rho_{\mathsf{P}}, \rho_{\mathsf{V}}) \to \top/\bot$: The certification algorithm takes as input a prover's state ρ_{P} and a verifier's state ρ_{V}. It outputs $\top$ (accept) or $\bot$ (reject).

A CE-NIZK must satisfy the following properties:

- *Completeness:* The following holds for every $x \in \mathcal{L}$ and $\omega \in R_{\mathcal{L}}(x)$:
$$\Pr\left[\mathsf{V}\big(\mathsf{crs}, x, \sigma\big)[1] \to 0 \ : \ \begin{array}{l}\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ (\sigma, \rho_{\mathsf{P}}) \leftarrow \mathsf{P}(\mathsf{crs}, x, \omega)\end{array}\right] \leq \mathsf{negl}(\lambda)$$
- *Deletion Correctness:* The following holds for every $x \in \mathcal{L}$ and $\omega \in R_{\mathcal{L}}(x)$:
$$\Pr\left[\mathsf{Cert}\big(\rho_{\mathsf{P}}, \rho_{\mathsf{V}}\big) \to \bot \ : \ \begin{array}{l}\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ (\sigma, \rho_{\mathsf{P}}) \leftarrow \mathsf{P}(\mathsf{crs}, x, \omega) \\ (b, \rho_{\mathsf{V}}) \leftarrow \mathsf{V}(\mathsf{crs}, x, \sigma)\end{array}\right] \leq \mathsf{negl}(\lambda)$$

- *Statistical Soundness:* For every $\widetilde{x} \notin \mathcal{L}$, the following holds:

$$\Pr_{\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda)} [\exists \widetilde{\sigma} : \mathsf{V}(\mathsf{crs}, \widetilde{x}, \widetilde{\sigma})[1] \to 1] \leq \mathsf{negl}(\lambda).$$

- *Computational Zero-Knowledge:* There exists a simulator Sim such that for all $x \in \mathcal{L}$ and $\omega \in R_{\mathcal{L}}(x)$, the following holds:

$$\left[(x, \mathsf{crs}, \sigma) \approx_c (x, \mathsf{crs}', \sigma') \ : \begin{array}{l} \mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ (\sigma, \rho_{\mathsf{P}}) \leftarrow \mathsf{P}(\mathsf{crs}, x, \omega) \\ (\mathsf{crs}', \sigma') \leftarrow \mathsf{Sim}(x) \end{array} \right]$$

- *Certified-Everlasting Zero-Knowledge (CE-ZK):* For every malicious QPT verifier V^*, there exists a QPT simulator $\mathsf{Sim}_{\mathsf{V}^*}$ such that the following holds for every $x \in \mathcal{L}$ and $\omega \in R_{\mathcal{L}}(x)$:

$$\left[\widetilde{\rho} \approx_s \rho' \ : \begin{array}{l} \mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ (\sigma, \rho_{\mathsf{P}}) \leftarrow \mathsf{P}(\mathsf{crs}, x, \omega) \\ \rho_{\mathrm{A}\otimes\mathrm{B}} \leftarrow \mathsf{V}^*(\mathsf{crs}, x, \sigma) \\ \text{If } \mathsf{Cert}(\rho_{\mathsf{P}}, \rho_{\mathrm{A}}) = \top \text{ then } \widetilde{\rho} \coloneqq \rho_{\mathrm{B}}. \\ \text{Else } \widetilde{\rho} \coloneqq \bot. \\ \rho' \leftarrow \mathsf{Sim}_{\mathsf{V}^*}(x) \end{array} \right]$$

  Here, $\rho_{\mathrm{A}\otimes\mathrm{B}}$ is a possibly entangled state on two registers A and B. We denote the residual state on A (likewise B) at any point as ρ_{A} (likewise ρ_{B}).

  Remark 4.2. *The CE-ZK definition describes non-black-box simulation (for generality in Theorem 2.2), but our feasibility uses only black-box simulation.*

Definition 4.3 (CE-NIZK in the Shared EPR Model)**.** Let $(\mathrm{P}_1, \mathrm{V}_1) \dots (\mathrm{P}_\ell, \mathrm{V}_\ell)$ be registers denoting halves of ℓ-many EPR pairs. Let the prover's register be defined as $\mathrm{P} \coloneqq \mathrm{P}_1 \otimes \dots \otimes \mathrm{P}_\ell$ and the verifier's register as $\mathrm{V} \coloneqq \mathrm{V}_1 \otimes \dots \otimes \mathrm{V}_\ell$. A CE-NIZK in the shared EPR model for a language $\mathcal{L} \in \mathsf{NP}$ is a tuple of five algorithms $(\mathsf{Setup}, \mathsf{P}, \mathsf{V}, \mathsf{Del}, \mathsf{Cert})$ with the following syntax:

- $\mathsf{Setup}(1^\lambda) \to \mathsf{crs}$: The classical setup algorithm takes as input a security parameter, and outputs a common reference string (CRS) crs.

- $\mathsf{P}(\mathsf{crs}, \mathrm{P}, x, \omega) \to (\pi, \rho_{\mathsf{P}})$: The quantum prover's algorithm takes as input a CRS crs, the prover register P, a statement $x \in \mathcal{L}$, and a witness $\omega \in R_{\mathcal{L}}(x)$. It outputs a classical proof π and a residual state ρ_{P} on a register $\mathrm{P}' \otimes \mathrm{P}$ where P is the EPR register and P' is a new register.

- $\mathsf{V}(\mathsf{crs}, \mathrm{V}, x, \pi) \to (b, \rho_{\mathsf{V}})$: The quantum verifier's algorithm takes as input a CRS crs, the verifier register V, a statement $x \in \mathcal{L}$, and a classical proof π. It outputs a bit b along with a residual state ρ_{V} on register $\mathrm{V}' \otimes \mathrm{V}$ where V is the EPR register and V' is a new register.

- $\mathsf{Del}(\rho_{\mathsf{V}}) \to (\mathsf{cert}, \widetilde{\rho}_{\mathsf{V}})$: The quantum deletion algorithm takes as input a state ρ_{V} on register $\mathrm{V}' \otimes \mathrm{V}$, and outputs a classical certificate cert, and state $\widetilde{\rho}_{\mathsf{V}}$.

- $\mathsf{Cert}(\mathsf{cert}, \rho_{\mathsf{P}}) \to \top/\bot$: The certification algorithm takes as input a classical certificate cert and a quantum state ρ_{P} on $\mathrm{P}' \otimes \mathrm{P}$. It outputs $\top$ or $\bot$.

A CE-NIZK in the shared EPR model must satisfy the following properties:

- *Completeness:* The following holds for every $x \in \mathcal{L}$ and $\omega \in R_{\mathcal{L}}(x)$:

$$\Pr\left[\mathsf{V}\big(\mathsf{crs}, \mathrm{V}, x, \pi\big)[1] \to 0 \ : \ \begin{array}{l} \mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ (\pi, \rho_{\mathsf{P}}) \leftarrow \mathsf{P}(\mathsf{crs}, \mathrm{P}, x, \omega) \end{array}\right] \leq \mathsf{negl}(\lambda)$$

- *Deletion Correctness:* The following holds for every $x \in \mathcal{L}$ and $\omega \in R_{\mathcal{L}}(x)$:

$$\Pr\left[\mathsf{Cert}\big(\mathsf{cert}, \rho_{\mathsf{P}}\big) \to \bot \ : \ \begin{array}{l} \mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ (\pi, \rho_{\mathsf{P}}) \leftarrow \mathsf{P}(\mathsf{crs}, \mathrm{P}, x, \omega) \\ (b, \rho_{\mathsf{V}}) \leftarrow \mathsf{V}(\mathsf{crs}, \mathrm{V}, x, \pi) \\ (\mathsf{cert}, \widetilde{\rho}_{\mathsf{V}}) \leftarrow \mathsf{Del}(\rho_{\mathsf{V}}) \end{array}\right] \leq \mathsf{negl}(\lambda)$$

- *Statistical Soundness:* For every $\widetilde{x} \notin \mathcal{L}$ and every unbounded P^*, we have:

$$\Pr\left[\mathsf{V}\big(\mathsf{crs}, \mathrm{V}, \widetilde{x}, \pi^*\big)[1] \to 1 \ : \ \begin{array}{l} \mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ (\pi^*, \rho^*_{\mathsf{P}}) \leftarrow \mathsf{P}^*(\mathsf{crs}, \mathrm{P}, \widetilde{x}) \end{array}\right] \leq \mathsf{negl}(\lambda)$$

- *Computational Zero-Knowledge:* There exists a simulator Sim such that for all $x \in \mathcal{L}$ and $\omega \in R_{\mathcal{L}}(x)$, the following holds:

$$\left[(x, \mathsf{crs}, \mathrm{V}, \pi) \approx_c (x, \mathsf{crs}', \mathrm{V}', \pi') \ : \ \begin{array}{l} \mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ (\pi, \rho_{\mathsf{P}}) \leftarrow \mathsf{P}(\mathsf{crs}, \mathrm{P}, x, \omega) \\ (\mathsf{crs}', \mathrm{V}', \pi') \leftarrow \mathsf{Sim}(x) \end{array}\right]$$

- *Certified-Everlasting Zero-Knowledge (CE-ZK):* There exists a QPT simulator Sim such that for every malicious verifier V^*, statement $x \in \mathcal{L}$ and witness $\omega \in R_{\mathcal{L}}(x)$, the following holds:

$$\left[\widetilde{\rho} \approx_s \rho' \ : \ \begin{array}{l} \mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ (\pi, \rho_{\mathsf{P}}) \leftarrow \mathsf{P}(\mathsf{crs}, \mathrm{P}, x, \omega) \\ (\mathsf{cert}^*, \rho^*) \leftarrow \mathsf{V}^*(\mathsf{crs}, \mathrm{V}, x, \pi) \\ \text{If } \mathsf{Cert}(\mathsf{cert}^*, \rho_{\mathsf{P}}) = \top \text{ then } \widetilde{\rho} := \rho^\star. \\ \text{Else } \widetilde{\rho} := \bot. \\ \rho' \leftarrow \mathsf{Sim}^{\mathsf{V}^*}(x) \end{array}\right]$$

  Here, ρ^* is a state on $\mathrm{V} \otimes \mathrm{V}'$ where V' is a new register, and $\mathsf{Sim}^{\mathsf{V}^*}$ denotes black-box access to V^* on its input and setup registers.

5 CE-NIZK in the CRS Model

5.1 Impossibility of Deletion-Resistant CE-NIZKs

Theorem 5.1. *If there exists a deletion-resistant CE-NIZK (Definition 2.1) for a language $\mathcal{L}$, then there exists a quantum NIZK for $\mathcal{L}$ satisfying statistical soundness and statistical zero-knowledge.*

Proof. Let Π be a deletion-resistant CE-NIZK for a language $\mathcal{L}$. This means it has algorithms $(\mathsf{Setup}, \mathsf{P}, \mathsf{V}, \mathsf{Cert})$. Furthermore, there exist QPT algorithms $\widetilde{\mathsf{V}}^*, \widetilde{\mathsf{V}}$ satisfying the conditions of the theorem by assumption. We now define the following prover $\widetilde{\mathsf{P}}$. Note that we can assume $\widetilde{\mathsf{P}}$ to be initialized with polynomially-many copies of the initial quantum states of P and V^*.

$\widetilde{\mathsf{P}}(\mathsf{crs}, x, \omega)$:

1. Compute $(\sigma, \rho_{\mathsf{P}}) \leftarrow \mathsf{P}(\mathsf{crs}, x, \omega)$.
2. Compute $\rho_{\mathrm{A}\otimes\mathrm{B}} \leftarrow \mathsf{V}^*(\sigma)$.
3. If $\mathsf{Cert}(\rho_{\mathsf{P}}, \rho_{\mathrm{A}}) = \top$, then output ρ_{B}. Otherwise, repeat from Step 1.

Observe now that the NIZK proof $\widetilde{\Pi} := (\mathsf{Setup}, \widetilde{\mathsf{P}}, \widetilde{\mathsf{V}})$ is statistically-sound due to condition (3) of Definition 2.1. Also, completeness follows from the fact that $\widetilde{\mathsf{P}}$ will obtain $\rho_{\mathrm{A}\otimes\mathrm{B}}$ such that $\mathsf{Cert}(\rho_{\mathsf{P}}, \rho_{\mathrm{A}}) = \top$ with overwhelming probability in polynomial time (due to condition (1) of Definition 2.1), and the fact that $\widetilde{\mathsf{V}}$ accepts ρ_{B} from condition (2). Finally, observe that by viewing V^* as a malicious QPT verifier attacking the CE-ZK property of Π, the state ρ_{B} output by $\widetilde{\mathsf{P}}$ can be statistically simulated. This means that $\widetilde{\Pi}$ satisfies statistical zero-knowledge. □

5.2 Construction from LWE

We now construct a CE-NIZK $\Pi = (\mathsf{Setup}, \mathsf{P}, \mathsf{V}, \mathsf{Cert})$ as follows:

Certified-Everlasting NIZK in the CRS Model

Building Blocks:

1. NIZK $\Pi_{\mathsf{nizk}} := (\mathsf{Setup}_{\mathsf{nizk}}, \mathsf{P}_{\mathsf{nizk}}, \mathsf{V}_{\mathsf{nizk}})$ for NP in the CRS model with:
   - Statistical Adaptive Soundness
   - Superposition-Secure Adaptive Zero-Knowledge
   - Non-Adaptive ZK with Statistical CRS Indistinguishability
2. BB84 States $|y\rangle^{\theta} := H^{\theta_1}|y_1\rangle \otimes \ldots \otimes H^{\theta_\lambda}|y_\lambda\rangle$, where $y, \theta \in \{0,1\}^\lambda$.
3. Quantum-Secure PRF $\{\mathsf{F}_{\mathsf{prfk}} : \{0,1\}^{\mathsf{poly}(\lambda)} \to \{0,1\}^{\mathsf{poly}(\lambda)}\}_{\mathsf{prfk}\in\{0,1\}^\lambda}$.

$\mathsf{Setup}(1^\lambda)$:

- Compute $\mathsf{crs}_{\mathsf{in}} \leftarrow \mathsf{Setup}_{\mathsf{nizk}}(1^\lambda)$ and $\mathsf{crs}_{\mathsf{out}} \leftarrow \mathsf{Setup}_{\mathsf{nizk}}(1^\lambda)$.
- Output $\mathsf{crs} := \mathsf{crs}_{\mathsf{in}} \| \mathsf{crs}_{\mathsf{out}}$.

$\mathsf{P}(\mathsf{crs}, x, \omega)$:

1. Let ℓ be an upper bound on the proof size of Π_{nizk}. For each $i \in [2\ell]$, sample $y^i, \theta^i \leftarrow \{0,1\}^\lambda$. Then, construct the following BB84 state on register $\mathrm{R}^{\mathrm{i}} := \mathrm{R}^{\mathrm{i}}_1 \otimes \ldots \otimes \mathrm{R}^{\mathrm{i}}_\lambda$.

$$|y^i\rangle^{\theta^i}_{\mathrm{R}^{\mathrm{i}}} := H^{\theta^i_1}|y^i_1\rangle_{\mathrm{R}^{\mathrm{i}}_1} \otimes \ldots \otimes H^{\theta^i_\lambda}|y^i_\lambda\rangle_{\mathrm{R}^{\mathrm{i}}_\lambda}$$

2. Parse $\mathsf{crs} = \mathsf{crs}_{\mathsf{in}} \| \mathsf{crs}_{\mathsf{out}}$. Then, compute $\pi_{\mathsf{in}} \leftarrow \mathsf{P}_{\mathsf{nizk}}(\mathsf{crs}_{\mathsf{in}}, x, \omega)$.

3. Sample $\mathsf{k}^0, \mathsf{k}^1 \leftarrow \{0,1\}^\ell$ and compute $\mathsf{ct}^0 := \pi_{\mathsf{in}} \oplus \mathsf{pad}^0 \oplus \mathsf{k}^0$ and $\mathsf{ct}^1 := 0^\ell \oplus \mathsf{pad}^1 \oplus \mathsf{k}^1$, where:

$$\mathsf{pad}^1 := \bigoplus_{i:\theta_i^1=0} y_i^1 \ \| \ldots \| \bigoplus_{i:\theta_i^\ell=0} y_i^\ell; \ \ \mathsf{pad}^0 := \bigoplus_{i:\theta_i^{\ell+1}=0} y_i^{\ell+1} \ \| \ldots \| \bigoplus_{i:\theta_i^{2\ell}=0} y_i^{2\ell}$$

4. For each $\mathsf{z} \in \{0,1\}^{2\ell\lambda}$, let $\mathsf{z} := \mathsf{z}^1 \| \ldots \| \mathsf{z}^{2\ell}$ where $|\mathsf{z}^i| = \lambda$ for each $i \in [2\ell]$. Then, define the following NP statement x^{z}:

   Statement x^{z}: There exist $\{\theta^i\}_{i\in[2\ell]}$ and $\mathsf{k}^0, \mathsf{k}^1 \in \{0,1\}^\ell$ s.t.

$$\mathsf{V}_{\mathsf{nizk}}\left(\mathsf{crs}_{\mathsf{in}}, x, \mathsf{ct}^0 \oplus \mathsf{k}^0 \oplus \widetilde{\mathsf{pad}}^0\right) = 1 \bigvee \mathsf{V}_{\mathsf{nizk}}\left(\mathsf{crs}_{\mathsf{in}}, x, \mathsf{ct}^1 \oplus \mathsf{k}^1 \oplus \widetilde{\mathsf{pad}}^1\right) = 1$$

$$\widetilde{\mathsf{pad}}^1 := \bigoplus_{i:\theta_i^1=0} \mathsf{z}_i^1 \ \| \ldots \| \bigoplus_{i:\theta_i^\ell=0} \mathsf{z}_i^\ell; \ \ \widetilde{\mathsf{pad}}^0 := \bigoplus_{i:\theta_i^{\ell+1}=0} \mathsf{z}_i^{\ell+1} \ \| \ldots \| \bigoplus_{i:\theta_i^{2\ell}=0} \mathsf{z}_i^{2\ell}$$

5. For each $i \in [2\ell\lambda]$, sample $s^{i,0}, s^{i,1} \leftarrow \{0,1\}^\lambda$. For each $\mathsf{z} \in \{0,1\}^{2\ell\lambda}$:

$$\mathsf{sig}^{\mathsf{z}} := s^{1,\mathsf{z}[1]} \| \ldots \| s^{2\ell\lambda,\mathsf{z}[2\ell\lambda]}$$

6. Define the register $\mathrm{R} := \mathrm{R}^1 \otimes \ldots \otimes \mathrm{R}^{2\ell}$. Let the state on register R be $\sum_{\mathsf{z}} \alpha^{\mathsf{z}} |\mathsf{z}\rangle$. Then, compute the following state $|\psi\rangle$ on registers R, P and S:

$$|\psi\rangle_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}} := \sum_{\mathsf{z}} \alpha^{\mathsf{z}} |\mathsf{z}\rangle_{\mathrm{R}} |\pi^{\mathsf{z}}_{\mathsf{out}}\rangle_{\mathrm{P}} |\mathsf{sig}^{\mathsf{z}}\rangle_{\mathrm{S}}$$

   where $\pi^{\mathsf{z}}_{\mathsf{out}} \leftarrow \mathsf{P}_{\mathsf{nizk}}(\mathsf{crs}_{\mathsf{out}}, x^{\mathsf{z}}, \omega_{\mathsf{out}}; r^{\mathsf{z}})$, $\omega_{\mathsf{out}} := (\{\theta^i\}_{i\in[2\ell]}, \mathsf{k}^0, \mathsf{k}^1)$ and explicit randomness $r^{\mathsf{z}} := \mathsf{F}_{\mathsf{prfk}}(\mathsf{z})$ for a PRF key sampled as $\mathsf{prfk} \leftarrow \{0,1\}^\lambda$.

7. Output the quantum proof $\sigma := \left(|\psi\rangle_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}}, \mathsf{ct}^0, \mathsf{ct}^1\right)$ and the verification state $\rho_{\mathsf{P}} := \left(\omega_{\mathsf{out}}, \mathsf{crs}, \{y^i\}_{i\in[2\ell]}, \mathsf{prfk}, \{s^{i,0}, s^{i,1}\}_{i\in[2\ell\lambda]}, \mathsf{ct}^0, \mathsf{ct}^1\right)$.

$\underline{\mathsf{V}(\mathsf{crs}, x, \sigma) :}$

1. Parse crs as $\mathsf{crs} = \mathsf{crs}_{\mathsf{in}} \| \mathsf{crs}_{\mathsf{out}}$ and $\sigma = (|\psi\rangle_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}}, \mathsf{ct}^0, \mathsf{ct}^1)$.
2. Let $|\psi\rangle_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}} = \sum_{\mathsf{z}} \alpha^{\mathsf{z}} |\mathsf{z}\rangle_{\mathrm{R}} |\pi^{\mathsf{z}}_{\mathsf{out}}\rangle_{\mathrm{P}} |\mathsf{sig}^{\mathsf{z}}\rangle_{\mathrm{S}}$. Compute the following state on registers $\mathrm{R}, \mathrm{P}, \mathrm{S}, \mathrm{OUT}$:

$$|\psi\rangle'_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}\otimes\mathrm{OUT}} := \sum_{\mathsf{z}} \alpha^{\mathsf{z}} |\mathsf{z}\rangle_{\mathrm{R}} |\pi^{\mathsf{z}}_{\mathsf{out}}\rangle_{\mathrm{P}} |\mathsf{sig}^{\mathsf{z}}\rangle_{\mathrm{S}} |\mathsf{V}_{\mathsf{nizk}}(\mathsf{crs}_{\mathsf{out}}, x^{\mathsf{z}}, \pi^{\mathsf{z}}_{\mathsf{out}})\rangle_{\mathrm{OUT}}$$

3. Measure the register OUT in the computational basis to obtain outcome $b \in \{0,1\}$ and a residual state $|\widetilde{\psi}\rangle_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}}$.
4. Set $\rho_{\mathsf{V}} := |\widetilde{\psi}\rangle_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}}$. Finally, output (b, ρ_{V}).

$\underline{\mathsf{Cert}(\rho_{\mathsf{P}}, \rho_{\mathsf{V}}) :}$

1. Parse $\rho_{\mathsf{P}} = \left(\omega_{\mathsf{out}}, \mathsf{crs}, \{y^i\}_{i\in[2\ell]}, \mathsf{prfk}, \{s^{i,0}, s^{i,1}\}_{i\in[2\ell\lambda]}, \mathsf{ct}^0, \mathsf{ct}^1\right)$ and $\rho_{\mathsf{V}} = |\widetilde{\psi}\rangle_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}}$.
2. Perform $|\mathsf{z}\rangle_{\mathrm{R}}\,|\widetilde{\mathsf{sig}}^{\mathsf{z}}\rangle_{\mathrm{S}}\,|0\rangle_{\mathrm{T}} \mapsto |\mathsf{z}\rangle_{\mathrm{R}}\,|\widetilde{\mathsf{sig}}^{\mathsf{z}}\rangle_{\mathrm{S}}\,|0 \oplus \mathsf{Test}(\mathsf{z}, \widetilde{\mathsf{sig}}^{\mathsf{z}})\rangle_{\mathrm{T}}$ where T is initialized to 0 and $\mathsf{Test}(\mathsf{z}, \widetilde{\mathsf{sig}}^{\mathsf{z}})$ is a function that outputs 1 if $\mathsf{sig}^{\mathsf{z}} = \widetilde{\mathsf{sig}}^{\mathsf{z}}$ and 0 otherwise. Measure register T and output $\perp$ if the outcome is 0.
3. Perform $|\mathsf{z}\rangle_{\mathrm{R}}\,|\mathsf{t}\rangle_{\mathrm{P}}\,|\mathsf{s}\rangle_{\mathrm{S}} \mapsto |\mathsf{z}\rangle_{\mathrm{R}}\,|\mathsf{t} \oplus \mathsf{P}_{\mathsf{nizk}}(\mathsf{crs}_{\mathsf{out}}, x^{\mathsf{z}}, \omega_{\mathsf{out}}; r^{\mathsf{z}})\rangle_{\mathrm{P}}\,|\mathsf{s} \oplus \mathsf{sig}^{\mathsf{z}}\rangle_{\mathrm{S}}$ where $r^{\mathsf{z}} = \mathsf{F}_{\mathsf{prfk}}(\mathsf{z})$ to get the state $|\widetilde{\phi}\rangle_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}}$.
4. Measure the R register of $|\widetilde{\phi}\rangle_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}}$ in the Hadamard basis to get cert.
5. Let $\theta := \theta^1\|\ldots\|\theta^{2\ell}$ and $y := y^1\|\ldots\|y^{2\ell}$. If $\mathsf{cert}_j = y_j$ for each $j \in [2\ell\lambda]$ such that $\theta_j = 1$, then output $\top$. Otherwise, output $\perp$.

Theorem 5.2. *Assuming Π_{nizk} is a NIZK proof for NP with statistical adaptive soundness, superposition-secure (adaptive) zero-knowledge and computational (non-adaptive) zero-knowledge with statistical CRS indistinguishability (Definition 3.6), and $\{\mathsf{F}_{\mathsf{prfk}}\}_{\mathsf{prfk}}$ is a quantum-secure PRF family, then $\Pi := (\mathsf{Setup}, \mathsf{P}, \mathsf{V}, \mathsf{Cert})$ is a secure CE-NIZK scheme for NP (Definition 4.1).*

Remark 5.3. *Definition 4.1 only enforces non-adaptive soundness and computational ZK. However, it is easy to see that the stronger adaptive notions follow for Π from those of the underlying NIZK Π_{nizk}, by nature of the construction.*

Proof.

Completeness: Observe that $\mathsf{V}_{\mathsf{nizk}}(\mathsf{crs}_{\mathsf{in}}, x, \mathsf{ct}^0 \oplus \widetilde{\mathsf{pad}}^0 \oplus \mathsf{k}^0)$ outputs 1 for each superposition term z. This is because unmasking ct^0 with k^0 will result in $\pi_{\mathsf{in}} \oplus \mathsf{pad}^0$. Moreover, $\widetilde{\mathsf{pad}}^0$ corresponding to every superposition term z is equal to pad^0. Next, observe that $\mathsf{V}_{\mathsf{nizk}}(\mathsf{crs}_{\mathsf{in}}, x, \pi_{\mathsf{in}})$ outputs 1 with overwhelming probability by the completeness of Π_{nizk}. Consequently, with all but negligible probability, all the statements x^{z} in superposition are valid. Finally, by completeness of Π_{nizk}, we have that $\mathsf{V}_{\mathsf{nizk}}(\mathsf{crs}_{\mathsf{out}}, x^{\mathsf{z}}, \pi^{\mathsf{z}}_{\mathsf{out}})$ outputs 1 with overwhelming probability for each term z, assuming the randomness for computing $\pi^{\mathsf{z}}_{\mathsf{out}}$ is obtained from a quantum-secure PRF $\mathsf{F}_{\mathsf{prfk}}$. Consequently, computing the result of verification on register OUT and measuring the result produces outcome 1 with overwhelming probability.

Deletion-Correctness: Recall that V is sent the state $|\psi\rangle_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}} = \sum_{\mathsf{z}} \alpha^{\mathsf{z}}\,|\mathsf{z}\rangle_{\mathrm{R}}\,|\pi^{\mathsf{z}}_{\mathsf{out}}\rangle_{\mathrm{P}}\,|\mathsf{sig}^{\mathsf{z}}\rangle_{\mathrm{S}}$ and the post-verification state $|\widetilde{\psi}\rangle$ is close in trace distance to $|\psi\rangle$ due to the above completeness argument and the gentle measurement lemma. Next, observe that the signature checks using Test pass in superposition for every term z and hence the state is undisturbed. Then, Cert uncomputes the proofs on registers P, S, resulting in the state $|\widetilde{\phi}\rangle$ that is close in trace distance to the state $|y\rangle^{\theta}$. Consequently, measuring the registers of $|\widetilde{\phi}\rangle$ in the Hadamard basis results in a value cert that is consistent with y at all the Hadamard basis positions with overwhelming probability. As a result, Cert outputs 1 with overwhelming probability.

Computational Zero-Knowledge: This follows immediately from the fact that Π_{nizk} has a simulator $\mathsf{Sim}_{\mathsf{nizk}}$ such that the values $(\mathsf{crs}_{\mathsf{Sim}}, \pi_{\mathsf{Sim}}) \leftarrow \mathsf{Sim}_{\mathsf{nizk}}(x)$ are computationally indistinguishable

from $\mathsf{crs}_{\mathsf{in}}, \pi_{\mathsf{in}}$. Consequently, we can replace π_{in} inside ct^0 with π_{Sim} and $\mathsf{crs}_{\mathsf{in}}$ (which is part of crs) with $\mathsf{crs}_{\mathsf{Sim}}$.

Statistical Soundness: For every false statement x', with overwhelming probability over the choice of $\mathsf{crs}_{\mathsf{in}}$, we know that there doesn't exist any π' such that $\mathsf{V}_{\mathsf{nizk}}(\mathsf{crs}_{\mathsf{in}}, x', \pi') = 1$ due to the statistical soundness of Π_{nizk}. Consequently, with overwhelming probability over $\mathsf{crs}_{\mathsf{in}}$, all the corresponding statements x_{z} in superposition are false as well. Now, these statements are decided after $\mathsf{crs}_{\mathsf{out}}$ is chosen. Consequently, we rely on the adaptive soundness property to ensure that with overwhelming probability over $\mathsf{crs}_{\mathsf{out}}$, there do not exist pairs $(x^{\mathsf{z}}, \pi^{\mathsf{z}}_{\mathsf{out}})$ that would be accepted by $\mathsf{V}_{\mathsf{nizk}}(\mathsf{crs}_{\mathsf{out}}, x^{\mathsf{z}}, \pi^{\mathsf{z}}_{\mathsf{out}})$. Consequently, verification outputs 0 with overwhelming probability, ensuring statistical soundness.

Certified-Everlasting Zero-Knowledge: We now proceed to prove the CE-ZK property using the following sequence of hybrids:

$\underline{\mathsf{Hyb}_0}$:

- V^* is initialized with x and $\mathsf{crs} = \mathsf{crs}_{\mathsf{in}} \| \mathsf{crs}_{\mathsf{out}} \leftarrow \mathsf{Setup}(1^\lambda)$.
- Then, V^* is provided with $\sigma = (|\psi\rangle_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}}, \mathsf{ct}^0, \mathsf{ct}^1)$ computed as $(\sigma, \rho_{\mathsf{P}}) \leftarrow \mathsf{P}(\mathsf{crs}, x, \omega)$, where $\mathsf{ct}^0 = \pi_{\mathsf{in}} \oplus \mathsf{pad}^0 \oplus \mathsf{k}^0$ and $\mathsf{ct}^1 = 0^\ell \oplus \mathsf{pad}^1 \oplus \mathsf{k}^1$.
- Then, V^* outputs a bi-partite state $\rho_{\mathrm{A}\otimes\mathrm{B}}$. If $\mathsf{Cert}(\rho_{\mathsf{P}}, \rho_{\mathrm{A}}) = \top$, output ρ_{B} where ρ_{A} is the state on register A and ρ_{B} is the residual state on register B. Output $\perp$ otherwise.

$\underline{\mathsf{Hyb}_1}$:

- V^* is initialized with x and $\mathsf{crs}_{\mathsf{Sim}} \| \mathsf{crs}_{\mathsf{out}}$, where $(\mathsf{crs}_{\mathsf{Sim}}, \pi_{\mathsf{Sim}}) \leftarrow \mathsf{Sim}_{\mathsf{nizk}}(x)$ and $\mathsf{crs}_{\mathsf{out}} \leftarrow \mathsf{Setup}_{\mathsf{nizk}}(1^\lambda)$, with $\mathsf{Sim}_{\mathsf{nizk}}$ being the QPT simulator of Π_{nizk}.
- Then, V^* is provided with $\sigma = (|\psi\rangle_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}}, \mathsf{ct}^0, \mathsf{ct}^1)$ computed as $(\sigma, \rho_{\mathsf{P}}) \leftarrow \mathsf{P}(\mathsf{crs}_{\mathsf{Sim}} \| \mathsf{crs}_{\mathsf{out}}, x, \omega)$, where $\mathsf{ct}^0 = \pi_{\mathsf{in}} \oplus \mathsf{pad}^0 \oplus \mathsf{k}^0$ and $\mathsf{ct}^1 = 0^\ell \oplus \mathsf{pad}^1 \oplus \mathsf{k}^1$.
- Then, V^* outputs a bi-partite state $\rho_{\mathrm{A}\otimes\mathrm{B}}$. If $\mathsf{Cert}(\rho_{\mathsf{P}}, \rho_{\mathrm{A}}) = \top$, output ρ_{B} where ρ_{A} is the state on register A and ρ_{B} is the residual state on register B. Output $\perp$ otherwise.

$\mathsf{Hyb}_0 \approx_s \mathsf{Hyb}_1$ follows from the statistical CRS indistinguishability property corresponding to $\mathsf{Sim}_{\mathsf{nizk}}$, which ensures that $(x, \mathsf{crs}_{\mathsf{in}}) \approx_s (x, \mathsf{crs}_{\mathsf{Sim}})$.

Consider now the hybrids, $\mathsf{Hyb}^1_2, \ldots, \mathsf{Hyb}^\ell_2$ where Hyb^i_2 corresponds to changing the i-th bit of 0^ℓ inside ct^1 into the i-th bit of π_{Sim}. Note that this also changes the statements x^{z} for the superposition terms z corresponding to the R register of $|\psi\rangle_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}}$. Consequently, the hybrids also involve a change to the proofs $\pi^{\mathsf{z}}_{\mathsf{out}}$ in superposition on register P accordingly. Now, let $\mathsf{Hyb}^0_2 := \mathsf{Hyb}_1$ and observe that Hyb^ℓ_2 is the following hybrid:

$\underline{\mathsf{Hyb}^\ell_2}$:

- V^* is initialized with x and $\mathsf{crs}_{\mathsf{Sim}} \| \mathsf{crs}_{\mathsf{out}}$, where $(\mathsf{crs}_{\mathsf{Sim}}, \pi_{\mathsf{Sim}}) \leftarrow \mathsf{Sim}_{\mathsf{nizk}}(x)$ and $\mathsf{crs}_{\mathsf{out}} \leftarrow \mathsf{Setup}_{\mathsf{nizk}}(1^\lambda)$, with $\mathsf{Sim}_{\mathsf{nizk}}$ being the QPT simulator of Π_{nizk}.
- Then, V^* is provided with $\sigma = (|\psi\rangle_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}}, \mathsf{ct}^0, \mathsf{ct}^1)$ computed as $(\sigma, \rho_{\mathsf{P}}) \leftarrow \mathsf{P}(\mathsf{crs}_{\mathsf{Sim}} \| \mathsf{crs}_{\mathsf{out}}, x, \omega)$, except that $\mathsf{ct}^0 = \pi_{\mathsf{in}} \oplus \mathsf{pad}^0 \oplus \mathsf{k}^0$ and $\mathsf{ct}^1 = \pi_{\mathsf{Sim}} \oplus \mathsf{pad}^1 \oplus \mathsf{k}^1$.
- Then, V^* outputs a bi-partite state $\rho_{\mathrm{A}\otimes\mathrm{B}}$. If $\mathsf{Cert}(\rho_{\mathsf{P}}, \rho_{\mathrm{A}}) = \top$, output ρ_{B} where ρ_{A} is the state on register A and ρ_{B} is the residual state on register B. Output $\perp$ otherwise.

Lemma 5.4. $\forall i \in [\ell] : \mathsf{Hyb}_2^{i-1} \approx_s \mathsf{Hyb}_2^i$

Proof. Let us now focus on the i-th BB84 state $|y^i\rangle_{\mathrm{R}^i}^{\theta^i}$, and the distribution $\mathcal{Z}^i$:

$\underline{\mathcal{Z}^i(\theta^i, b', \mathrm{D}) :}$

- If $\theta^i \neq 0^\lambda$, then execute the following:
  1. Compute $\mathsf{ct}^0, \mathsf{ct}^1$ in the same way as in Hyb_2^{i-1} except b' is set as the i-th bit inside ct^1. In particular, we have:
  $$\mathsf{ct}^0 := \pi_{\mathsf{in}} \oplus \mathsf{pad}^0 \oplus \mathsf{k}^0, \quad \mathsf{ct}^1 := \pi_{\mathsf{Sim}}[1] \| \ldots \| \pi_{\mathsf{Sim}}[i-1] \| b' \| 0^{\ell - i} \oplus \mathsf{pad}^1 \oplus \mathsf{k}^1$$
  2. Compute $|\psi\rangle_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}}$ in the same way as in Hyb_2^{i-1}, except using θ^i from the input (instead of sampling it) and register D instead of sampling the state $|y^i\rangle^{\theta^i}$.
  3. Initialize V^* with input $x, \mathsf{crs}_{\mathsf{Sim}} \| \mathsf{crs}_{\mathsf{out}}$ and provide it with the proof $\sigma := (|\psi\rangle_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}}, \mathsf{ct}^0, \mathsf{ct}^1)$.
  4. When V^* outputs a bi-partite state $\rho_{\mathrm{A}\otimes\mathrm{B}}$, parse it is as a state $\rho_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}\otimes\mathrm{B}}$ on registers $\mathrm{R}, \mathrm{P}, \mathrm{S}$ and B.
  5. Perform the map $|\mathsf{z}\rangle_{\mathrm{R}} |\widetilde{\mathsf{sig}}^{\mathsf{z}}\rangle_{\mathrm{S}} |0\rangle_{\mathrm{T}} \mapsto |\mathsf{z}\rangle_{\mathrm{R}} |\widetilde{\mathsf{sig}}^{\mathsf{z}}\rangle_{\mathrm{S}} |0 \oplus \mathsf{Test}(\mathsf{z}, \widetilde{\mathsf{sig}}^{\mathsf{z}})\rangle_{\mathrm{T}}$ where Test is as defined in the construction. Then, measure register T and output $\perp$ if the result is not 1.
  6. Perform $|\mathsf{z}\rangle_{\mathrm{R}} |\mathsf{t}\rangle_{\mathrm{P}} |\mathsf{s}\rangle_{\mathrm{S}} \mapsto |\mathsf{z}\rangle_{\mathrm{R}} |\mathsf{t} \oplus \mathsf{P}_{\mathsf{nizk}}(\mathsf{crs}_{\mathsf{out}}, x^{\mathsf{z}}, \omega_{\mathsf{out}}; r^{\mathsf{z}})\rangle_{\mathrm{P}} |\mathsf{s} \oplus \mathsf{sig}^{\mathsf{z}}\rangle_{\mathrm{S}}$ on registers R, P and S where $r^{\mathsf{z}} = \mathsf{F}_{\mathsf{prfk}}(\mathsf{z})$ to get the state $\rho'_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}\otimes\mathrm{B}}$.
  7. Output the state $\rho'_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}\otimes\mathrm{B}}$ along with $\{\theta^j, y^j\}_{j\in[2\ell]\setminus\{i\}}$.
- If $\theta^i = 0^\lambda$, then execute the following:
  1. Compute the state $\sum_{\mathsf{z}} \alpha^{\mathsf{z}} |\mathsf{z}\rangle_{\mathrm{R}}$ in the same way as in Hyb_2^{i-1}, except that register D is used instead of the state $|y^i\rangle^{\theta^i}$.
  2. Compute $\mathsf{ct}^0, \mathsf{ct}^1$ as random strings, i.e.:
  $$\mathsf{ct}^0 \leftarrow \{0,1\}^\ell, \quad \mathsf{ct}^1 \leftarrow \{0,1\}^\ell$$
  3. Then, compute the state $|\psi\rangle_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}}$ as:
  $$|\psi\rangle_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}} := \sum_{\mathsf{z}} \alpha^{\mathsf{z}} |\mathsf{z}\rangle_{\mathrm{R}} |\pi_{\mathsf{S}}^{\mathsf{z}}\rangle_{\mathrm{P}} |\mathsf{sig}^{\mathsf{z}}\rangle_{\mathrm{S}}$$
  where the NIZK proof $\pi_{\mathsf{S}}^{\mathsf{z}}$ is computed as $\pi_{\mathsf{S}}^{\mathsf{z}} \leftarrow \mathsf{S}_2(\mathsf{td}, x^{\mathsf{z}}; r^{\mathsf{z}})$ where $r^{\mathsf{z}} := \mathsf{F}_{\mathsf{prfk}}(\mathsf{z})$ for $\mathsf{prfk} \leftarrow \{0,1\}^\lambda$ and $(\mathsf{crs}_{\mathsf{S}}, \mathsf{td}) \leftarrow \mathsf{S}_1(1^\lambda)$ where $\mathsf{S} := (\mathsf{S}_1, \mathsf{S}_2)$ is the superposition-secure adaptive zero-knowledge simulator of Π_{nizk}.
  4. Initialize V^* with input $x, \mathsf{crs}_{\mathsf{Sim}} \| \mathsf{crs}_{\mathsf{S}}$ and provide it with the proof $\sigma := (|\psi\rangle_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}}, \mathsf{ct}^0, \mathsf{ct}^1)$.
  5. When V^* outputs a bi-partite state $\rho_{\mathrm{A}\otimes\mathrm{B}}$, parse it as a state $\rho_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}\otimes\mathrm{B}}$.
  6. Perform the map $|\mathsf{z}\rangle_{\mathrm{R}} |\widetilde{\mathsf{sig}}^{\mathsf{z}}\rangle_{\mathrm{S}} |0\rangle_{\mathrm{T}} \mapsto |\mathsf{z}\rangle_{\mathrm{R}} |\widetilde{\mathsf{sig}}^{\mathsf{z}}\rangle_{\mathrm{S}} |0 \oplus \mathsf{Test}(\mathsf{z}, \widetilde{\mathsf{sig}}^{\mathsf{z}})\rangle_{\mathrm{T}}$ where Test is as defined in the construction. Then, measure register T and output $\perp$ if the result is not 1.

7. Perform the map $|\mathsf{z}\rangle_{\mathrm{R}}|\mathsf{t}\rangle_{\mathrm{P}}|\mathsf{s}\rangle_{\mathrm{S}} \mapsto |\mathsf{z}\rangle_{\mathrm{R}}|\mathsf{t} \oplus \mathsf{S}_2(\mathsf{td}, x^{\mathsf{z}}; r^{\mathsf{z}}))\rangle_{\mathrm{P}}|\mathsf{s} \oplus \mathsf{sig}^{\mathsf{z}}\rangle_{\mathrm{S}}$ on registers R, P and S where $r^{\mathsf{z}} = \mathsf{F}_{\mathsf{prfk}}(\mathsf{z})$ to get the state $\rho'_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}\otimes\mathrm{B}}$.
8. Output the state $\rho'_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}\otimes\mathrm{B}}$ along with $\{\theta^j, y^j\}_{j\in[2\ell]\setminus\{i\}}$.

Lemma 5.5. *For all $\theta^i \in \{0,1\}^\lambda, b' \in \{0,1\}$ and states $|\tau\rangle_{\mathrm{C}\otimes\mathrm{D}}$ on registers C, D where D is of size λ, the following holds:*

$$\left(\mathrm{C}, \mathcal{Z}^i(\theta^i, b', \mathrm{D})\right) \approx_c \left(\mathrm{C}, \mathcal{Z}^i(0^\lambda, b', \mathrm{D})\right)$$

Proof. Consider the following hybrids that act on inputs $(\theta^i, b', \mathrm{D})$:

- $\underline{\mathcal{Z}^i_0}$: Same as $\mathcal{Z}^i$.
- $\underline{\mathcal{Z}^i_1}$: Similar to $\mathcal{Z}^i_0$, except that after the signature check in Step 5., another check is performed before Step 6., where the terms z in superposition are post-selected to be valid statements wrt the witness $\omega_{\mathsf{out}} = (\{\theta^i\}_{i\in[2\ell]}, \mathsf{k}^0, \mathsf{k}^1)$. In other words, the map $|\mathsf{z}\rangle_{\mathrm{R}}|0\rangle_{\mathrm{W}} \mapsto |\mathsf{z}\rangle_{\mathrm{R}}|\mathsf{Check}(\mathsf{z})\rangle_{\mathrm{W}}$ is first performed, where $\mathsf{Check}(\mathsf{z})$ is a function that outputs 1 iff $\omega_{\mathsf{out}} \in \mathsf{R}_{\mathsf{out}}(x^{\mathsf{z}})$, where $\mathsf{R}_{\mathsf{out}}$ is the NP relation specified by the description of x^{z}. Then, the W register is measured and $\perp$ is output if the result is not 1.
- $\underline{\mathcal{Z}^i_2}$: Similar to $\mathcal{Z}^i_1$, except that it generates simulated proofs $\pi^{\mathsf{z}}_{\mathsf{S}}$ on register P as in the distribution $\mathcal{Z}(0^\lambda, b', \mathrm{D})$, instead of the real proofs $\pi^{\mathsf{z}}_{\mathsf{out}}$, apart from replacing $\mathsf{crs}_{\mathsf{out}}$ with a simulated reference string $\mathsf{crs}_{\mathsf{S}}$.
- $\underline{\mathcal{Z}^i_3}$: Similar to $\mathcal{Z}^i_2$, except that $\mathsf{ct}^0, \mathsf{ct}^1$ are sampled as truly random values, and the simulated proofs $\pi^{\mathsf{z}}_{\mathsf{S}}$ are generated according to the new ciphertexts.
- $\underline{\mathcal{Z}^i_4}$: Similar to $\mathcal{Z}^i_3$, except that the post-selection step introduced in $\mathcal{Z}^i_1$ is no longer performed.

We begin by arguing the indistinguishability of $\mathcal{Z}^i_0$ and $\mathcal{Z}^i_1$ as follows:

Lemma 5.6. *For all $\theta^i \in \{0,1\}^\lambda$, $b' \in \{0,1\}$ and states $|\tau\rangle_{\mathrm{C}\otimes\mathrm{D}}$ with a λ-size register D, we have:*

$$(\mathrm{C}, \mathcal{Z}^i_0(\theta^i, b', \mathrm{D})) \approx_s (\mathrm{C}, \mathcal{Z}^i_1(\theta^i, b', \mathrm{D}))$$

Proof. Recall first how the two distributions differ:

- $\underline{\mathcal{Z}^i_0}$: Same as $\mathcal{Z}^i$.
- $\underline{\mathcal{Z}^i_1}$: Similar to $\mathcal{Z}^i_0$, except that after the signature check in Step 5., another check is performed before Step 6., where the terms z in superposition are post-selected to be valid statements wrt the witness $\omega_{\mathsf{out}} = (\{\theta^i\}_{i\in[2\ell]}, \mathsf{k}^0, \mathsf{k}^1)$. In other words, the map $|\mathsf{z}\rangle_{\mathrm{R}}|0\rangle_{\mathrm{W}} \mapsto |\mathsf{z}\rangle_{\mathrm{R}}|\mathsf{Check}(\mathsf{z})\rangle_{\mathrm{W}}$ is first performed, where $\mathsf{Check}(\mathsf{z})$ is a function that outputs 1 iff $\omega_{\mathsf{out}} \in \mathsf{R}_{\mathsf{out}}(x^{\mathsf{z}})$, where $\mathsf{R}_{\mathsf{out}}$ is the NP relation specified by the description of x^{z}. Then, the W register is measured and $\perp$ is output if the result is not 1.

Consider now the case when $\theta^i \neq 0^\lambda$ and assume that there exists an unbounded distinguisher D that has non-negligible advantage in distinguishing between the distributions. Consider the following functions G, H:

$\underline{G[\mathsf{aux}](\mathsf{z}, \widetilde{\mathsf{sig}^{\mathsf{z}}}) :}$

- Output $(1, \{s^{i,0}, s^{i,1}\}_{i\in[2\ell\lambda]})$ if $\mathsf{Test}(\mathsf{z}, \widetilde{\mathsf{sig}}^{\mathsf{z}}) = 1$ AND $\omega_{\mathsf{out}} \in \mathsf{R}_{\mathsf{out}}(x^{\mathsf{z}})$, where $\mathsf{Test}, x^{\mathsf{z}}, \omega_{\mathsf{out}}$, $\{s^{i,0}, s^{i,1}\}_{i\in[2\ell\lambda]}$ and $\mathsf{R}_{\mathsf{out}}$ are all as defined in Section 5, and are parsed from aux that is part of the description of G.
- Otherwise, output $(0, \{s^{i,0}, s^{i,1}\}_{i\in[2\ell\lambda]})$.

$\underline{H[\mathsf{aux}](\mathsf{z}, \widetilde{\mathsf{sig}}^{\mathsf{z}}) :}$

- Output $(1, \{s^{i,0}, s^{i,1}\}_{i\in[2\ell\lambda]})$ if $\mathsf{Test}(\mathsf{z}, \widetilde{\mathsf{sig}}^{\mathsf{z}}) = 1$ where $\mathsf{Test}, x^{\mathsf{z}}, \omega_{\mathsf{out}}$, $\{s^{i,0}, s^{i,1}\}_{i\in[2\ell\lambda]}$ and $\mathsf{R}_{\mathsf{out}}$ are all as defined in Section 5, and are parsed from aux that is part of the description of H.
- Otherwise, output $(0, \{s^{i,0}, s^{i,1}\}_{i\in[2\ell\lambda]})$.

We now construct an algorithm $\mathcal{A}$ that makes use of D in order to distinguish between the oracles G and H.

$\underline{\mathcal{A}^{\mathcal{O}}(\rho_{\mathsf{aux}}) :}$

- From the input ρ_{aux}, obtain the state $\rho_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}\otimes\mathrm{B}}$, which corresponds to the state computed after Step 4. of $\mathcal{Z}^i$.
- Evaluate $\mathcal{O}$ (which is either G or H) in superposition of registers R, S and post-select on the outcome being 1. Also, obtain $\{s^{i,0}, s^{i,1}\}_{i\in[2\ell\lambda]}$ from $\mathcal{O}$ and all other values corresponding to $\mathcal{Z}^i$ from the auxiliary input state ρ_{aux}.
- Perform the uncomputation step, i.e., Step 6. of $\mathcal{Z}^i$ to get the state $\rho'_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}\otimes\mathrm{B}}$.
- Finally, run D on input $\rho'_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}\otimes\mathrm{B}}$ and $\{\theta^j, y^j\}_{j\in[2\ell]\setminus\{i\}}$ and output the bit output by D.

Observe that $\mathcal{A}^{\mathcal{O}}$ provides D with either the output of $\mathcal{Z}_0^i$ or $\mathcal{Z}_1^i$ depending on whether $\mathcal{O}$ is G or H.

Consider now the following algorithm $\mathcal{R}$ which is given a state ρ_{aux} which consists of the state $\rho_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}\otimes\mathrm{B}}$ that is computed after Step 4. of $\mathcal{Z}^i$ along with other values sampled in $\mathcal{Z}^i$. Importantly, $\mathcal{R}$ is only given $s^{j,y[j]}$ for each $j \in [(\ell+1)\lambda, 2\ell\lambda]$ such that $\theta_j = 0$. Then, $\mathcal{R}$ runs the algorithm $\mathcal{A}^H(\rho_{\mathsf{aux}})$, except it doesn't respond to the query and simply measures the oracle query made by $\mathcal{A}$ to obtain $(\widetilde{\mathsf{z}}, \widetilde{\mathsf{sig}}^{\widetilde{\mathsf{z}}})$. Observe that $\mathcal{R}$ here simply corresponds to the algorithm $\mathcal{B}^H$ guaranteed by the O2H Lemma in the special case when $q = 1$. Hence, with non-negligible probability, $\widetilde{\mathsf{z}}$ is such that $\omega_{\mathsf{out}} \notin \mathsf{R}_{\mathsf{out}}(x^{\widetilde{\mathsf{z}}})$ and yet $\widetilde{\mathsf{sig}}^{\widetilde{\mathsf{z}}}$ is a valid signature, i.e., $\widetilde{\mathsf{sig}}^{\widetilde{\mathsf{z}}} = \mathsf{sig}^{\widetilde{\mathsf{z}}}$. Clearly, this means there exists some $j \in [(\ell+1)\lambda, 2\ell\lambda]$ where $\theta_j = 0$ and $\widetilde{\mathsf{z}}[j] \neq y[j]$. Notice however, that $\mathcal{R}$ was given no information about $s^{j,1-y[j]}$. Consequently, it is easy to see by a union bound analysis that $\widetilde{\mathsf{sig}}^{\widetilde{\mathsf{z}}} = \mathsf{sig}^{\widetilde{\mathsf{z}}}$ holds with at most negligible probability, a contradiction. This means no such distinguisher D exists, and the proof of the $\theta = 0^\lambda$ case is essentially the same, finishing the proof of Lemma 5.6. □

Next, we argue that $\mathcal{Z}_1^i$ and $\mathcal{Z}_2^i$ are computationally close. Recall that by definition, the statements x^{z} corresponding to superposition terms z are all true wrt ω_{out} before computation of the proofs on register P. Moreover, the same holds before uncomputation as well, due to the post-selection step introduced in $\mathcal{Z}_1^i$. Hence, we can invoke superposition-secure adaptive zero-knowledge (Definition 3.6) of Π_{nizk} to guarantee that $\left(\mathrm{C}, \mathcal{Z}_1^i(\theta^i, b', \mathrm{D})\right) \approx_c \left(\mathrm{C}, \mathcal{Z}_2^i(\theta^i, b', \mathrm{D})\right)$, as the two only differ in that one uses real proofs, while the other uses simulated proofs. In particular, both the computation and uncomputation of the proofs in the former (likewise latter) distribution can be performed using quantum oracle access (as in Definition 3.6) to $\mathsf{P}_{\mathsf{nizk}}(\mathsf{crs}_{\mathsf{out}}, \cdot, \cdot)$ (likewise

$\mathsf{S}_2(\mathsf{td},\cdot)$). Note that by quantum-security of $\mathsf{F}_{\mathsf{prfk}}$, we can assume wlog that the oracle access to $\mathsf{P}_{\mathsf{nizk}}(\mathsf{crs}_{\mathsf{out}},\cdot,\cdot),\mathsf{S}_2(\mathsf{td},\cdot)$ in the superposition-secure adaptive zero-knowledge definition to be simulated via randomness derived from the PRF.

Observe now that $\left(\mathrm{C},\mathcal{Z}_2^i(\theta^i,b',\mathrm{D})\right)\equiv\left(\mathrm{C},\mathcal{Z}_3^i(\theta^i,b',\mathrm{D})\right)$ because in the former, the ciphertexts $\mathsf{ct}^0,\mathsf{ct}^1$ are masked by uniform and independent values $\mathsf{k}^0,\mathsf{k}^1$, while they are directly sampled from the uniform distribution in the latter. The equivalence of the distributions then follows from the fact that the simulated proofs are directly computed using the statements x^{z} (which depend on $\mathsf{ct}^0,\mathsf{ct}^1$).

Finally, $\left(\mathrm{C},\mathcal{Z}_3^i(\theta^i,b',\mathrm{D})\right)\approx_s\left(\mathrm{C},\mathcal{Z}_4^i(\theta^i,b',\mathrm{D})\right)$ follows from essentially the same argument as that of Lemma 5.6. Notice now that $\mathcal{Z}_4^i(\theta^i,b',\mathrm{D})$ is the same distribution as $\mathcal{Z}^i(0^\lambda,b',\mathrm{D})$, and recall that $\mathcal{Z}_0^i(\theta^i,b',\mathrm{D})$ is the same distribution as $\mathcal{Z}^i(\theta^i,b',\mathrm{D})$. Therefore, we have $\left(\mathrm{C},\mathcal{Z}^i(\theta^i,b',\mathrm{D})\right)\approx_c\left(\mathrm{C},\mathcal{Z}^i(0^\lambda,b',\mathrm{D})\right)$. This concludes the proof of Lemma 5.5. □

Next, assume for contradiction that there exists a malicious QPT verifier V^* and an unbounded distinguisher D that distinguishes the output of Hyb_2^{i-1} and Hyb_2^i with non-negligible probability. We will now construct a QPT adversary $\mathcal{A}$ (along with an unbounded distinguisher) that participates in experiment $\mathsf{Exp}_{\mathcal{Z}^i,\mathcal{A}}$ specified by Theorem 3.9 and breaks its certified-deletion security.

$\underline{\mathsf{Exp}_{\mathcal{Z}^i,\mathcal{A}}(b):}$

1. Sample $y^i,\theta^i\leftarrow\{0,1\}^\lambda$ and initialize $\mathcal{A}$ with 1^λ along with the following input (or the input $\perp$ if $\mathcal{Z}^i$ does output $\perp$):

$$\left(\rho'_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}\otimes\mathrm{B}},\{\theta^j,y^j\}_{j\in[2\ell]\setminus\{i\}}\right):=\mathcal{Z}^i\left(\theta^i,b\oplus\bigoplus_{j:\theta_j^i=0}y_j^i,|y^i\rangle^{\theta^i}\right)$$

2. $\mathcal{A}$ executes as follows:

   - Measure the R register of $\rho'_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}\otimes\mathrm{B}}$ in the Hadamard basis to get outcome cert and leftover state $\rho''_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}\otimes\mathrm{B}}$.
   - Parse $\mathsf{cert}=\mathsf{cert}^1\|\ldots\|\mathsf{cert}^{2\ell}$ where $\mathsf{cert}^j\in\{0,1\}^\lambda$ for each $j\in[2\ell]$.
   - If for each $j\in[2\ell]\setminus\{i\}$, it holds that $\mathsf{cert}_k^j=y_k^j$ for each $k\in[\lambda]$ such that $\theta_k^j=1$, then output cert^i along with the state ρ''_{B} on register B. Otherwise, set $\rho''_{\mathrm{B}}:=\perp$ and output cert^i along with ρ''_{B}.

3. If $\forall j\in[\lambda]$ such that $\theta_j^i=1$, $\mathsf{cert}_j^i=y_j^i$, output ρ''_B (and $\perp$ otherwise).

Observe now that the output of Hyb_2^{i-1} is identically distributed to the output of $\mathsf{Exp}_{\mathcal{Z}^i,\mathcal{A}}(0)$ (conditioned on $\theta^i\neq 0^\lambda$ which occurs with $\mathsf{negl}(\lambda)$ probability). Likewise, Hyb_2^i is statistically close to $\mathsf{Exp}_{\mathcal{Z}^i,\mathcal{A}}(1)$, assuming a bit flip occurs from Hyb_2^{i-1} to Hyb_2^i (otherwise, the hybrids are identically distributed). Hence, the unbounded distinguisher D can also distinguish between $\mathsf{Exp}_{\mathcal{Z}^i,\mathcal{A}}(0)$ and $\mathsf{Exp}_{\mathcal{Z}^i,\mathcal{A}}(1)$ with non-negligible advantage. However, Theorem 3.9 guarantees that $\mathsf{Exp}_{\mathcal{Z}^i,\mathcal{A}}(0)\approx_s\mathsf{Exp}_{\mathcal{Z}^i,\mathcal{A}}(1)$, a contradiction. Therefore, $\mathsf{Hyb}_2^{i-1}\approx_s\mathsf{Hyb}_2^i$. This concludes the proof of Lemma 5.4. □

As a consequence of the above lemma, we have that $\mathsf{Hyb}_1\approx_s\mathsf{Hyb}_2^\ell$.

Next, consider the hybrids $\mathsf{Hyb}_2^{\ell+1},\ldots,\mathsf{Hyb}_2^{2\ell}$, in which the proof π_{in} inside ct^0 will be replaced bit-by-bit with $0s$. In other words, $\mathsf{Hyb}_2^{2\ell}$ is the following:

$\underline{\mathsf{Hyb}_2^{2\ell}:}$

- V^* is initialized with $\mathsf{crs}_{\mathsf{Sim}}\|\mathsf{crs}_{\mathsf{out}}$, where $(\mathsf{crs}_{\mathsf{Sim}}, \pi_{\mathsf{Sim}}) \leftarrow \mathsf{Sim}_{\mathsf{nizk}}(x)$ and $\mathsf{crs}_{\mathsf{out}} \leftarrow \mathsf{Setup}_{\mathsf{nizk}}(1^\lambda)$, with $\mathsf{Sim}_{\mathsf{nizk}}$ being the QPT simulator of Π_{nizk}.
- Then, V^* is provided with $\sigma = (|\psi\rangle_{\mathrm{R}\otimes\mathrm{P}\otimes\mathrm{S}}, \mathsf{ct}^0, \mathsf{ct}^1)$ computed as $(\sigma, \rho_{\mathsf{P}}) \leftarrow \mathsf{P}(\mathsf{crs}_{\mathsf{Sim}}\|\mathsf{crs}_{\mathsf{out}}, x, \omega)$, except $\mathsf{ct}^0 = 0^\ell \oplus \mathsf{pad}^0 \oplus \mathsf{k}^0$ and $\mathsf{ct}^1 = \pi_{\mathsf{Sim}} \oplus \mathsf{pad}^1 \oplus \mathsf{k}^1$.
- Then, V^* outputs a bi-partite state $\rho_{\mathrm{A}\otimes\mathrm{B}}$. If $\mathsf{Cert}(\rho_{\mathsf{P}}, \rho_{\mathrm{A}}) = \top$, output ρ_{B} where ρ_{A} is the state on register A and ρ_{B} is the residual state on register B. Output $\bot$ otherwise.

Lemma 5.7. $\forall i \in [\ell] : \mathsf{Hyb}_2^{\ell+i-1} \approx_s \mathsf{Hyb}_2^{\ell+i}$

Proof. This follows from a similar argument as the proof of Lemma 5.4. This is due to the fact that the statement x^{z} remains true whether a particular bit inside ct^0 is flipped or not. This is because ct^1 already consists of π_{Sim} which is sufficient to satisfy the second clause of the OR condition of statement x^{z}, due to the computational zero-knowledge property of Π_{nizk}. Consequently, we can invoke the superposition-secure adaptive zero-knowledge guarantee for the outer proofs as in the proof of Lemma 5.4. This finishes the proof of Lemma 5.7. □

Using this lemma, we have $\mathsf{Hyb}_2^{\ell} \approx_s \mathsf{Hyb}_2^{2\ell}$, which gives us that $\mathsf{Hyb}_0 \approx_s \mathsf{Hyb}_2^{2\ell}$. Notice that in $\mathsf{Hyb}_2^{2\ell}$, the witness ω is never used and that the view of V^* can be simulated entirely with the statement x. This proves Theorem 5.2. □

6 CE-NIZK in the Shared EPR Model

6.1 The Hidden-Bits Paradigm

Definition 6.1 (Hidden Bits Generator [QRW19])**.** A hidden bits generator (HBG) is a triple of algorithms $(\mathsf{Setup}, \mathsf{GenBits}, \mathsf{Ver})$ with the following syntax:

- $\mathsf{Setup}(1^\lambda, 1^k) \to \mathsf{crs}$: The setup algorithm takes as input a security parameter and the number of hidden-bits k, and outputs a common reference string crs.
- $\mathsf{GenBits}(\mathsf{crs}) \to \left(\mathsf{com}, r, \{\pi_i\}_{i\in[k]}\right)$: The hidden bits generation algorithm outputs a commitment com, hidden bits $r \in \{0,1\}^k$ and openings $\{\pi_i\}_{i\in[k]}$.
- $\mathsf{Ver}(\mathsf{crs}, \mathsf{com}, i, r_i, \pi_i) \to \top/\bot$: The verification algorithm takes an index i, hidden bit r_i and proof π_i along with a CRS crs and commitment com. It outputs $\top$ (accept) or $\bot$ (reject).

An HBG must satisfy the following properties:

- *Correctness:* For all $i \in [k]$, the following holds:

$$\Pr\left[\mathsf{Ver}(\mathsf{crs}, \mathsf{com}, i, r_i, \pi_i) = \bot \ : \ \begin{array}{l} \mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda, 1^k) \\ (\mathsf{com}, r, \pi_{[k]}) \leftarrow \mathsf{GenBits}(\mathsf{crs}) \end{array}\right] \leq \mathsf{negl}(\lambda)$$

- *Succinct Commitment:* There exists a set of commitments $\mathcal{COM}(\lambda)$ such that for all $\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda, 1^k)$ and every com output by $\mathsf{GenBits}(\mathsf{crs})$, $\mathsf{com} \in \mathcal{COM}(\lambda)$. Moreover, there exists a constant $\delta < 1$ such that $|\mathcal{COM}(\lambda)| \leq 2^{k^\delta \mathsf{poly}(\lambda)}$. Furthermore, $\forall \mathsf{com} \notin \mathcal{COM}(\lambda)$, $\mathsf{Ver}(\mathsf{crs}, \mathsf{com}, \cdot, \cdot, \cdot)$ outputs $\bot$.

- *Statistical Binding:* There exists a (possibly inefficient) deterministic algorithm $\mathsf{Open}(1^k, \mathsf{crs}, \mathsf{com})$ such that for all $k = \mathsf{poly}(\lambda), \mathsf{crs}$ and com, it outputs r such that for every every unbounded $\widetilde{\mathsf{P}}$, the following holds:

$$\Pr\left[\begin{array}{ll} r_i^* \neq r_i & \mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda, 1^k) \\ \wedge & : \quad (\mathsf{com}, i, r_i^*, \pi_i) \leftarrow \widetilde{\mathsf{P}}(\mathsf{crs}) \\ \mathsf{Ver}(\mathsf{crs}, \mathsf{com}, i, r_i^*, \pi_i) = \top & r \leftarrow \mathsf{Open}(1^k, \mathsf{crs}, \mathsf{com}) \end{array}\right] \leq \mathsf{negl}(\lambda)$$

- *Computational Hiding:* For all $k = \mathsf{poly}(\lambda)$ and $I \subseteq [k]$, the following holds:

$$\left[\begin{array}{lll} (\mathsf{crs}, \mathsf{com}, & I, r_I, \pi_I, r_{\overline{I}}) & \mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda, 1^k) \\ & \approx_c & : \quad (\mathsf{com}, r, \pi_{[k]}) \leftarrow \mathsf{GenBits}(\mathsf{crs}) \\ (\mathsf{crs}, \mathsf{com}, & I, r_I, \pi_I, r'_{\overline{I}}) & r' \leftarrow \{0,1\}^k \end{array}\right]$$

Theorem 6.2 ([WWW25])**.** *Assuming the polynomial hardness of LWE, there exists an HBG with statistical binding and computational hiding in the CRS model.*

Definition 6.3 (NIZK in the Hidden Bits Model [FLS90, FLS99])**.** A NIZK proof in the hidden bits model (HBM) for a language $\mathcal{L} \in \mathsf{NP}$ is a pair of algorithms (P, V) with the following syntax, where $k(\lambda, n) = \mathsf{poly}(\lambda, n)$:

- $\mathsf{P}(r, x, \omega) \to (I, \pi)$: The prover takes a hidden bit string $r \in \{0,1\}^{k(\lambda,n)}$, a statement x with size $|x| = n$, and a witness ω as input. It outputs a set of indices $I \subseteq [k]$ and a proof π.
- $\mathsf{V}(I, r_I, x, \pi) \to \top/\bot$: The verifier on input a set of indices I, hidden bits $r_I = \{r_i\}_{i \in I}$, statement x and proof π, outputs $\top/\bot$.

Such a NIZK must satisfy the following properties:

- *Completeness:* For every $x \in \mathcal{L}$ of size $|x| = n$ and $\omega \in R_{\mathcal{L}}(x)$, we have:

$$\Pr\left[\mathsf{V}(I, r_I, x, \pi) \to 0 \ : \ \begin{array}{l} r \leftarrow \{0,1\}^{k(\lambda,n)} \\ (I, \pi) \leftarrow \mathsf{P}(r, x, \omega) \end{array}\right] \leq \mathsf{negl}(\lambda)$$

- *Soundness:* For all $n = \mathsf{poly}(\lambda)$ and every unbounded P^*, the following holds:

$$\Pr\left[\mathsf{V}(I, r_I, x, \pi) \to 1 \wedge |x| = n \wedge x \notin \mathcal{L} \ : \ \begin{array}{l} r \leftarrow \{0,1\}^{k(\lambda,n)} \\ (x, \pi, I) \leftarrow \mathsf{P}^*(r) \end{array}\right] \leq \mathsf{negl}(\lambda)$$

- *Zero-Knowledge:* There exists a simulator Sim such that for every $x \in \mathcal{L}$ and $\omega \in R_{\mathcal{L}}(x)$, the following holds:

$$\left[(x, I, r_I, \pi) \equiv (x, I', r'_I, \pi') \ : \ \begin{array}{l} r \leftarrow \{0,1\}^{k(\lambda,n)} \\ (I, \pi) \leftarrow \mathsf{P}(r, x, \omega) \\ (I', r'_I, \pi') \leftarrow \mathsf{Sim}(x) \end{array}\right]$$

Theorem 6.4 ([FLS90, FLS99])**.** *NIZKs exist in the HBM for every language $\mathcal{L} \in \mathsf{NP}$.*

Remark 6.5 (Amplifying Soundess [QRW19])**.** *For any polynomial $q(\lambda, n)$, given a NIZK in the HBM (requiring k' many hidden bits) with $\mathsf{negl}(\lambda)$ soundness error, we can build one with $2^{-q(\lambda,n)} \cdot \mathsf{negl}(\lambda)$ soundness error. This is obtained by a q-fold parallel repetition of the base NIZK, requiring $q \cdot k'$ many hidden bits.*

6.2 Construction from any Hidden Bits Generator

Certified-Everlasting NIZK in the Shared EPR Model

Building Blocks:

1. Hidden Bits Generator $\mathsf{HBG} \coloneqq (\mathsf{Setup}_{\mathsf{bg}}, \mathsf{GenBits}_{\mathsf{bg}}, \mathsf{Ver}_{\mathsf{bg}})$.
2. NIZK in the Hidden Bits Model $\Pi_{\mathsf{hb}} \coloneqq (\mathsf{P}_{\mathsf{hb}}, \mathsf{V}_{\mathsf{hb}})$.
3. Let $k \coloneqq \lambda$ and ℓ be a sufficiently large polynomial in λ (the exact size is determined in the proof of statistical soundness, based on the number of hidden bits k' required by Π_{hb}). For each $(i,j) \in [\ell] \times [k]$, consider EPR pairs of the form $\mathrm{P}^{\mathrm{i}}_{\mathrm{j}}, \mathrm{V}^{\mathrm{i}}_{\mathrm{j}}$. For each $i \in [\ell]$, set $\mathrm{P}^{\mathrm{i}} \coloneqq \mathrm{P}^{\mathrm{i}}_{1} \otimes \ldots \otimes \mathrm{P}^{\mathrm{i}}_{\mathrm{k}}$ and likewise for V^{i}. Finally, set $\mathrm{P} \coloneqq \mathrm{P}^1 \otimes \ldots \otimes \mathrm{P}^\ell$ and likewise for V.

$\underline{\mathsf{Setup}(1^\lambda) :}$

- Sample $\mathsf{crs}_{\mathsf{bg}} \leftarrow \mathsf{Setup}_{\mathsf{bg}}(1^\lambda, 1^{k\ell})$ and $s \leftarrow \{0,1\}^\ell$.
- Output $\mathsf{crs} \coloneqq (\mathsf{crs}_{\mathsf{bg}}, s)$.

$\underline{\mathsf{P}(\mathsf{crs}, \mathrm{P}, x, \omega) :}$

1. Parse $\mathsf{crs} = (\mathsf{crs}_{\mathsf{bg}}, s)$ and $\mathrm{P} = \mathrm{P}^1 \otimes \ldots \otimes \mathrm{P}^\ell$.
2. Compute $(\mathsf{com}, \theta, \{\mathsf{op}^{(i,j)}\}_{(i,j)\in[\ell]\times[k]}) \leftarrow \mathsf{GenBits}_{\mathsf{bg}}(\mathsf{crs}_{\mathsf{bg}})$, where $\theta = \{\theta^i\}_{i\in[\ell]}$ and $\theta^i \in \{0,1\}^k$ for each $i \in [\ell]$, while $\mathsf{op}^{(i,j)}$ denotes the opening for bit θ^i_j.
3. For each $i \in [\ell]$, execute the following:
   (a) Measure P^{i} in the θ^i basis to get outcome y^i.
   (b) Compute $t_i \coloneqq \bigoplus_{j:\theta^i[j]=0} y^i_j$.
4. Compute $r_i \coloneqq t_i \oplus s_i$ for each $i \in [\ell]$.
5. Execute $(I \subseteq [\ell], \pi_{\mathsf{hb}}) \leftarrow \mathsf{P}_{\mathsf{hb}}(r, x, \omega)$.
6. Output $\pi \coloneqq (I, \pi_{\mathsf{hb}}, \mathsf{com}, \theta^I = \{\theta^i\}_{i\in I}, \{\mathsf{op}^{(i,j)}\}_{(i,j)\in I\times[k]})$ and $\rho_{\mathsf{P}} \coloneqq (y \coloneqq y^1 \| \ldots \| y^\ell, \theta, I, \mathrm{P})$.

$\underline{\mathsf{V}(\mathsf{crs}, \mathrm{V}, x, \pi) :}$

1. Parse $\pi = (I, \pi_{\mathsf{hb}}, \mathsf{com}, \theta^I = \{\theta^i\}_{i\in I}, \{\mathsf{op}^{(i,j)}\}_{i\in I, j\in[k]})$, $\mathsf{crs} = (\mathsf{crs}_{\mathsf{bg}}, s)$ and $\mathrm{V} = \mathrm{V}^1 \otimes \ldots \otimes \mathrm{V}^\ell$.
2. For each $i \in I$, do the following:
   - For each $j \in [k]$, if $\mathsf{Ver}_{\mathsf{bg}}(\mathsf{crs}_{\mathsf{bg}}, \mathsf{com}, (i,j), \theta^i_j, \mathsf{op}^{(i,j)}) \neq \top$, output $\bot$.
   - Measure V^{i} in the θ^i basis to get y^i.

- Compute $t_i \coloneqq \bigoplus_{j:\theta^i[j]=0} y^i_j$.

3. Compute $r_I \coloneqq t_I \oplus s_I$.
4. Output the bit $b \coloneqq \mathsf{V}_{\mathsf{hb}}(I, r_I, x, \pi_{\mathsf{hb}})$ along with $\rho_{\mathsf{V}} \coloneqq (I, \mathrm{V})$.

$\underline{\mathsf{Del}(\rho_{\mathsf{V}})}$:

1. Parse $\rho_{\mathsf{V}} = (I, \mathrm{V})$ and $\mathrm{V} = \mathrm{V}^1 \otimes \ldots \otimes \mathrm{V}^\ell$.
2. For each $i \in [\ell] \setminus I$, measure V^i in the Hadamard basis to get cert^i.
3. Output $\mathsf{cert} \coloneqq \{\mathsf{cert}^i\}_{i\in[\ell]\setminus I}$.

$\underline{\mathsf{Cert}(\mathsf{cert}, \rho_{\mathsf{P}})}$:

1. Parse $\mathsf{cert} = \{\mathsf{cert}^i\}_{i\in[\ell]\setminus I}$, $\rho_{\mathsf{P}} = (y^1 \| \ldots \| y^\ell, \theta, I, \mathrm{P})$.
2. If there exists $i \in [\ell] \setminus I$ and $j \in [k]$ such that $\theta^i_j = 1$ and $\mathsf{cert}^i_j \neq y^i_j$, output $\bot$. Otherwise, output $\top$.

Theorem 6.6. *Let* HBG *be any statistical-binding hidden bits generator and* Π_{hb} *be a NIZK for* NP *in the hidden bits model. Then,* $\Pi = (\mathsf{Setup}, \mathsf{P}, \mathsf{V}, \mathsf{Del}, \mathsf{Cert})$ *is a CE-NIZK for* NP *in the shared EPR model.*

Hence, from Theorem 6.2 and Theorem 6.4, we have the following:

Corollary 6.7. *Assuming the polynomial hardness of LWE, there exists a CE-NIZK for* NP *in the shared EPR model.*

Proof. We show that Π satisfies the properties of CE-NIZK as follows:

6.2.1 Completeness and Deletion Correctness

Completeness. The hidden-bits NIZK proof π_{hb} is computed as $(I, \pi_{\mathsf{hb}}) \leftarrow \mathsf{P}_{\mathsf{hb}}(r, x, \omega)$ where $r = t \oplus s$. Clearly, the distribution of r is uniform. By correctness of the HBG HBG, all of the $\mathsf{Ver}_{\mathsf{bg}}(\mathsf{crs}_{\mathsf{bg}}, \mathsf{com}, (i,j), \theta^i_j, \mathsf{op}^{(i,j)})$ checks will pass. Moreover, the verifier V computes the correct r_I values as s_I is available as part of crs and the t_I values are common due to the EPR pairs being measured in the same basis on both ends for the indices corresponding to I. Finally, $\mathsf{V}_{\mathsf{hb}}(I, r_I, x, \pi_{\mathsf{hb}})$ outputs 1 by the completeness of the hidden-bits NIZK Π_{hb}. Consequently, V outputs 1.

Deletion Correctness. This follows directly from the fact that for each $i \in [\ell] \setminus I$ and each $j \in [k]$ such that $\theta^i_j = 1$, both parties measure the EPR halves in the Hadamard basis to get $\mathsf{cert}^i_j = y^i_j$.

6.2.2 Computational Zero-Knowledge

For every $x \in L$ and $\omega \in R_L(x)$, consider the following hybrids.

$\underline{\mathsf{Hyb}^{\mathsf{zk}}_0}$: This is the real experiment, which proceeds as follows:

1. $\mathsf{crs}_{\mathsf{bg}} \leftarrow \mathsf{Setup}_{\mathsf{bg}}(1^\lambda, 1^{k\ell})$ and $s \leftarrow \{0,1\}^\ell$.
2. $(\mathsf{com}, \theta, \{\mathsf{op}^{(i,j)}\}_{(i,j)\in[\ell]\times[k]}) \leftarrow \mathsf{GenBits}_{\mathsf{bg}}(\mathsf{crs}_{\mathsf{bg}})$.
3. For each $i \in [\ell]$, measure P^{i} in basis θ^i to obtain y^i, and compute
$$t_i := \bigoplus_{j:\theta^i_j=0} y^i_j.$$
Let $t = t_1\| \cdots \|t_\ell$.
4. Compute $r := t \oplus s$ and $(I, \pi_{\mathsf{hb}}) \leftarrow \mathsf{P}_{\mathsf{hb}}(r, x, \omega)$.
5. Output
$$\mathsf{crs} := (\mathsf{crs}_{\mathsf{bg}}, s), \qquad \pi := (I, \pi_{\mathsf{hb}}, \mathsf{com}, \{\theta^i\}_{i\in I}, \{\mathsf{op}^{(i,j)}\}_{(i,j)\in I\times[k]})$$
and the verifier register V.

$\underline{\mathsf{Hyb}^{\mathsf{zk}}_1}$: Here, r and the opening set I are determined before the HBG basis string is generated.

1. $\mathsf{crs}_{\mathsf{bg}} \leftarrow \mathsf{Setup}_{\mathsf{bg}}(1^\lambda, 1^{k\ell})$ and $r \leftarrow \{0,1\}^\ell$.
2. Compute $(I, \pi_{\mathsf{hb}}) \leftarrow \mathsf{P}_{\mathsf{hb}}(r, x, \omega)$.
3. $(\mathsf{com}, \theta, \{\mathsf{op}^{(i,j)}\}_{(i,j)\in[\ell]\times[k]}) \leftarrow \mathsf{GenBits}_{\mathsf{bg}}(\mathsf{crs}_{\mathsf{bg}})$.
4. For each $i \in [\ell]$, measure P^{i} in basis θ^i to obtain y^i and compute $t_i := \bigoplus_{j:\theta^i_j=0} y^i_j$. Let $t = t_1\| \cdots \|t_\ell$.
5. Compute $s := t \oplus r$ and output $(\mathsf{crs}, \pi, \mathrm{V})$ as in $\mathsf{Hyb}^{\mathsf{zk}}_0$.

$\underline{\mathsf{Hyb}^{\mathsf{zk}}_2}$: This is the same as $\mathsf{Hyb}^{\mathsf{zk}}_1$, except that the unopened EPR blocks are measured in fresh random bases.

1. $\mathsf{crs}_{\mathsf{bg}} \leftarrow \mathsf{Setup}_{\mathsf{bg}}(1^\lambda, 1^{k\ell})$ and $r \leftarrow \{0,1\}^\ell$.
2. Compute $(I, \pi_{\mathsf{hb}}) \leftarrow \mathsf{P}_{\mathsf{hb}}(r, x, \omega)$.
3. $(\mathsf{com}, \theta, \{\mathsf{op}^{(i,j)}\}_{(i,j)\in[\ell]\times[k]}) \leftarrow \mathsf{GenBits}_{\mathsf{bg}}(\mathsf{crs}_{\mathsf{bg}})$.
4. For every $i \in [\ell] \setminus I$, sample $\widetilde{\theta}^i \leftarrow \{0,1\}^k$ independently. For every $i \in [\ell]$, measure P^{i} in basis θ^i if $i \in I$, and in basis $\widetilde{\theta}^i$ if $i \notin I$. Denote the outcome by y^i and set
$$t_i := \begin{cases} \bigoplus_{j:\theta^i_j=0} y^i_j, & i \in I, \\ \bigoplus_{j:\widetilde{\theta}^i_j=0} y^i_j, & i \notin I. \end{cases}$$
Let $t = t_1\| \cdots \|t_\ell$.
5. Compute $s := t \oplus r$ and output $(\mathsf{crs}, \pi, \mathrm{V})$ as in $\mathsf{Hyb}^{\mathsf{zk}}_0$.

$\underline{\mathsf{Hyb}_3^{\mathsf{zk}}}$: This is the same as $\mathsf{Hyb}_2^{\mathsf{zk}}$, except that the CRS bits corresponding to unopened indices are sampled uniformly. Namely, after computing t, set

$$s_i := \begin{cases} t_i \oplus r_i, & i \in I, \\ u_i, & i \notin I, \end{cases} \qquad u_i \leftarrow \{0,1\},$$

where the values $\{u_i\}_{i \notin I}$ are mutually independent. The output is $(\mathsf{crs}, \pi, \mathrm{V})$ as before.

$\underline{\mathsf{Hyb}_4^{\mathsf{zk}}}$: Finally, the tuple $(I, r_I, \pi_{\mathsf{hb}})$ is generated using the hidden-bits simulator.

1. $\mathsf{crs}_{\mathsf{bg}} \leftarrow \mathsf{Setup}_{\mathsf{bg}}(1^\lambda, 1^{k\ell})$ and $(I, r_I, \pi_{\mathsf{hb}}) \leftarrow \mathsf{Sim}_{\mathsf{hb}}(x)$.
2. $(\mathsf{com}, \theta, \{\mathsf{op}^{(i,j)}\}_{(i,j)\in[\ell]\times[k]}) \leftarrow \mathsf{GenBits}_{\mathsf{bg}}(\mathsf{crs}_{\mathsf{bg}})$.
3. For every $i \in I$, prepare k EPR pairs, measure the prover halves $\mathrm{P}^{\mathrm{i}} := \mathrm{P}^{\mathrm{i}}_1 \otimes \ldots \otimes \mathrm{P}^{\mathrm{i}}_{\mathrm{k}}$ in basis θ^i to obtain y^i, compute $t_i := \bigoplus_{j:\theta^i_j=0} y^i_j$, and set $s_i := t_i \oplus r_i$.
4. For every $i \notin I$, sample $s_i \leftarrow \{0,1\}$ and prepare the register V^{i} in the maximally mixed state $\mathbb{I}/2^k$.
5. Output
$$\mathsf{crs} := (\mathsf{crs}_{\mathsf{bg}}, s), \qquad \pi := (I, \pi_{\mathsf{hb}}, \mathsf{com}, \{\theta^i\}_{i\in I}, \{\mathsf{op}^{(i,j)}\}_{(i,j)\in I\times[k]})$$
and the verifier register $\mathrm{V} := \mathrm{V}^1 \otimes \cdots \otimes \mathrm{V}^\ell$.

We now prove the indistinguishability of the hybrids.

Lemma 6.8. $\mathsf{Hyb}_0^{\mathsf{zk}} \equiv \mathsf{Hyb}_1^{\mathsf{zk}}$.

Proof. Sampling s uniformly and setting $r := t \oplus s$ is identical to sampling r uniformly and setting $s := t \oplus r$. Hence, the claim follows. □

Lemma 6.9. $\mathsf{Hyb}_1^{\mathsf{zk}} \approx_c \mathsf{Hyb}_2^{\mathsf{zk}}$.

Proof. Observe that I is determined before the hidden-bits are generated, so applying computational hiding of the HBG with revealed set $I \times [k]$ gives us:

$$\begin{aligned} &(\mathsf{crs}_{\mathsf{bg}}, \mathsf{com}, I, \theta_I, \{\mathsf{op}^{(i,j)}\}_{(i,j)\in I\times[k]}, \{\theta^i\}_{i\in[\ell]\setminus I}) \\ &\qquad\qquad \approx_c \\ &(\mathsf{crs}_{\mathsf{bg}}, \mathsf{com}, I, \theta_I, \{\mathsf{op}^{(i,j)}\}_{(i,j)\in I\times[k]}, \{\widetilde{\theta}^i\}_{i\in[\ell]\setminus I}), \end{aligned}$$

where the values $\{\widetilde{\theta}^i\}_{i\in[\ell]\setminus I}$ are uniformly random and independent of other values. The remainder of either hybrid is efficient post-processing of this tuple. The last components of the tuples are used as measurement bases for the "unopened" EPR blocks. Then, the measurement outcomes are used to compute t, s and the resulting view. Hence, these hybrids are computationally indistinguishable. □

We first establish the following lemma, before showing the indistinguishability of $\mathsf{Hyb}_2^{\mathsf{zk}}$ and $\mathsf{Hyb}_3^{\mathsf{zk}}$.

Lemma 6.10. *Let*

$$|\Phi_k\rangle_{\mathrm{P}\otimes\mathrm{V}} := 2^{-k/2} \sum_{z\in\{0,1\}^k} |z\rangle_{\mathrm{P}}|z\rangle_{\mathrm{V}}.$$

Sample $\Theta \leftarrow \{0,1\}^k$, measure P in the basis Θ to obtain $Y \in \{0,1\}^k$, and write

$$T := \bigoplus_{j:\Theta_j=0} Y_j,$$

on a 1-qubit register T. Note that the XOR over the empty set is defined to be 0. Let $\rho_{\mathrm{T}\otimes\mathrm{V}}$ denote the resulting state on registers T,V. Then, if Θ is not included in the output, the following holds:

$$\mathsf{TD}\left(\rho_{\mathrm{T}\otimes\mathrm{V}}, \frac{\mathbb{I}_{\mathrm{T}}}{2}\otimes\frac{\mathbb{I}_{\mathrm{V}}}{2^k}\right) = 2^{-k-1}.$$

The same bound holds after adjoining arbitrary side information independent of $(\Theta,\mathrm{P},\mathrm{V})$ and after replacing T by $T\oplus R$, where R is a classical bit computed from that side information.

Proof. For $\theta\in\{0,1\}^k$, write $H^\theta := \bigotimes_{j=1}^k H^{\theta_j}$ and $p_\theta(y) := \bigoplus_{j:\theta_j=0} y_j$. Conditioned on $(\Theta,Y)=(\theta,y)$, register V is in state $H^\theta|y\rangle$. For $b\in\{0,1\}$, let

$$\sigma_b := 4^{-k} \sum_{\theta,y\in\{0,1\}^k:\ p_\theta(y)=b} H^\theta\,|y\rangle\langle y|\,H^\theta.$$

Thus,

$$\rho_{\mathrm{T}\otimes\mathrm{V}} = |0\rangle\langle 0|\otimes\sigma_0 + |1\rangle\langle 1|\otimes\sigma_1 \quad \text{and} \quad \sigma_0+\sigma_1 = \frac{\mathbb{I}_{\mathrm{V}}}{2^k}.$$

Moreover,

$$\begin{aligned}
\sigma_0-\sigma_1 &= 4^{-k}\sum_{\theta,y}(-1)^{p_\theta(y)}H^\theta\,|y\rangle\langle y|\,H^\theta \\
&= 4^{-k}\bigotimes_{j=1}^k\left(\sum_{\theta_j\in\{0,1\}}\sum_{y_j\in\{0,1\}}(-1)^{(1-\theta_j)y_j}H^{\theta_j}\,|y_j\rangle\langle y_j|\,H^{\theta_j}\right) \\
&= 4^{-k}(Z+\mathbb{I})^{\otimes k} = 2^{-k}\,|0^k\rangle\langle 0^k|.
\end{aligned}$$

The second equality follows coordinate-wise. If $\theta_j=0$, the contribution is $\sum_{y_j}(-1)^{y_j}|y_j\rangle\langle y_j| = Z$. If $\theta_j=1$, the contribution is $\sum_{y_j} H|y_j\rangle\langle y_j|H = \mathbb{I}$. Since the two diagonal blocks of $\rho_{\mathrm{T}\otimes\mathrm{V}} - (\mathbb{I}_{\mathrm{T}}/2)\otimes(\mathbb{I}_{\mathrm{V}}/2^k)$ are $(\sigma_0-\sigma_1)/2$ and its negative, respectively,

$$\mathsf{TD}\left(\rho_{\mathrm{T}\otimes\mathrm{V}}, \frac{\mathbb{I}_{\mathrm{T}}}{2}\otimes\frac{\mathbb{I}_{\mathrm{V}}}{2^k}\right) = \frac{1}{2}\|\sigma_0-\sigma_1\|_1 = 2^{-k-1}.$$

Also, tensoring an independent state and applying a controlled XOR are trace-distance preserving, while discarding registers cannot increase trace distance. This proves the final statement. □

Lemma 6.11. $\mathsf{Hyb}_2^{\mathsf{zk}} \approx_s \mathsf{Hyb}_3^{\mathsf{zk}}$.

Proof. Enumerate the indices in $[\ell]\setminus I$ and replace the corresponding bits of s one at a time. At a target index $i\notin I$, condition on all classical and quantum side information outside block i, including r_i and the hidden-bits proof. The basis $\widetilde{\theta}^i$ and the EPR block $(\mathrm{P}^{\mathrm{i}},\mathrm{V}^{\mathrm{i}})$ are independent of this side

information. Lemma 6.10, followed by the map $s_i = t_i \oplus r_i$, bounds the trace distance of this switch by 2^{-k-1}. Therefore,

$$\mathsf{TD}(\mathsf{Hyb}_2^{\mathsf{zk}}, \mathsf{Hyb}_3^{\mathsf{zk}}) \leq \ell\, 2^{-k-1} = \mathsf{negl}(\lambda)$$

since $\ell = \mathsf{poly}(\lambda)$ and $k = \lambda$. □

Lemma 6.12. $\mathsf{Hyb}_3^{\mathsf{zk}} \approx_s \mathsf{Hyb}_4^{\mathsf{zk}}$.

Proof. In $\mathsf{Hyb}_3^{\mathsf{zk}}$, the unopened values s_i are uniform and no longer depend on r_i or t_i. After the hidden random measurement basis and outcome of an unopened EPR block are discarded, the verifier's marginal state is exactly $\mathbb{I}/2^k$. Consequently, the remainder of the hybrid depends on the hidden-bits execution only via $(I, r_I, \pi_{\mathsf{hb}}) \leftarrow \mathsf{P}_{\mathsf{hb}}(r, x, \omega)$. By perfect zero-knowledge of the NIZK in the hidden-bits model, we have

$$(x, I, r_I, \pi_{\mathsf{hb}}) \equiv (x, I', r'_{I'}, \pi'_{\mathsf{hb}}), \qquad (I', r'_{I'}, \pi'_{\mathsf{hb}}) \leftarrow \mathsf{Sim}_{\mathsf{hb}}(x).$$

Hence, the indistinguishability follows. □

Observe that $\mathsf{Hyb}_4^{\mathsf{zk}}$ can be computed using the statement x alone. This finishes the proof of computational zero-knowledge.

6.2.3 Certified-Everlasting Zero-Knowledge

Fix a QPT malicious verifier V^*, $x \in L$, and $\omega \in R_L(x)$. Consider the following hybrids.

$\mathsf{Hyb}_0^{\mathsf{cezk}}$: This is the real experiment, which proceeds as follows:

1. $\mathsf{crs}_{\mathsf{bg}} \leftarrow \mathsf{Setup}_{\mathsf{bg}}(1^\lambda, 1^{k\ell})$ and $s \leftarrow \{0,1\}^\ell$.
2. $(\mathsf{com}, \theta, \{\mathsf{op}^{(i,j)}\}_{(i,j)\in[\ell]\times[k]}) \leftarrow \mathsf{GenBits}_{\mathsf{bg}}(\mathsf{crs}_{\mathsf{bg}})$.
3. For every $i \in [\ell]$, measure P^{i} in basis θ^i to obtain y^i, and compute $t_i := \bigoplus_{j:\theta^i_j=0} y^i_j$. Let $t = t_1 \| \cdots \| t_\ell$.
4. Compute $r := t \oplus s$ and execute $(I, \pi_{\mathsf{hb}}) \leftarrow \mathsf{P}_{\mathsf{hb}}(r, x, \omega)$.
5. Set
$$\mathsf{crs} := (\mathsf{crs}_{\mathsf{bg}}, s), \qquad \pi := (I, \pi_{\mathsf{hb}}, \mathsf{com}, \theta_I, \{\mathsf{op}^{(i,j)}\}_{(i,j)\in I\times[k]}).$$
6. Run $(\mathsf{cert}^*, \rho^*) \leftarrow \mathsf{V}^*(\mathsf{crs}, \mathrm{V}, x, \pi)$.
7. Let $\rho_{\mathsf{P}} := (y := y^1 \| \ldots \| y^\ell, \theta, I, \mathrm{P})$.
8. Output ρ^* if $\mathsf{Cert}(\mathsf{cert}^*, \rho_{\mathsf{P}}) = \top$ and $\bot$ otherwise.

$\mathsf{Hyb}_1^{\mathsf{cezk}}$: Here, r and the opening set I are determined before the HBG basis is generated.

1. $\mathsf{crs}_{\mathsf{bg}} \leftarrow \mathsf{Setup}_{\mathsf{bg}}(1^\lambda, 1^{k\ell})$ and $r \leftarrow \{0,1\}^\ell$.
2. Execute $(I, \pi_{\mathsf{hb}}) \leftarrow \mathsf{P}_{\mathsf{hb}}(r, x, \omega)$.
3. $(\mathsf{com}, \theta, \{\mathsf{op}^{(i,j)}\}_{(i,j)\in[\ell]\times[k]}) \leftarrow \mathsf{GenBits}_{\mathsf{bg}}(\mathsf{crs}_{\mathsf{bg}})$.
4. For every $i \in [\ell]$, measure P^{i} in basis θ^i to obtain y^i, and compute $t_i := \bigoplus_{j:\theta^i_j=0} y^i_j$. Let $t = t_1 \| \cdots \| t_\ell$.

5. Set $s := t \oplus r$, and compute (crs, π) similar to $\mathsf{Hyb}_0^{\mathsf{cezk}}$.

6. Run $(\mathsf{cert}^*, \rho^*) \leftarrow \mathsf{V}^*(\mathsf{crs}, \mathrm{V}, x, \pi)$.

7. Output ρ^* if $\mathsf{Cert}(\mathsf{cert}^*, \rho_{\mathsf{P}}) = \top$ and $\bot$ otherwise.

For each execution, write $[\ell] \setminus I = \{j_1 < \cdots < j_m\}$. For $q \in \{0, \ldots, \ell\}$, let $J_q := \{j_1, \ldots, j_{\min\{q,m\}}\}$.

$\underline{\mathsf{Hyb}_2^{\mathsf{cezk},q}}$: This is the same as $\mathsf{Hyb}_1^{\mathsf{cezk}}$, except that after t is computed, the CRS mask is set as

$$s_i := \begin{cases} U_i, & i \in J_q, \\ t_i \oplus r_i, & i \notin J_q, \end{cases}$$

where the U_i are independent uniform bits. Thus, $\mathsf{Hyb}_2^{\mathsf{cezk},0} \equiv \mathsf{Hyb}_1^{\mathsf{cezk}}$, while in $\mathsf{Hyb}_2^{\mathsf{cezk},\ell}$ every s_i with $i \notin I$ is uniform and independent.

$\underline{\mathsf{Hyb}_3^{\mathsf{cezk}}}$: Finally, the hidden-bits transcript is generated using the hidden-bits simulator.

1. $\mathsf{crs}_{\mathsf{bg}} \leftarrow \mathsf{Setup}_{\mathsf{bg}}(1^\lambda, 1^{k\ell})$ and $(I, r_I, \pi_{\mathsf{hb}}) \leftarrow \mathsf{Sim}_{\mathsf{hb}}(x)$.

2. Compute
$$(\mathsf{com}, \theta, \{\mathsf{op}^{(i,j)}\}_{(i,j)\in[\ell]\times[k]}) \leftarrow \mathsf{GenBits}_{\mathsf{bg}}(\mathsf{crs}_{\mathsf{bg}}).$$

3. For every $i \in [\ell]$, measure $\mathrm{P^i}$ in basis θ^i to obtain y^i, and compute $t_i := \bigoplus_{j:\theta^i_j=0} y^i_j$.

4. Set $s_I := t_I \oplus r_I$ and sample $s_i \leftarrow \{0,1\}$ independently for every $i \notin I$.

5. Set
$$\mathsf{crs} := (\mathsf{crs}_{\mathsf{bg}}, s), \qquad \pi := (I, \pi_{\mathsf{hb}}, \mathsf{com}, \theta_I, \{\mathsf{op}^{(i,j)}\}_{(i,j)\in I\times[k]}).$$

6. Run $(\mathsf{cert}^*, \rho^*) \leftarrow \mathsf{V}^*(\mathsf{crs}, \mathrm{V}, x, \pi)$.

7. Let $\rho_{\mathsf{P}} := (y := y^1 \| \ldots \| y^\ell, \theta, I, \mathrm{P})$.

8. Output ρ^* if $\mathsf{Cert}(\mathsf{cert}^*, \rho_{\mathsf{P}}) = \top$ and $\bot$ otherwise.

We next describe the two auxiliary experiments used to compare adjacent $\mathsf{Hyb}_2^{\mathsf{cezk},q}$ hybrids. Fix $q \in [\ell]$ with $q \le m$, and let $a := j_q$.

$\underline{\mathsf{Exp}_q(b)}$: This experiment is the same as $\mathsf{Hyb}_2^{\mathsf{cezk},q-1}$, except that the target coordinate is set to

$$s_a := t_a \oplus r_a \oplus b.$$

The experiment then runs V^* and applies certification as usual, and outputs ρ^* or $\bot$ accordingly.

$\underline{\widetilde{\mathsf{Exp}}_q(b)}$ This is the deferred measurement version of $\mathsf{Exp}_q(b)$.

1. Generate $\mathsf{crs}_{\mathsf{bg}}$, r, (I, π_{hb}), and the HBG output as in $\mathsf{Hyb}_1^{\mathsf{cezk}}$. For every $i \neq a$, measure register $\mathrm{P^i}$, as in that hybrid, but leave $\mathrm{P^a}$ unmeasured.

2. Sample $u \leftarrow \{0,1\}$ and set $s_a := u$. Set all other coordinates of s as in $\mathsf{Exp}_q(b)$, compute (crs, π), and run $(\mathsf{cert}^*, \rho^*) \leftarrow \mathsf{V}^*(\mathsf{crs}, \mathrm{V}, x, \pi)$.

3. Measure P^{a} in basis θ^a to obtain y^a, and compute $t_a := \bigoplus_{j:\theta^a_j=0} y^a_j$.

4. Output ρ^* if certification passes and $u = t_a \oplus r_a \oplus b$. Otherwise, output $\perp$.

We now prove indistinguishability of the hybrids.

Lemma 6.13. $\mathsf{Hyb}_0^{\mathsf{cezk}} \equiv \mathsf{Hyb}_1^{\mathsf{cezk}}$.

Proof. For every fixed $t \in \{0,1\}^\ell$, the map $s \mapsto r = t \oplus s$ is a bijection on $\{0,1\}^\ell$. Thus, one can first sample $r \leftarrow \{0,1\}^\ell$ and later set $s = t \oplus r$. The HBG generation and EPR measurements are independent of r and may be moved after the computation of $\mathsf{P}_{\mathsf{hb}}(r, x, \omega)$. □

Lemma 6.14. *For every $q \le m$ and $b \in \{0,1\}$,*

$$\mathsf{Hyb}_2^{\mathsf{cezk},q-1} \equiv \mathsf{Exp}_q(0), \qquad \mathsf{Hyb}_2^{\mathsf{cezk},q} \equiv \tfrac{1}{2}\mathsf{Exp}_q(0) + \tfrac{1}{2}\mathsf{Exp}_q(1),$$

and

$$\widetilde{\mathsf{Exp}}_q(b) \equiv \tfrac{1}{2}\mathsf{Exp}_q(b) + \tfrac{1}{2}\perp.$$

Proof. The first identity is immediate from the definition. In $\mathsf{Hyb}_2^{\mathsf{cezk},q}$, the target value s_a is uniform, which is equivalent to sampling $b \leftarrow \{0,1\}$ and setting $s_a = t_a \oplus r_a \oplus b$, giving the second identity. For the last identity, the measurement of P^{a} commutes with the action of V^*, as they act on different registers. Hence, this measurement can be done before running V^*. The consistency equality in $\widetilde{\mathsf{Exp}}_q(b)$ then holds with probability exactly $1/2$, and conditioned on holding, the resulting experiment is $\mathsf{Exp}_q(b)$. □

Now, fix $q \le m$ and let $a = j_q$. In the common state of $\widetilde{\mathsf{Exp}}_q(0)$ and $\widetilde{\mathsf{Exp}}_q(1)$ immediately after the output of V^*, let $d := \mathsf{cert}^{*,a}$ denote the certificate for the corresponding EPR block. For a candidate basis $\vartheta \in \{0,1\}^k$, write $\vartheta_0 := \{j : \vartheta_j = 0\}$ and $\vartheta_1 := \{j : \vartheta_j = 1\}$. Define the following projector on P^{a}:

$$\Pi_{d,\vartheta} := \left(H^{\otimes|\vartheta_1|}\,|d_{\vartheta_1}\rangle\,\langle d_{\vartheta_1}|\,H^{\otimes|\vartheta_1|}\right)^{\mathrm{P}^{\mathrm{a}}_{\vartheta_1}} \otimes \sum_{\substack{z\in\{0,1\}^{|\vartheta_0|}:\\ \Delta(z,d_{\vartheta_0})\ge 1/2}} \left(H^{\otimes|\vartheta_0|}\,|z\rangle\,\langle z|\,H^{\otimes|\vartheta_0|}\right)^{\mathrm{P}^{\mathrm{a}}_{\vartheta_0}},$$

where Δ denotes relative Hamming distance. To include the event that there are too few computational-basis positions, define

$$\widehat{\Pi}_{d,\vartheta} := \begin{cases} \mathbb{I}, & |\vartheta_0| < k/3, \\ \Pi_{d,\vartheta}, & |\vartheta_0| \ge k/3. \end{cases}$$

Lemma 6.15. *In the common pre-measurement state of $\widetilde{\mathsf{Exp}}_q(0)$ and $\widetilde{\mathsf{Exp}}_q(1)$,*

$$\Pr[\widehat{\Pi}_{d,\theta^a} \text{ accepts}] = \mathsf{negl}(\lambda)$$

Proof. First, replace the candidate basis used in the projector by an independent $\widetilde{\theta}^a \leftarrow \{0,1\}^k$. It is independent of the state held by V^*, its output d, and the unmeasured register P^{a}. The projector is diagonal in the Hadamard basis. Hence, for every fixed outcome z and fixed d, the probability over $\widetilde{\theta}^a$ that

$$z_{\widetilde{\theta}^a_1} = d_{\widetilde{\theta}^a_1} \quad \text{and} \quad \Delta(z_{\widetilde{\theta}^a_0}, d_{\widetilde{\theta}^a_0}) \ge 1/2$$

is $\mathsf{negl}(\lambda)$ by the same Hoeffding argument as Claim 3.3 in the work of Bartusek and Khurana [BK23]. Also, $\Pr[|\widetilde{\theta}^a_0| < k/3] = \mathsf{negl}(\lambda)$. Thus, the bad projector accepts with negligible probability for an independent candidate basis.

We transfer this bound to the actual HBG basis. Since $a \notin I$ and I was fixed before HBG generation, apply HBG hiding with all coordinates outside block a revealed. A reduction receives the HBG commitment, all outside basis bits and openings, and either the actual target block θ^a or an independent uniform block. It uses the outside information to generate the non-target part of the experiment, leaves P^{a} unmeasured, runs V^*, and measures $\widehat{\Pi}_{d,\vartheta}$ using the challenge block ϑ. This measurement is efficient. A non-negligible difference between the two acceptance probabilities would therefore violate computational hiding of the HBG. This argument is similar to the one in Claim 3.4 of [BK23]. □

Lemma 6.16. *For every $q \le m$, $\widetilde{\mathsf{Exp}}_q(0) \approx_s \widetilde{\mathsf{Exp}}_q(1)$.*

Proof. By the previous lemma and gentle measurement, the common global state after V^* outputs is negligibly close to a state in the image of $\mathbb{I} - \widehat{\Pi}_{d,\theta^a}$. Consider the final measurements and tests on this ideal state. If the target check fails, both experiments output $\perp$. Suppose it passes. Let $C := \theta^a_0$ and $n := |C| \ge k/3$. A purification of the remaining state has the form

$$|\psi\rangle = \sum_{z\in\{0,1\}^n:\ \mathrm{HD}(z,d_C)<n/2} H^{\otimes n}|z\rangle_{P^a_C}|\phi_z\rangle_{\mathrm{E}},$$

where HD denotes Hamming distance and E contains all side information and the vectors $|\phi_z\rangle$ may be unnormalized.

Consider applying $Z^{d_C} := \bigotimes_{j\in C} Z^{d_j}$ for the sake of analysis. Since

$$Z^{d_C} H^{\otimes n}|z\rangle = H^{\otimes n}|z \oplus d_C\rangle, \qquad \mathrm{wt}(z \oplus d_C) < n/2,$$

where wt denotes Hamming weight. The transformed register has low-weight support in the Hadamard basis. Hence, the XOR extractor of [BK23] implies that its computational-basis parity $t_a = \bigoplus_{j\in C} y^a_j$ is uniform and independent of E. The same holds for the original register: Z^{d_C} is diagonal in the computational basis, so it does not affect the computational basis parity. This proves $\widetilde{\mathsf{Exp}}_q(0) \approx_s \widetilde{\mathsf{Exp}}_q(1)$. □

Lemma 6.17. *For every $q \in [\ell]$, $\mathsf{Hyb}_2^{\mathsf{cezk},q-1} \approx_s \mathsf{Hyb}_2^{\mathsf{cezk},q}$.*

Proof. If $q > m$, the hybrids are identical. Otherwise, Lemmas 6.14 and 6.16 show that $\mathsf{Exp}_q(0) \approx_s \mathsf{Exp}_q(1)$. Since $\mathsf{Hyb}_2^{\mathsf{cezk},q-1} \equiv \mathsf{Exp}_q(0)$ and $\mathsf{Hyb}_2^{\mathsf{cezk},q}$ is the equal mixture of $\mathsf{Exp}_q(0)$ and $\mathsf{Exp}_q(1)$, the claim follows. By a hybrid argument, we get:

$$\mathsf{Hyb}_1^{\mathsf{cezk}} \equiv \mathsf{Hyb}_2^{\mathsf{cezk},0} \approx_s \mathsf{Hyb}_2^{\mathsf{cezk},\ell}.$$

□

Lemma 6.18. $\mathsf{Hyb}_2^{\mathsf{cezk},\ell} \approx_s \mathsf{Hyb}_3^{\mathsf{cezk}}$.

Proof. In $\mathsf{Hyb}_2^{\mathsf{cezk},\ell}$, every s_i for $i \notin I$ is uniform and independent. Hence, we can replace $(I, r_I, \pi_{\mathsf{hb}})$ by $(I', r'_I, \pi'_{\mathsf{hb}}) \leftarrow \mathsf{Sim}_{\mathsf{hb}}(x)$ using the perfect zero-knowledge of the hidden bits NIZK. This yields the hybrid $\mathsf{Hyb}_3^{\mathsf{cezk}}$. □

Combining the lemmas gives $\mathsf{Hyb}_0^{\mathsf{cezk}} \approx_s \mathsf{Hyb}_3^{\mathsf{cezk}}$. The last hybrid is generated by a QPT simulator using only x and black-box access to V^*. Hence, certified-everlasting zero-knowledge follows.

6.2.4 Statistical Soundness

Let P^* be an unbounded malicious prover that breaks soundness with some $\mu(\lambda)$ probability:

$$\Pr\left[\begin{array}{ll} \mathsf{V}(\mathsf{crs}, \mathrm{V}, x, \pi^*) = \top & \\ \wedge & : \quad \begin{array}{l}\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ (x, \pi^*) \leftarrow \mathsf{P}^*(\mathsf{crs}, \mathrm{P})\end{array} \\ x \notin \mathcal{L} & \end{array}\right] \geq \mu(\lambda)$$

where $\pi^* = (I, \pi_{\mathsf{hb}}, \mathsf{com}^*, \{\theta^i\}_{i\in I}, \{\mathsf{op}^{(i,j)}\}_{(i,j)\in I\times[k]})$.

Consider now a verifier $\widetilde{\mathsf{V}}_0$ that behaves similarly to V, except that it measures all the registers $\mathrm{V}^1 \dots \mathrm{V}^\ell$ in the computational basis. It is easy to see that if P^* convinces V with probability $\mu(\lambda)$, then it also convinces $\widetilde{\mathsf{V}}_0$ with probability $\mu(\lambda)$. This is because only the computational basis positions are used for verification by V.

Next, consider a verifier $\widetilde{\mathsf{V}}$ that behaves similar to $\widetilde{\mathsf{V}}_0$ except that the registers $\mathrm{V}^1, \dots, \mathrm{V}^\ell$ are measured in the computational basis before P^* is initialized. Since operations on different registers commute, P^* also convinces $\widetilde{\mathsf{V}}$ with probability $\mu(\lambda)$:

$$\Pr\left[\begin{array}{ll} \widetilde{\mathsf{V}}(\mathsf{crs}, \mathrm{V}, x, \pi^*) = \top & \\ \wedge & : \quad \begin{array}{l}\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ (x, \pi^*) \leftarrow \mathsf{P}^*(\mathsf{crs}, \mathrm{P})\end{array} \\ x \notin \mathcal{L} & \end{array}\right] \geq \mu(\lambda)$$

where $\pi^* = (I, \pi_{\mathsf{hb}}, \mathsf{com}^*, \{\theta^i\}_{i\in I}, \{\mathsf{op}^{(i,j)}\}_{(i,j)\in I\times[k]})$.

Now, let $y^1, \dots, y^\ell$ be the values measured by $\widetilde{\mathsf{V}}$ corresponding to registers $\mathrm{V}^1, \dots, \mathrm{V}^\ell$. Consequently, once the basis choices $\theta^1, \dots, \theta^\ell$ are specified, the values $t_1, \dots, t_\ell$ are fixed as they are computed as the parity of the computational basis values. Observe now that by statistical binding of HBG, except with negligible probability, we have that r_I computed by $\widetilde{\mathsf{V}}$ satisfies $r_I = t_I \oplus s_I$ where $t_i = \bigoplus_{j:\widetilde{\theta}^i_j = 0} y^i_j$ for each $i \in I$ and $\widetilde{\theta}^1 \| \dots \| \widetilde{\theta}^\ell := \mathsf{Open}(1^{k\ell}, \mathsf{crs}_{\mathsf{bg}}, \mathsf{com}^*)$. Moreover, succinctness of HBG guarantees that $\mathsf{com}^* \in \mathcal{COM}(\lambda)$. Hence, we have the following, where $\mathsf{op}^{I,[k]}$ denotes $\{\mathsf{op}^{(i,j)}\}_{(i,j)\in I\times[k]}$:

$$v(\lambda) := \Pr\left[\begin{array}{ll} & \mathsf{crs} = (\mathsf{crs}_{\mathsf{bg}}, s) \leftarrow \mathsf{Setup}(1^\lambda) \\ \mathsf{V}_{\mathsf{hb}}(I, r_I, x, \pi_{\mathsf{hb}}) = \top & (x, \pi^*) \leftarrow \mathsf{P}^*(\mathsf{crs}, \mathrm{P}) \\ \wedge\ x \notin \mathcal{L} & : \quad \widetilde{\theta}^1 \| \dots \| \widetilde{\theta}^\ell := \mathsf{Open}(1^{k\ell}, \mathsf{crs}_{\mathsf{bg}}, \mathsf{com}^*) \\ \wedge\ \mathsf{com}^* \in \mathcal{COM}(\lambda) & \forall i \in [\ell] : t_i := \bigoplus_{j:\widetilde{\theta}^i_j = 0} y^i_j \\ & r = t \oplus s \end{array}\right] \geq \mu(\lambda) - \mathsf{negl}(\lambda)$$

where $\pi^* = (I, \pi_{\mathsf{hb}}, \mathsf{com}^*, \{\theta^i\}_{i\in I}, \{\mathsf{op}^{(i,j)}\}_{(i,j)\in I\times[k]})$.

Next, we define a similar probability for any fixed $\mathsf{com} \in \mathcal{COM}(\lambda)$:

$$
\begin{aligned}
v_{\mathsf{com}}(\lambda) &\coloneqq \Pr\left[\begin{array}{l|l} \mathsf{V}_{\mathsf{hb}}(I, r_I, x, \pi_{\mathsf{hb}}) = \top & \mathsf{crs} = (\mathsf{crs}_{\mathsf{bg}}, s) \leftarrow \mathsf{Setup}(1^\lambda) \\ \wedge\ x \notin \mathcal{L} & (x, \pi^*) \leftarrow \mathsf{P}^*(\mathsf{crs}, \mathrm{P}) \\ \wedge\ \mathsf{com}^* = \mathsf{com} & \{\widetilde{\theta}^i\}_{i\in[\ell]} \coloneqq \mathsf{Open}(1^{k\ell}, \mathsf{crs}_{\mathsf{bg}}, \mathsf{com}^*) \\ & \forall i \in [\ell] : t_i \coloneqq \bigoplus_{j:\widetilde{\theta}^i_j = 0} y^i_j \\ & r = t \oplus s \end{array}\right] \\
&\leq \Pr\left[\begin{array}{l|l} \mathsf{V}_{\mathsf{hb}}(I, r_I, x, \pi_{\mathsf{hb}}) = \top & \mathsf{crs} = (\mathsf{crs}_{\mathsf{bg}}, s) \leftarrow \mathsf{Setup}(1^\lambda) \\ \wedge & \widetilde{\theta}^1 \| \ldots \| \widetilde{\theta}^\ell \coloneqq \mathsf{Open}(1^{k\ell}, \mathsf{crs}_{\mathsf{bg}}, \mathsf{com}) \\ x \notin \mathcal{L} & \forall i \in [\ell] : t_i \coloneqq \bigoplus_{j:\widetilde{\theta}^i_j = 0} y^i_j \\ & r = t \oplus s \\ & (x, I, \pi_{\mathsf{hb}}) \leftarrow \widehat{\mathsf{P}}_{\mathsf{crs}_{\mathsf{bg}}, \mathsf{com}}(r) \end{array}\right] \\
&\leq 2^{-q(\lambda,n)} \cdot \mathsf{negl}(\lambda)
\end{aligned}
$$

where $\pi^* = (I, \pi_{\mathsf{hb}}, \mathsf{com}^*, \{\theta^i\}_{i\in I}, \{\mathsf{op}^{(i,j)}\}_{(i,j)\in I\times[k]})$.

The first inequality follows because r is already fixed by com and the second follows from the soundness of the hidden bits NIZK by viewing $\widehat{\mathsf{P}}_{\mathsf{crs}_{\mathsf{bg}}, \mathsf{com}}(r)$ as a malicious hidden bits model prover and from the fact that r is uniformly random (because s is uniform).

Next, we set $q(\lambda, n)^{1-\delta} = k \cdot k' \cdot p(\lambda)$ so that:

$$
(k\ell)^\delta \cdot p(\lambda) = (k \cdot k' \cdot q(\lambda, n))^\delta \cdot p(\lambda) \leq q^\delta(\lambda, n) \cdot q^{1-\delta}(\lambda, n)
$$

This ensures that $|\mathcal{COM}(\lambda)| \cdot v_{\mathsf{com}}(\lambda) \leq \mathsf{negl}(\lambda)$. Therefore, a union bound gives us the following:

$$
\mu(\lambda) - \mathsf{negl}(\lambda) \leq v(\lambda) \leq \sum_{\mathsf{com}\in\mathcal{COM}(\lambda)} v_{\mathsf{com}}(\lambda) \leq \mathsf{negl}(\lambda)
$$

This proves that $\mu(\lambda) \leq \mathsf{negl}(\lambda)$, ensuring soundness.

□

Acknowledgements: We would like to thank Fang Song for helpful discussions and the anonymous reviewers of Crypto 2026 and TCC 2026 for their useful feedback.

AI Statement: The proofs of the EPR construction for computational ZK (Section 6.2.2) and CE-ZK (Section 6.2.3) have been rewritten using ChatGPT, fixing the errors from previous versions.

References

[ABKK23] Amit Agarwal, James Bartusek, Dakshita Khurana, and Nishant Kumar. A new framework for quantum oblivious transfer. In *EUROCRYPT 2023, Part I*, volume 14004 of *LNCS*, pages 363–394. Springer, Cham, April 2023.

[ACE+24] Behzad Abdolmaleki, Céline Chevalier, Ehsan Ebrahimi, Giulio Malavolta, and Quoc-Huy Vu. On quantum simulation-soundness. *CiC*, 1(4):18, 2024.

[AGL24] Prabhanjan Ananth, Aditya Gulati, and Yao-Ting Lin. Cryptography in the common Haar state model: Feasibility results and separations. In *TCC 2024, Part II*, volume 15365 of *LNCS*, pages 94–125. Springer, Cham, December 2024.

[AHH24] Prabhanjan Ananth, Zihan Hu, and Zikuan Huang. Quantum key-revocable dual-regev encryption, revisited. In *TCC 2024, Part III*, volume 15366 of *LNCS*, pages 257–288. Springer, Cham, December 2024.

[AHU19] Andris Ambainis, Mike Hamburg, and Dominique Unruh. Quantum security proofs using semi-classical oracles. In *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 269–295. Springer, Cham, August 2019.

[AJL+12] Gilad Asharov, Abhishek Jain, Adriana López-Alt, Eran Tromer, Vinod Vaikuntanathan, and Daniel Wichs. Multiparty computation with low communication, computation and interaction via threshold FHE. In *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 483–501. Springer, Berlin, Heidelberg, April 2012.

[AK25] Kasra Abbaszadeh and Jonathan Katz. Non-interactive zero-knowledge arguments with certified deletion. In *ASIACRYPT 2025, Part VIII*, LNCS, pages 381–410. Springer, Singapore, December 2025.

[AKN+23] Shweta Agrawal, Fuyuki Kitagawa, Ryo Nishimaki, Shota Yamada, and Takashi Yamakawa. Public key encryption with secure key leasing. In *EUROCRYPT 2023, Part I*, volume 14004 of *LNCS*, pages 581–610. Springer, Cham, April 2023.

[AL21] Prabhanjan Ananth and Rolando L. La Placa. Secure software leasing. In *EUROCRYPT 2021, Part II*, volume 12697 of *LNCS*, pages 501–530. Springer, Cham, October 2021.

[AMP25] Prabhanjan Ananth, Saachi Mutreja, and Alexander Poremba. Revocable Encryption, Programs, and More: The Case of Multi-Copy Security. In *6th Conference on Information-Theoretic Cryptography (ITC 2025)*, volume 343 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 9:1–9:23, Dagstuhl, Germany, 2025. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[APV23] Prabhanjan Ananth, Alexander Poremba, and Vinod Vaikuntanathan. Revocable cryptography from learning with errors. In *TCC 2023, Part IV*, volume 14372 of *LNCS*, pages 93–122. Springer, Cham, November / December 2023.

[AQY22] Prabhanjan Ananth, Luowen Qian, and Henry Yuen. Cryptography from pseudorandom quantum states. In *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 208–236. Springer, Cham, August 2022.

[BCC+09] Mira Belenkiy, Jan Camenisch, Melissa Chase, Markulf Kohlweiss, Anna Lysyanskaya, and Hovav Shacham. Randomizable proofs and delegatable anonymous credentials. In *Annual international cryptology conference*, pages 108–125. Springer, 2009.

[BCC+16] Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Essam Ghadafi, and Jens Groth. Foundations of fully dynamic group signatures. In *ACNS 2016*, volume 9696 of *LNCS*, pages 117–136. Springer, Cham, June 2016.

[BFM88] Manuel Blum, Paul Feldman, and Silvio Micali. Non-interactive zero-knowledge and its applications (extended abstract). In *20th ACM STOC*, pages 103–112. ACM Press, May 1988.

[BG90] Mihir Bellare and Shafi Goldwasser. New paradigms for digital signatures and message authentication based on non-interactive zero knowledge proofs. In *CRYPTO'89*, volume 435 of *LNCS*, pages 194–211. Springer, New York, August 1990.

[BGK+24] James Bartusek, Vipul Goyal, Dakshita Khurana, Giulio Malavolta, Justin Raizes, and Bhaskar Roberts. Software with certified deletion. In *EUROCRYPT 2024, Part IV*, volume 14654 of *LNCS*, pages 85–111. Springer, Cham, May 2024.

[BI20] Anne Broadbent and Rabib Islam. Quantum encryption with certified deletion. In *TCC 2020, Part III*, volume 12552 of *LNCS*, pages 92–122. Springer, Cham, November 2020.

[BJL+21] Anne Broadbent, Stacey Jeffery, Sébastien Lord, Supartha Podder, and Aarthi Sundaram. Secure software leasing without assumptions. In *TCC 2021, Part I*, volume 13042 of *LNCS*, pages 90–120. Springer, Cham, November 2021.

[BJRT26] James Bartusek, Ruta Jawale, Justin Raizes, and Kabir Tomer. A new approach to arguments of quantum knowledge. In *Annual International Cryptology Conference*, pages 95–124. Springer, 2026.

[BK23] James Bartusek and Dakshita Khurana. Cryptography with certified deletion. In *CRYPTO 2023, Part V*, volume 14085 of *LNCS*, pages 192–223. Springer, Cham, August 2023.

[BKM+23] James Bartusek, Dakshita Khurana, Giulio Malavolta, Alexander Poremba, and Michael Walter. Weakening assumptions for publicly-verifiable deletion. In *Theory of Cryptography Conference*, pages 183–197. Springer, 2023.

[BKP23] James Bartusek, Dakshita Khurana, and Alexander Poremba. Publicly-verifiable deletion via target-collapsing functions. In *Annual International Cryptology Conference*, pages 99–128. Springer, 2023.

[BKS23] James Bartusek, Dakshita Khurana, and Akshayaram Srinivasan. Secure computation with shared EPR pairs (or: How to teleport in zero-knowledge). In *CRYPTO 2023, Part V*, volume 14085 of *LNCS*, pages 224–257. Springer, Cham, August 2023.

[BOV03] Boaz Barak, Shien Jin Ong, and Salil P. Vadhan. Derandomization in cryptography. In *CRYPTO 2003*, volume 2729 of *LNCS*, pages 299–315. Springer, Berlin, Heidelberg, August 2003.

[BP15] Nir Bitansky and Omer Paneth. ZAPs and non-interactive witness indistinguishability from indistinguishability obfuscation. In *TCC 2015, Part II*, volume 9015 of *LNCS*, pages 401–427. Springer, Berlin, Heidelberg, March 2015.

[BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *ACM CCS 93*, pages 62–73. ACM Press, November 1993.

[BR24] James Bartusek and Justin Raizes. Secret sharing with certified deletion. In *CRYPTO 2024, Part VII*, volume 14926 of *LNCS*, pages 184–214. Springer, Cham, August 2024.

[Bra18] Zvika Brakerski. Quantum FHE (almost) as secure as classical. In *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 67–95. Springer, Cham, August 2018.

[CCH+19] Ran Canetti, Yilei Chen, Justin Holmgren, Alex Lombardi, Guy N. Rothblum, Ron D. Rothblum, and Daniel Wichs. Fiat-Shamir: from practice to theory. In *51st ACM STOC*, pages 1082–1090. ACM Press, June 2019.

[CGJL25] Orestis Chardouvelis, Vipul Goyal, Aayush Jain, and Jiahui Liu. Quantum key leasing for PKE and FHE with a classical lessor. In *EUROCRYPT 2025, Part III*, volume 15603 of *LNCS*, pages 248–277. Springer, Cham, May 2025.

[ÇGLZR24] Alper Çakan, Vipul Goyal, Chen-Da Liu-Zhang, and João Ribeiro. Unbounded leakage-resilience and intrusion-detection in a quantum world. In *TCC 2024, Part II*, volume 15365 of *LNCS*, pages 159–191. Springer, Cham, December 2024.

[CGR26] Alper Cakan, Vipul Goyal, and Justin Raizes. How to delete without a trace: Certified deniability in a quantum world. In *Annual International Cryptology Conference*, pages 183–214. Springer, 2026.

[CKNY25] Jeffrey Champion, Fuyuki Kitagawa, Ryo Nishimaki, and Takashi Yamakawa. Untelegraphable encryption and its applications. In *Theory of Cryptography Conference*, pages 3–35. Springer, 2025.

[CLLZ21] Andrea Coladangelo, Jiahui Liu, Qipeng Liu, and Mark Zhandry. Hidden cosets and applications to unclonable cryptography. In *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 556–584, Virtual Event, August 2021. Springer, Cham.

[CMP24] Andrea Coladangelo, Christian Majenz, and Alexander Poremba. Quantum copy-protection of compute-and-compare programs in the quantum random oracle model. *Quantum*, 8:1330, 2024.

[DN00] Cynthia Dwork and Moni Naor. Zaps and their applications. In *41st FOCS*, pages 283–293. IEEE Computer Society Press, November 2000.

[FLS90] U. Feige, D. Lapidot, and A. Shamir. Multiple non-interactive zero knowledge proofs based on a single random string. In *Proceedings [1990] 31st Annual Symposium on Foundations of Computer Science*, pages 308–317 vol.1, 1990.

[FLS99] Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple noninteractive zero knowledge proofs under general assumptions. *SIAM Journal on computing*, 29(1):1–28, 1999.

[For87] Lance Fortnow. The complexity of perfect zero-knowledge (extended abstract). In *19th ACM STOC*, pages 204–209. ACM Press, May 1987.

[FS87] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *CRYPTO'86*, volume 263 of *LNCS*, pages 186–194. Springer, Berlin, Heidelberg, August 1987.

[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 *54th FOCS*, pages 40–49. IEEE Computer Society Press, October 2013.

[GGP10] Rosario Gennaro, Craig Gentry, and Bryan Parno. Non-interactive verifiable computing: Outsourcing computation to untrusted workers. In *CRYPTO 2010*, volume 6223 of *LNCS*, pages 465–482. Springer, Berlin, Heidelberg, August 2010.

[GLR+25] Aparna Gupte, Jiahui Liu, Justin Raizes, Bhaskar Roberts, and Vinod Vaikuntanathan. Quantum one-time programs, revisited. In *57th ACM STOC*, pages 213–221. ACM Press, June 2025.

[GMR24] Vipul Goyal, Giulio Malavolta, and Justin Raizes. Unclonable commitments and proofs. In *TCC 2024, Part III*, volume 15366 of *LNCS*, pages 193–224. Springer, Cham, December 2024.

[GO94] Oded Goldreich and Yair Oren. Definitions and properties of zero-knowledge proof systems. *Journal of Cryptology*, 7(1):1–32, December 1994.

[GOS06] Jens Groth, Rafail Ostrovsky, and Amit Sahai. Non-interactive zaps and new techniques for NIZK. In *CRYPTO 2006*, volume 4117 of *LNCS*, pages 97–111. Springer, Berlin, Heidelberg, August 2006.

[Gro06] Jens Groth. Simulation-sound NIZK proofs for a practical language and constant size group signatures. In *ASIACRYPT 2006*, volume 4284 of *LNCS*, pages 444–459. Springer, Berlin, Heidelberg, December 2006.

[HKM+24] Taiga Hiroka, Fuyuki Kitagawa, Tomoyuki Morimae, Ryo Nishimaki, Tapas Pal, and Takashi Yamakawa. Certified everlasting secure collusion-resistant functional encryption, and more. In *EUROCRYPT 2024, Part III*, volume 14653 of *LNCS*, pages 434–456. Springer, Cham, May 2024.

[HMNY22] Taiga Hiroka, Tomoyuki Morimae, Ryo Nishimaki, and Takashi Yamakawa. Certified everlasting zero-knowledge proof for QMA. In *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 239–268. Springer, Cham, August 2022.

[HU19] Dennis Hofheinz and Bogdan Ursu. Dual-mode NIZKs from obfuscation. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part I*, volume 11921 of *LNCS*, pages 311–341. Springer, Cham, December 2019.

[HXY19] Minki Hhan, Keita Xagawa, and Takashi Yamakawa. Quantum random oracle model with auxiliary input. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part I*, volume 11921 of *LNCS*, pages 584–614. Springer, Cham, December 2019.

[JK24] Ruta Jawale and Dakshita Khurana. Unclonable non-interactive zero-knowledge. In *ASIACRYPT 2024, Part IX*, volume 15492 of *LNCS*, pages 94–128. Springer, Singapore, December 2024.

[JLS18] Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 126–152. Springer, Cham, August 2018.

[KLYY26] Fuyuki Kitagawa, Jiahui Liu, Shota Yamada, and Takashi Yamakawa. A unified approach to quantum key leasing with a classical lessor. In *Annual International Cryptology Conference*, pages 90–121. Springer, 2026.

[KMY25] Fuyuki Kitagawa, Tomoyuki Morimae, and Takashi Yamakawa. A simple framework for secure key leasing. In *EUROCRYPT 2025, Part III*, volume 15603 of *LNCS*, pages 217–247. Springer, Cham, May 2025.

[KN22] Fuyuki Kitagawa and Ryo Nishimaki. Functional encryption with secure key leasing. In *ASIACRYPT 2022, Part IV*, volume 13794 of *LNCS*, pages 569–598. Springer, Cham, December 2022.

[KNP25] Fuyuki Kitagawa, Ryo Nishimaki, and Nikhil Pappu. PKE and ABE with collusion-resistant secure key leasing. In *CRYPTO 2025, Part III*, volume 16002 of *LNCS*, pages 35–68. Springer, Cham, August 2025.

[KNP26] Fuyuki Kitagawa, Ryo Nishimaki, and Nikhil Pappu. Collusion-resistant quantum secure key leasing beyond decryption. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 571–602. Springer, 2026.

[KNY21] Fuyuki Kitagawa, Ryo Nishimaki, and Takashi Yamakawa. Secure software leasing from standard assumptions. In *TCC 2021, Part I*, volume 13042 of *LNCS*, pages 31–61. Springer, Cham, November 2021.

[KNY23] Fuyuki Kitagawa, Ryo Nishimaki, and Takashi Yamakawa. Publicly verifiable deletion from minimal assumptions. In *TCC 2023, Part IV*, volume 14372 of *LNCS*, pages 228–245. Springer, Cham, November / December 2023.

[KR19] Yael Tauman Kalai and Leonid Reyzin. A survey of leakage-resilient cryptography. In *Providing Sound Foundations for Cryptography: On the Work of Shafi Goldwasser and Silvio Micali*, pages 727–794. 2019.

[KS25] Jonathan Katz and Ben Sela. Secret sharing with publicly verifiable deletion. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 278–304. Springer, 2025.

[Mah20] Urmila Mahadev. Classical homomorphic encryption for quantum circuits. *SIAM Journal on Computing*, 52(6):FOCS18–189, 2020.

[MNY24] Tomoyuki Morimae, Barak Nehoran, and Takashi Yamakawa. Unconditionally secure commitments with quantum auxiliary inputs. In *CRYPTO 2024, Part VII*, volume 14926 of *LNCS*, pages 59–92. Springer, Cham, August 2024.

[MPY24] Tomoyuki Morimae, Alexander Poremba, and Takashi Yamakawa. Revocable Quantum Digital Signatures. In *19th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2024)*, volume 310 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 5:1–5:24, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[MW18] Sanketh Menda and John Watrous. Oracle separations for quantum statistical zero-knowledge. *arXiv preprint arXiv:1801.08967*, 2018.

[MY22a] Tomoyuki Morimae and Takashi Yamakawa. Classically verifiable NIZK for QMA with preprocessing. In *ASIACRYPT 2022, Part IV*, volume 13794 of *LNCS*, pages 599–627. Springer, Cham, December 2022.

[MY22b] Tomoyuki Morimae and Takashi Yamakawa. Quantum commitments and signatures without one-way functions. In *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 269–295. Springer, Cham, August 2022.

[Nao91] Moni Naor. Bit commitment using pseudorandomness. *Journal of cryptology*, 4(2):151–158, 1991.

[NY90] Moni Naor and Moti Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *22nd ACM STOC*, pages 427–437. ACM Press, May 1990.

[Pas13] Rafael Pass. Unprovable security of perfect NIZK and non-interactive non-malleable commitments. In *TCC 2013*, volume 7785 of *LNCS*, pages 334–354. Springer, Berlin, Heidelberg, March 2013.

[PS05] Rafael Pass and Abhi Shelat. Unconditional characterizations of non-interactive zero-knowledge. In *Annual International Cryptology Conference*, pages 118–134. Springer, 2005.

[PS19] Chris Peikert and Sina Shiehian. Noninteractive zero knowledge for NP from (plain) learning with errors. In *CRYPTO 2019, Part I*, volume 11692 of *LNCS*, pages 89–114. Springer, Cham, August 2019.

[PWYZ24] Duong Hieu Phan, Weiqiang Wen, Xingyu Yan, and Jinwei Zheng. Adaptive hardcore bit and quantum key leasing over classical channel from LWE with polynomial modulus. In *ASIACRYPT 2024, Part IX*, volume 15492 of *LNCS*, pages 185–214. Springer, Singapore, December 2024.

[QRW19] Willy Quach, Ron D. Rothblum, and Daniel Wichs. Reusable designated-verifier NIZKs for all NP from CDH. In *EUROCRYPT 2019, Part II*, volume 11477 of *LNCS*, pages 593–621. Springer, Cham, May 2019.

[Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *37th ACM STOC*, pages 84–93. ACM Press, May 2005.

[SCG+14] Eli Ben Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *2014 IEEE symposium on security and privacy*, pages 459–474. IEEE, 2014.

[Shm21] Omri Shmueli. Multi-theorem designated-verifier NIZK for QMA. In *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 375–405, Virtual Event, August 2021. Springer, Cham.

[Shm22a] Omri Shmueli. Public-key quantum money with a classical bank. In *54th ACM STOC*, pages 790–803. ACM Press, June 2022.

[Shm22b] Omri Shmueli. Semi-quantum tokenized signatures. In *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 296–319. Springer, Cham, August 2022.

[SW14] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In *46th ACM STOC*, pages 475–484. ACM Press, May / June 2014.

[Unr14] Dominique Unruh. Revocable quantum timed-release encryption. In *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 129–146. Springer, Berlin, Heidelberg, May 2014.

[Wat24] Brent Waters. A new approach for non-interactive zero-knowledge from learning with errors. In *56th ACM STOC*, pages 399–410. ACM Press, June 2024.

[WWW25] Brent Waters, Hoeteck Wee, and David J. Wu. New techniques for preimage sampling: Improved NIZKs and more from LWE. In *EUROCRYPT 2025, Part IV*, volume 15604 of *LNCS*, pages 3–33. Springer, Cham, May 2025.

[Zha12] Mark Zhandry. How to construct quantum random functions. In *53rd FOCS*, pages 679–687. IEEE Computer Society Press, October 2012.