How to Define Expected Quantum Polynomial-Time Zero Knowledge Simulation

Zhengnan Lai* zl345@cornell.edu
Nicholas Spooner* nspooner@cornell.edu
Max Tromanhauser* mft55@cornell.edu

Abstract

Zero knowledge is formalized via a simulator — i.e., an efficient computation which simulates the view of a (malicious) verifier. The foundational results in constant-round zero knowledge [GMW86, FS90, GK96] all use expected polynomial-time (EPT) simulators, and there is evidence that strict poly-time simulators do not exist for these protocols [BL02]. In the post-quantum setting, we must upgrade the simulator to at least quantum polynomial time (QPT) in order to properly simulate quantum verifiers. Chia et al. [CCLY22] proved a surprising negative result which precludes non-trivial ZK for constant-round protocols with both (strict) QPT and a natural notion of expected quantum polynomial-time (EQPT) black-box simulation. In light of this, Lombardi, Spooner, and Ma [LMS22] introduced a novel EQPT notion, coherent-runtime EQPT or EQPT_c, and showed that the [GMW86, FS90, GK96] protocols all allow for EQPT_c simulation.

In this work, we identify a fundamental issue with the definition of EQPT_c simulation, and propose a resolution. In particular, we demonstrate that EQPT_c computation is not necessarily efficient and can, in fact, decide any classical decision problem. This is possible through a freedom of choice in selecting a unitary dilation for an efficient quantum channel. We propose an revised definition which carefully restricts this choice, and prove that the definition preserves the zero knowledge of the [GMW86, FS90, GK96] protocols. Additionally, by upgrading the [GK96] framework to the fully quantum setting, we demonstrate for the first time a constant-round (malicious verifier) zero knowledge proof system for QMA (with EQPT_c simulation).

1 Introduction

This paper is about zero knowledge proof systems (ZKPs) which allow a prover to convince an efficient verifier that a statement is valid without revealing any additional information. Zero knowledge is formalized via a *simulator*: an efficient computation which simulates the view of an efficient (malicious) verifier.

An important aspect of that formalization is the notion of computational efficiency of the simulator. In the classical setting, the most common cryptographic and complexity-theoretic definition of efficient computation is strict probabilistic polynomial time (strict PPT), but many seminal constructions of (constant-round) ZKPs are not known to have a strict PPT simulator [GK96, GMW86]. Moreover, there is evidence that strict PPT simulators do not exist for these protocols: constant-

*Cornell University

round ZKPs with *black-box* strict PPT simulation and negligible soundness error exist only for languages in BPP [BL02]. While there are ways to circumvent this barrier via non-black-box simulation techniques [Bar01, BP12, BS20], this requires tailoring the protocol to the simulation technique, and it seems unlikely that "simple" protocols admit such simulators.

To remedy this issue, the standard classical definition of zero knowledge permits the simulator to run in *expected* polynomial time (EPT), a variable-time randomized computation which halts after a polynomial number of steps in expectation (but may take any finite number of steps with positive probability).

In recent years, advances in quantum computing have led to a reevaluation of classical notions of efficiency and security. For zero knowledge, showing security against quantum adversaries entails exhibiting a simulator that can simulate the view of any efficient *quantum* verifier [Wat06]. Naturally, this requires that the simulator be at least a QPT machine. Similarly to the classical setting, however, there is evidence that *strict* QPT simulation is insufficient for natural constant-round ZKPs [CCLY22].[1] This raises the question of whether there is a natural computational model for (black-box) zero knowledge simulation in the quantum setting. Given the classical situation, a quantum version of EPT would seem a natural candidate. Below, we survey existing attempts at defining quantum EPT and where they fall short.

1.1 Expected Quantum Polynomial-Time (EQPT) Computation

There is surprising subtlety in formalizing quantum computation with variable running time. How does a quantum Turing machine (QTM) know when to halt? In Deutsch's original formulation [Deu85], QTMs evolve unitarily *except* that at each step a measurement is made to determine if the computation has halted. One can define EQPT in this model as those computations where the expected time before obtaining a halt outcome is polynomial. Lombardi, Ma, and Spooner [LMS22] refer to this definition as "measured-runtime" EQPT or EQPT_m (Fig. 1).

Issues with this definition were initially observed decades ago in the context of the quantum halting problem. As pointed out by Myers [Mye97], if different branches of computation halt at different times, then this measurement will necessarily collapse the state of the machine, preventing useful interference between those branches. This nuance turns out to be a critical issue for zero knowledge simulation. Specifically, Chia et al. [CCLY22] showed that this phenomenon leads to a striking impossibility: constant-round ZKPs with black-box EQPT_m simulation exist only for languages in BQP.

This raised the question of whether this obstacle can be circumvented by choosing a different "halting scheme". After some debate [Oza98a, LP98, Oza98b], it became clear that this effect is inherent. In particular, [LP98] noted that regardless of whether we actually perform the halt measurement or not, each branch of computation will necessarily become entangled with some description of its runtime, preventing interference between branches with different runtimes.[2] Thus, in the context of zero knowledge simulation, we will need to look beyond standard notions of EQPT.

[1] This does not follow immediately from the classical impossibility, as the simulator is potentially more powerful.

[2] In light of this issue, subsequent definitions of efficient quantum computation have typically focused on the circuit model, in which the running time is inherently fixed [BV97, BBBV97].

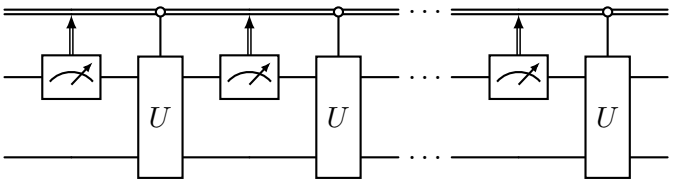

Figure 1: EQPT_m Computation

Given the discussion so far, one possible way forward is to settle for a relaxed definition of zero knowledge called ε-ZK [DNS98], in which the simulator receives a simulation quality parameter ε and is permitted to run in time $\mathsf{poly}(1/\varepsilon)$. This definition admits strict polynomial time simulation in both the classical and quantum settings, and is generally sufficient for applications of zero knowledge, since the parameter ε can be tuned in the reduction to a sufficiently small inverse polynomial.

In contrast, in this work we will aim to preserve as closely as possible the original definition of ZK, and in particular the negligible simulation error, at the cost of relaxing the notion of efficiency for the simulator. This direction was first explored by Lombardi et al. [LMS22], who proposed an alternative model for variable-runtime computation, *coherent-runtime* expected quantum polynomial-time (EQPT_c), where the halting measurement is "purified" in the following sense. Recall that any quantum channel M can be realized by the following steps: (1) expand the Hilbert space with an ancilla register, (2) apply a unitary dilation $\widehat{M}$ of M on that larger system, and (3) trace out the ancillas [Sti55]. The EQPT_c model bestows the simulator with the power to compute and (crucially) uncompute the dilation $\widehat{M}$ of an EQPT_m computation M (in addition to standard strict QPT computation). Concretely, an EQPT_c computation has the structure depicted in Fig. 2, where (C_1, C_2, C_3) are polynomial-sized circuits, and $\widehat{M}$ is a dilation of an EQPT_m computation.

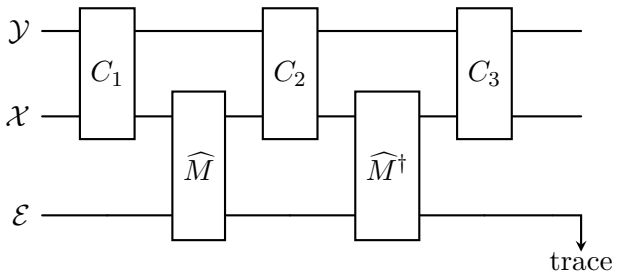

Figure 2: EQPT_c Computation (Simplified)

In order to justify this model, it should possess at least the following two properties:

1. There should exist post-quantum ZKPs with EQPT_c simulation for non-trivial languages, and
2. EQPT_c simulators should in some way be "limited" to performing efficient computations.

Both properties are highly nontrivial to establish. Lombardi et al. demonstrated (1) by proving that seminal constant-round ZKPs [GMW86, FS90, GK96] are post-quantum zero knowledge

with EQPT_c simulation. As for (2), note that because the running time of M is not polynomially bounded, the unitary $\widehat{M}$ is in general implemented by a circuit of superpolynomial (even unbounded) size. Hence it is not clear *a priori* that EQPT_c simulators cannot perform, say, exponential-time computations.

Lombardi et al. attempt to address this issue via *truncation*. They exhibit, for any EQPT_m computation M, a dilation $\widehat{M}$ such that circuits of the form presented in Fig. 2 can be approximated to within trace distance ε via a circuit of size $\mathsf{poly}(1/\varepsilon)$. This construction is central to their theory of EQPT_c, and they rely on it heavily in their arguments for both (1) and (2).

1.2 Our Contributions

In this work, we point out a flaw in the definition of EQPT_c simulation and propose a resolution. In this section, we give a brief overview of that flaw and that resolution. With those established, we then give results on the post-quantum zero knowledge of [GMW86, FS90, GK96] under our revised definition.

The problem with the prior definition. In [LMS22], the authors describe a "canonical" dilation for an EQPT_m computation, which they prove can be efficiently approximated; moreover, they show that any EQPT_c circuit as in Fig. 2 which utilizes this dilation can also be efficiently approximated. They then claim that, because unitary dilations are equivalent up to a local isometry on the ancilla, this property will hold for EQPT_c circuits regardless of the chosen dilation. Recall that the EQPT_c circuit cannot access the ancilla except through the dilation, so it may seem natural that the isometry on the ancilla will not affect the circuit output.

In fact, it is not the case that computations of the structure in Fig. 2 are invariant for different choices of dilation $\widehat{M}$ of the same channel M. As a minimal example, consider the following.

Example 1.1. Let $M(\rho) = |0\rangle\langle 0|\,\rho\,|0\rangle\langle 0| + |1\rangle\langle 1|\,\rho\,|1\rangle\langle 1|$ be the single-qubit channel that measures ρ in the computational basis and discards the result. Then both

$$\widehat{M}_0 = \mathrm{CNOT}_{\mathcal{X}\to\mathcal{E}}, \qquad \text{and} \qquad \widehat{M}_1 = \mathrm{SWAP}_{\mathcal{X}\leftrightarrow\mathcal{E}} \circ \mathrm{CNOT}_{\mathcal{X}\to\mathcal{E}}$$

are valid dilations of M in that $M(\rho) = \mathrm{Tr}_{\mathcal{E}}(\widehat{M}_0\rho\widehat{M}_0^\dagger) = Tr_{\mathcal{E}}(\widehat{M}_1\rho\widehat{M}_1^\dagger)$. Now let $C_1 = C_3 = I$ act trivially and $C_2 = H$ be the Hadamard gate on $\mathcal{X}$. If we instantiate an EQPT_c computation with $\widehat{M}_0$ and $\widehat{M}_1$ as in Fig. 2, we get the following circuits D_0 and D_1, respectively.

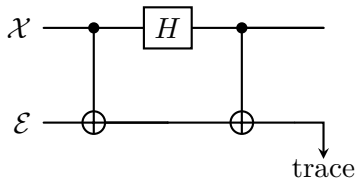

(a) D_0

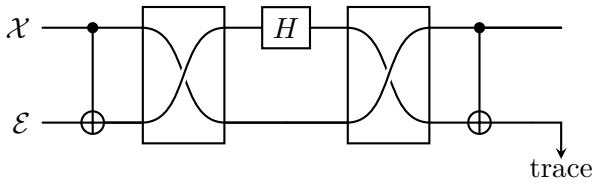

(b) D_1

For the input $|0\rangle$ on $\mathcal{X}$ (with $\mathcal{E}$ initialized to $|0\rangle$), the output of D_0 is maximally mixed but D_1

will always output $|0\rangle$.

The above already demonstrates an inconsistency in the definition. This particular example may seem benign, but in Section 3.1 we describe a more involved example where a particular choice of dilation allows an EQPT_c computation to solve an arbitrary classical decision problem. Thus not only is the computation sensitive to the choice of dilation, but some choices of dilation allow the model to perform (presumably) inefficient computations.

Our proposed definition. In order to resolve this issue, we redefine EQPT_c simulation in such a way that preserves the main results and spirit of [LMS22] while avoiding the above problems. In the new definition, an EQPT_c zero knowledge simulator is correct only if, for all "reasonable" dilations of the underlying EQPT_m procedure, its output is indistinguishable from the real view of the verifier. This requires some theory for what constitutes a "reasonable" dilation, and we introduce a new framework for describing dilations which we call *approximate dilation schemes*.

An approximate dilation scheme is an oracle circuit such that, when the oracle is instantiated with an EQPT_m computation M, the resulting unitary is a dilation of M. In this way, a dilation scheme describes a black-box strategy for dilating variable-runtime procedures. We focus on dilation schemes which satisfy the following desirable properties. First, we require that the dilation can be *approximated* by a series of circuits whose size scales polynomially in the inverse of the approximation error. In particular, we want that any EQPT_c computation using the resulting unitary dilation can be ε-approximated by replacing the unitary with a circuit of size $\mathsf{poly}(\varepsilon^{-1})$. Second, we require that the scheme *and* its approximations are black-box in M. This prevents perverse behaviour where the dilation "knows" the computation and can perform it independently of the transition unitary. In the context of zero knowledge simulation, it ensures that the dilation respects the black-box access model; and if M is using an oracle, the black-box restriction allows the approximation to use it as well. This is vital for use in reductions, as will be seen in Section 4.

We describe a dilation as being *efficiently approximable* if it is the result of applying an approximate dilation scheme to an EQPT_m computation. Efficiently approximable dilations attempt to capture the key attributes of [LMS22]'s canonical dilation while avoiding pitfalls where simulatability could strongly depend on the choice of dilation. We can now redefine EQPT_c simulation as follows.

Definition 1.2 (Informal)**.** A protocol is statistical (resp. computational) zero knowledge with EQPT_c simulation if the simulator is correct for all (resp. efficiently approximable) dilations.

Note that in the statistical zero knowledge case the [LMS22] definition is essentially recovered. This could be expected because proving statistical indistinguishability does not require reductions to be efficient. Our first result re-establishes the connection between EQPT_c ZK and ε-ZK.

Theorem 1.3. *Under our revised definition, ZK with EQPT_c simulation implies ε-ZK.*

This result is obtained by observing that [LMS22]'s "canonical" dilation for EQPT_m computations satisfies the properties of an approximate dilation scheme. With the definition established, the remainder of the work is dedicated to reproving the EQPT_c-ZK results of [LMS22] under the new definition.

Theorem 1.4. *Under our revised definition, the [GMW86] graph nonisomorphism protocol, the Feige-Shamir [FS90] 4-message argument for NP (as instantiated in [LMS22]), and the Goldreich-Kahan [GK96] 5-message proof for NP are post-quantum zero knowledge with EQPT_c simulation, under the same assumptions as in [LMS22].*

As an additional contribution, we show how to upgrade the [GK96] protocol and its analysis to the fully quantum setting, yielding the first constant-round (malicious verifier) zero knowledge **proof** system for QMA (even with ε-ZK).

Theorem 1.5. *Assuming quantum commitments exist, QMA admits a 4-message zero knowledge proof system with black-box EQPT_c simulation.*

Discussion. The flaw we identify in the original definition of EQPT_c, and the complexity of our revised definition, point to substantial difficulties in defining expected quantum polynomial time for the purposes of (computational) zero knowledge simulation. Many reasonable candidate definitions we tried led to subtle issues when used in security proofs. Thus we view our main question as open: whether there is a truly natural definition of efficient simulation in the quantum setting that recovers negligible simulation error for constant-round protocols. Indeed, this work perhaps serves as further evidence that the "right" notion of ZK in the quantum setting is ε-ZK: it is much easier to specify and prove, and suffices for all applications of which the authors are aware.

2 Preliminaries

We use $|\psi\rangle$ to denote pure quantum states, and ρ to denote density matrices, which represent mixed quantum states. For a quantum state ρ, its trace norm is denoted as $\|\rho\|_1$. For a quantum channel $\Phi : \mathcal{X} \to \mathcal{Y}$, its diamond norm is defined as $\|\Phi\|_\diamond = \max_\rho \|(\Phi \otimes I)(\rho)\|_1$.

Definition 2.1. Let ρ, σ be quantum states indexed by a security parameter λ. They are statistically (resp. computationally) indistinguishable if for all unbounded (resp. non-uniform QPT) adversary Adv,

$$|\Pr[\mathsf{Adv}(\rho) = 1] - \Pr[\mathsf{Adv}(\sigma) = 1]| \leq \mathsf{negl}(\lambda).$$

Our simulator in Section 4 uses quantum commitments (Appendix A), and the following indistinguishability lemma from [Yan23].

Lemma 2.2 ([Yan23])**.** *Let $|\phi_0\rangle, |\phi_1\rangle$ be quantum states on $\mathcal{X}, \mathcal{Y}$, and $U_b = |\phi_b\rangle\langle\phi_b| \otimes X + (I - |\phi_b\rangle\langle\phi_b|) \otimes I$. If there exists a polynomial distinguisher S^{U_b} with oracle access to U_b and without direct access to $\mathcal{X}$ such that*

$$\Pr\left[S^{U_b}(|\phi_0\rangle) = 1\right] - \Pr\left[S^{U_b}(|\phi_1\rangle) = 1\right] \geq 1/\mathsf{poly}(\lambda).$$

Then there exists a quantum algorithm S that distinguishes $|\phi_0\rangle$ from $|\phi_1\rangle$.

2.1 Quantum Interactive Protocols

Definition 2.3. Let (P, V) be an interactive protocol for a language $L \subseteq \{0,1\}^n$, representing prover and verifier algorithms. (P, V) is a *zero-knowledge proof* if:

- Completeness: $\forall x \in L$, $\Pr[\langle P(x), V(x)\rangle \to 1] = 1$.
- Soundness: $\forall x \notin L$ and $\tilde{P}$, $\Pr\left[\langle \tilde{P}(x), V(x)\rangle \to 1\right] \leq \mathsf{negl}(\lambda)$.
- Zero Knowledge: $\exists$ an "efficient" simulator $\mathsf{Sim}^{\tilde{V}}$ such that $\forall x \in L$ and $\tilde{V}$, $\mathrm{View}_{\tilde{V}}(P(x), \tilde{V}(x)) \approx \mathsf{Sim}^{\tilde{V}}(x)$.

Recall the definition of a Σ-protocol, which consists of three rounds: commitment, challenge, and answer. A quantum Σ-protocol is a Σ-protocol in the normal sense, but in which provers can send quantum states to verifiers, instead of only classical messages.

Remark 2.4. *Our definitions in this work refer only to ZK proof systems; i.e., protocols with statistical soundness. It is straightforward to modify them for computationally sound protocols (argument systems); the zero knowledge property is unaffected.*

Definition 2.5 ([BG20])**.** A *quantum Σ-protocol* consists of a three-round protocol between a prover and a verifier, which takes the following form:

- In the first round, the prover sends an initial quantum register $\mathcal{A}$ as a commitment.
- In the second round, the verifier sends a uniformly random challenge r.
- In the third round, the prover answers the challenge by sending a register $\mathcal{Z}$.

Importantly, [BG20] gave a zero-knowledge quantum Σ-protocol for the class QMA and its simulator, with non-negligible soundness error.

Recall that [GK96] uses a Σ-protocol that satisfies special honest verifier zero knowledge (SHVZK) to build a constant-round zero-knowledge proof for NP with negligible soundness error. We similarly define the special honest verifier zero knowledge for quantum Σ-protocols.

Definition 2.6. A quantum Σ-protocol is *special honest verifier zero knowledge* if it is zero-knowledge, and there is a quantum algorithm $\mathsf{SHVZK.Sim}(r)$ that takes in a challenge r and outputs a quantum state $\rho'_{\mathcal{A},\mathcal{Z}}$ such that the distribution of $(\rho'_{\mathcal{A},\mathcal{Z}}, r)$ is indistinguishable from the distribution of an honest-prover interaction $(\rho_{\mathcal{A},\mathcal{Z}}, r)$.

We will use the fact that the zero-knowledge protocol for QMA in [BG20] satisfies the SHVZK property. This is because the simulator, given the challenge r as its input, can produce a quantum state in $\mathcal{A}$ that is indistinguishable from the one in the real execution by using local density matrices of a simulatable witness for the string x.

3 EQPT Computations

In this section, we recall the definitions for EQPT_m and EQPT_c computation as in [LMS22], demonstrate the flaw in that formalization of EQPT_c, and propose an revised definition for EQPT_c zero knowledge simulation. Starting with EQPT_m, we first formalize what is meant by a variable-runtime quantum computation (which is also depicted in Fig. 4).

Definition 3.1. A (T-bounded) *variable-runtime computation M* consists of an input size n and a

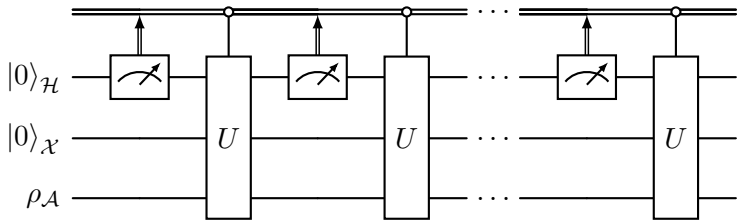

Figure 4: Variable-Runtime Computation

transition unitary circuit U acting on $\mathcal{A} \otimes \mathcal{X} \otimes \mathcal{H}$ where $\mathcal{A}$ is an n-qubit input/output register, $\mathcal{X}$ is a work register, and $\mathcal{H}$ is a single 'halt qubit'. Given an input state ρ on $\mathcal{A}$ and $\mathcal{X} \otimes \mathcal{H}$ initialized to $|0\rangle$, we execute M by repeating the following for at most T steps:

1. Measure the $\mathcal{H}$ in the standard basis.
2. If the outcome is 1, halt and discard all registers except $\mathcal{A}$.
3. Otherwise, apply U to $\mathcal{A} \otimes \mathcal{X} \otimes \mathcal{H}$.

If $\ell(\rho)$ is the random variable corresponding to the number of iterations performed for an input ρ, then we define M's *expected running time* to be $t_M = \max_\rho \mathbb{E}[\ell(\rho)] \cdot |U|$ where the maximization is over all states ρ on $\mathcal{A}$.

Remark 3.2. *Note that our definition differs slightly from [LMS22]'s. Their starting point is the formalism from [Deu85]'s quantum Turing machine. For clarity, we have dispensed with the QTM particulars (e.g., the head tape) but it is easy to see that these are equivalent up to aesthetic changes.*

Given that formalism, the definition for an EQPT_m computation is simply a family of variable-runtime computations which are efficient in the following sense.

Definition 3.3. Let $M = \{M_n\}_{n\in\mathbb{N}}$ be a family of ($T \leq 2^{O(n)}$-bounded) variable-runtime computations where the transition unitaries $\{U_n\}_{n\in\mathbb{N}}$ are a uniform family of circuits. Then we say M is an *EQPT_m computation* if there exists $n_0 \in \mathbb{N}$ and polynomial p such that $t_{M_n} \leq p(n)$ for all $n \geq n_0$.

Although this is a natural notion for EQPT, this is also the model of computation that [CCLY22] uses (implicitly) to garner the following negative result.

Theorem 3.4 ([CCLY22])**.** *There do not exist post-quantum, constant-round, black-box zero knowledge proofs for* NP *with EQPT_m simulation, unless* NP $\subseteq$ BQP*.*

3.1 [LMS22]'s EQPT_c and Its Issue

In light of this negative result, [LMS22] proposed EQPT_c, an alternative model of computation which allowed for the reversal of variable-runtime processes. Recall that any quantum channel M has many unitary dilations $\widehat{M}$. Given one such dilation, an EQPT_c computation is one which has access to $\widehat{M}$ and $\widehat{M}^\dagger$ in a limited fashion. In particular, it should have the form of a 'D-circuit', defined by [LMS22] as follows (and diagrammed in Fig. 5).

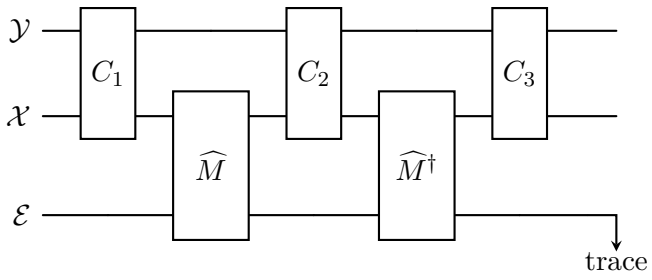

Figure 5: D-circuit $D[\widehat{M}]$ where $k = 1$.

Definition 3.5. A D-circuit is a quantum circuit D with special gates $\{G_i, G_i^{-1}\}_{i=1}^k$ with the following restriction: for each i, there is a single G_i gate and a single G_i^{-1} gate acting on a designated register $\mathcal{X}_i$, where G_i acts before G_i^{-1}. All other gates may act arbitrarily on $\mathcal{Y} \otimes \bigotimes_{i=1}^k \mathcal{X}_i$ for some register $\mathcal{Y}$. Given unitaries $\{U_i\}_{i=1}^k$ (where U_i acts on registers $\mathcal{X}_i \otimes \mathcal{E}_i$), let D' be the circuit on $\mathcal{Y} \otimes (\bigotimes_i \mathcal{X}_i \otimes \mathcal{E}_i)$ derived from D by replacing G_i with U_i and G_i^{-1} with U_i^{-1}. Then define the channel $D[U_1, \ldots, U_k]$ to be

$$D[U_1, \ldots, U_k](\rho) = \mathrm{Tr}_{\mathcal{E}} \left[D'(\rho \otimes |0\rangle\langle 0|_{\mathcal{E}}) \right]$$

Remark 3.6. *The original definition for D-circuits in [LMS22] is parameterized by CPTP maps $\{\Phi_i\}_{i=1}^k$ rather than unitary maps $\{U_i\}_{i=1}^k$. In that version, G_i is replaced by a unitary dilation $\widehat{\Phi}_i$ of Φ_i. This change is necessary because, as we will show later in this section, the output of a D-circuit may depend on that choice of unitary dilation. That is, the original definition does not uniquely specify a map.*

In this way, [LMS22]'s definition for EQPT_c computation is a channel which can be realized by a D-circuit instantiated with unitary dilations of EQPT_m computations.

Definition 3.7. A sequence of CPTP maps $\{\Phi_n\}_{n\in\mathbb{N}}$ is an *EQPT_c computation* if there exist a uniform family of poly-sized D-circuits $\{D_n\}_{n\in\mathbb{N}}$ and unitary dilations $\widehat{M}_1, \ldots, \widehat{M}_k$ of EQPT_m computations $M_1, \ldots, M_k$ such that $\Phi_n = D_n[\widehat{M}_1, \ldots \widehat{M}_k]$.

Lombardi et al. argued that an EQPT_c channel inherits the efficiency of the D-circuit and EQPT_m computations. The key to that argument is that — just as with EQPT_m — one can 'truncate' an EQPT_c computation to get a poly-sized circuit which approximates the channel with inverse-polynomial accuracy.

Claim 3.8 ([LMS22], Lemma 9.5)**.** *For any family of D-circuits $\{D_n\}_{n\in\mathbb{N}}$, EQPT_m computations $M_1, \ldots, M_k$, and $\varepsilon : \mathbb{N} \to (0, 1]$, there exist unitary dilations $\widehat{M}_1, \ldots \widehat{M}_k$ (as in Fig. 6) and a uniform sequence of (standard) quantum circuits $\{C_n\}_{n\in\mathbb{N}}$ of size $O(1/\varepsilon(n)^2)$ such that,*

$$\|D_n[\widehat{M}_1, \ldots \widehat{M}_k] - C_n\|_\diamond \leq \varepsilon(\lambda).$$

Observe that this claim only proves that a particular dilation can be approximated. That "canonical" dilation, depicted in Fig. 6, is key to proving the claim. Intuitively, this dilation records

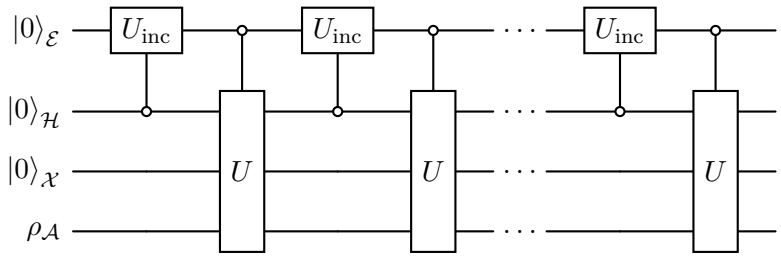

Figure 6: "Canonical" Dilation for a Variable-Runtime Computation.

the number of steps after halting in a 'clock' register which is incremented with a unitary U_{inc}. The transition unitary is then only applied to branches where the clock register has yet to be incremented. In this way, the dilation can be approximated simply by truncating the number of iterations.

[LMS22] then assert, because dilations are equivalent under a local isometry, Claim 3.8 holds for any dilation of the EQPT_m computations. This is not true, and we have already demonstrated in Example 1.1 that the output of the D-circuit can completely change depending on that choice. The following counterexample shows that this is a serious issue: by exploiting the choice of dilation, a D-circuit can compute hard classical decision problems.

Example 3.9. In this example, we design an EQPT_m computation M, a dilation $\widehat{M}$ of M, and a D-circuit D such that $D[\widehat{M}]$ can decide an arbitrary decision problem $f : \{0,1\}^n \to \{0,1\}$. In fact, the D-circuit will only require a simple single qubit measurement in the standard basis (depicted in Fig. 7).

Let $\mathcal{A}$ be M's input/output register, divided into $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3$ such that $\mathcal{A}_1$ and $\mathcal{A}_2$ are single qubit registers and $\mathcal{A}_3$ has n-qubits. Let $\mathcal{X}$ be the single qubit ancilla used by M and $\mathcal{H}$ the halt qubit. M induces the channel which replaces the state on $\mathcal{A}_1$ with the state $|+\rangle\langle+|$. It has the transition unitary

$$U = (X_{\mathcal{H}} \otimes I_{\mathcal{X}} \otimes H_{\mathcal{A}_1})(I_{\mathcal{H}} \otimes \text{SWAP}_{\mathcal{X} \leftrightarrow \mathcal{A}_1}) \otimes I_{\mathcal{A}_2\mathcal{A}_3}$$

which swaps the $\mathcal{A}_1$ qubit with a fresh ancilla and then applies H to map it to $|+\rangle$. It also immediately flips the halt qubit so all inputs will only undergo one iteration (independent of n). Now we define a unitary $\widehat{M}$ on $\mathcal{E} \otimes \mathcal{H} \otimes \mathcal{X} \otimes \mathcal{A}$:

$$\widehat{M} = (I_{\mathcal{E}} \otimes U_{\mathcal{H}\mathcal{A}\mathcal{X}})(|0\rangle\langle 0|_{\mathcal{X}} \otimes I + |1\rangle\langle 1|_{\mathcal{X}} \otimes U_f)$$

where U_f is the unitary on $\mathcal{A}_2 \otimes \mathcal{A}_3$ which performs the mapping $|a\rangle_{\mathcal{A}_2} |x\rangle_{\mathcal{A}_3} \mapsto |a \oplus f(x)\rangle_{\mathcal{A}_2} |x\rangle_{\mathcal{A}_3}$. Observe that $\widehat{M}$ is a unitary dilation of M — i.e., for any state ρ on $\mathcal{A}$,

$$\begin{aligned}
\text{Tr}_{\mathcal{E}\mathcal{H}\mathcal{X}}[\widehat{M}(|0\rangle\langle 0|_{\mathcal{E}\mathcal{H}\mathcal{X}} \otimes \rho_{\mathcal{A}})\widehat{M}^\dagger] &= \text{Tr}_{\mathcal{E}\mathcal{H}\mathcal{X}}[U(|0\rangle\langle 0|_{\mathcal{E}\mathcal{H}\mathcal{X}} \otimes \rho_{\mathcal{A}})U^\dagger] \\
&= \text{Tr}_{\mathcal{E}\mathcal{H}\mathcal{X}}[|0\rangle\langle 0|_{\mathcal{E}} \otimes |1\rangle\langle 1|_{\mathcal{H}} \otimes |+\rangle\langle +|_{\mathcal{A}_1} \otimes \rho_{\mathcal{X}\mathcal{A}_2\mathcal{A}_3}] \\
&= |+\rangle\langle +|_{\mathcal{A}_1} \otimes \text{Tr}_{\mathcal{A}_1}[\rho]
\end{aligned}$$

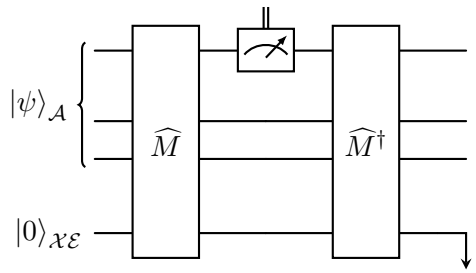

Figure 7: D-circuit $C[\widehat{M}]$.

Consider the D-circuit $D[\widehat{M}]$ given in Fig. 7 which runs $\widehat{M}$, measures $\mathcal{A}_1$ in the computational basis, and then applies $\widehat{M}^\dagger$. We will show that $D[\widehat{M}]$ can be used to decide $f(x)$ by providing the input state $|0\rangle_{\mathcal{A}_1\mathcal{A}_2}|x\rangle_{\mathcal{A}_3}$. After applying $\widehat{M}$ (with $\mathcal{E}$ and $\mathcal{X}$ initialized to $|0\rangle$), the state evolves to

$$\begin{aligned}\widehat{M}|0\rangle|x\rangle &= (I_{\mathcal{E}}\otimes U_{\mathcal{H}\mathcal{A}\mathcal{X}})(|0\rangle\langle 0|_{\mathcal{X}}\otimes I + |1\rangle\langle 1|_{\mathcal{X}}\otimes U_f)|0\rangle_{\mathcal{E}}|0\rangle_{\mathcal{X}}|0\rangle_{\mathcal{A}_1}|0\rangle_{\mathcal{A}_2}|x\rangle_{\mathcal{A}_3}\\ &= (I_{\mathcal{E}}\otimes U_{\mathcal{H}\mathcal{A}\mathcal{X}})|0\rangle_{\mathcal{E}}|0\rangle_{\mathcal{X}}|0\rangle_{\mathcal{A}_1}|0\rangle_{\mathcal{A}_2}|x\rangle_{\mathcal{A}_3}\\ &= |0\rangle_{\mathcal{E}}|1\rangle_{\mathcal{H}}|0\rangle_{\mathcal{X}}|+\rangle_{\mathcal{A}_1}|0\rangle_{\mathcal{A}_2}|x\rangle_{\mathcal{A}_3}\end{aligned}$$

The $\mathcal{A}_1$ qubit is then measured in the standard basis to get outcome b. This leaves us with the residual state

$$|\psi\rangle = |0\rangle_{\mathcal{E}}|1\rangle_{\mathcal{H}}|0\rangle_{\mathcal{X}}\frac{1}{\sqrt{2}}\left(|+\rangle + (-1)^b|-\rangle\right)_{\mathcal{A}_1}|0\rangle_{\mathcal{A}_2}|x\rangle_{\mathcal{A}_3}$$

Finally, we apply $\widehat{M}^\dagger$ to get:

$$\begin{aligned}\widehat{M}^\dagger|\psi\rangle &= (|0\rangle\langle 0|_{\mathcal{X}}\otimes I + |1\rangle\langle 1|_{\mathcal{X}}\otimes U_f)(\mathrm{SWAP}_{\mathcal{X}\mathcal{A}_1})(X_{\mathcal{H}}\otimes H_{\mathcal{A}_1})|\psi\rangle\\ &= \frac{1}{\sqrt{2}}(|0\rangle\langle 0|_{\mathcal{X}}\otimes I + |1\rangle\langle 1|_{\mathcal{X}}\otimes U_f)(\mathrm{SWAP}_{\mathcal{X}\mathcal{A}_1})|0\rangle_{\mathcal{E}}|0\rangle_{\mathcal{H}}|0\rangle_{\mathcal{X}}\left(|0\rangle + (-1)^b|1\rangle\right)_{\mathcal{A}_1}|0\rangle_{\mathcal{A}_2}|x\rangle_{\mathcal{A}_3}\\ &= \frac{1}{\sqrt{2}}(|0\rangle\langle 0|_{\mathcal{X}}\otimes I + |1\rangle\langle 1|_{\mathcal{X}}\otimes U_f)|0\rangle_{\mathcal{E}}|0\rangle_{\mathcal{H}}\left(|0\rangle + (-1)^b|1\rangle\right)_{\mathcal{X}}|0\rangle_{\mathcal{A}_1}|0\rangle_{\mathcal{A}_2}|x\rangle_{\mathcal{A}_3}\\ &= |0\rangle_{\mathcal{E}}|0\rangle_{\mathcal{H}}|0\rangle_{\mathcal{A}_1}\otimes\frac{1}{\sqrt{2}}\left(|0\rangle_{\mathcal{X}}|0\rangle_{\mathcal{A}_2} + (-1)^b|1\rangle_{\mathcal{X}}|f(x)\rangle_{\mathcal{A}_2}\right)\otimes|x\rangle_{\mathcal{A}_3}\end{aligned}$$

Once $\mathcal{E},\mathcal{X},\mathcal{H}$ are traced out, we are left with the mixed state:

$$|0\rangle\langle 0|_{\mathcal{A}_1}\otimes\frac{1}{2}\left(|0\rangle\langle 0| + |f(x)\rangle\langle f(x)|\right)_{\mathcal{A}_2}\otimes|x\rangle\langle x|_{\mathcal{A}_3}$$

To decide if $f(x)=1$, measure $\mathcal{A}_2$ in the computational basis. If $f(x)=0$, then the outcome will always be 0. Otherwise, with probability $\frac{1}{2}$ the outcome will be 1.

As a consequence, although [LMS22]'s definition for EQPT_c computation is still well defined, it captures channels which are certainly outside the realm of 'efficient computation'. Thus a zero knowledge protocol with EQPT_c simulation in this sense is not meaningful, in the same way that an unbounded simulator is trivially constructed.

3.2 A Revised Definition

Our goal is to restrict EQPT_c simulation such that the [LMS22] results can be recovered but our model of computation still excludes the 'degenerate' cases. One such approach would be to elevate the canonical dilation to be the definition — i.e., an EQPT_c simulator is a D-circuit instantiated with the canonical dilation of an EQPT_m computation. We feel that this is needlessly restrictive: in particular, it would prevent (e.g.) a dilation which encodes the clock in unary rather than binary.

We would like a definition which is both minimal, in the sense that it captures all reasonable dilation choices, and compatible with the security proofs that must use it. Towards that end, we introduce a new object we call an "approximate dilation scheme".

Definition 3.10. Let $C = \{C_\gamma^{(\cdot)}\}_{\gamma\in\mathbb{N}}$ be a uniform family of unitary oracle circuits acting on registers $\mathcal{AXHE}$. Given a variable-runtime computation M with transition unitary U, we use $C_\gamma[M]$ to denote the isometry

$$C_\gamma[M](\rho) := C_\gamma^U(\rho_\mathcal{A} \otimes |0\rangle\langle 0|_{\mathcal{XHE}})C_\gamma^{U\dagger}.$$

We say C is an *approximate dilation scheme* if the following three properties hold.

- **Black-Box.** C_γ^U may apply arbitrary gates on $\mathcal{E}$ but only acts on $(\mathcal{A},\mathcal{X},\mathcal{H})$ with the oracle gates $(\text{ctrl-}U, \text{ctrl-}U^\dagger)$ and coherent measurements of the halt qubit $\mathcal{H}$ in the standard basis (i.e., $\text{CNOT}_{\mathcal{H}\to\mathcal{E}}$ targeting some qubit in $\mathcal{E}$).
- **Dilation Correctness.** For all (T-bounded) variable-runtime computations M with transition unitary U, there exists $\Gamma(T)$ such that $C_\gamma[M] = C_{\gamma'}[M]$ for all $\gamma,\gamma' \geq \Gamma(T)$, and

$$M(\rho) = \text{Tr}_{\mathcal{XHE}}\left[C_{\Gamma(T)}[M](\rho)\right],$$

- **Efficient Approximation.** For some $\gamma_0 \in \mathbb{N}$, all $\gamma \geq \gamma_0$, and all (T-bounded) variable-runtime computations M with transition unitary U, $|C_\gamma| \leq \mathsf{poly}(\gamma, t_M)$ and, for all channels N on $\mathcal{A}$

$$\left\|\Phi(C_{\Gamma(T)}[M],N) - \Phi(C_\gamma[M],N)\right\|_\diamond \leq \frac{1}{\gamma}$$

where $\Phi(C_\gamma[M],A)$ is the D-circuit in Fig. 5 with $C_1, C_3 = I$, $C_2 = N$ corresponding to computing $C_\gamma[M]$, applying N, then uncomputing $C_\gamma[M]$ — i.e., for ρ on $\mathcal{A}$

$$\Phi(C_\gamma[M],N)(\rho) = \text{Tr}_\mathcal{E}\left[(C_\gamma^U)^\dagger N(C_\gamma[M](\rho))\, C_\gamma^U\right]$$

For brevity, given a (T-bounded) variable-runtime computation M, we denote $C[M] = C_{\Gamma(T)}[M]$.

Approximate dilation schemes allow us to classify the space of unitary dilations. In particular, we say a unitary dilation is "efficiently approximable" if it can be derived from some approximate dilation scheme.

Definition 3.11. Let M be a variable-runtime computation and $\widehat{M}$ be a unitary dilation of M. Then we say $\widehat{M}$ is *efficiently approximable* if there exists an approximate dilation scheme such that $C[M] = \widehat{M}$.

The next natural question is whether approximate dilation schemes exist at all; we show that this is indeed the case.

Lemma 3.12. *Approximate dilation schemes exist.*

Proof Sketch. The canonical dilation of [LMS22] (depicted in Fig. 6) fulfills Definition 3.10. Note that the black-box property is respected and only the transition unitary and a coherent measurement of the halt qubit are applied to the $\mathcal{A}, \mathcal{X}, \mathcal{H}$ registers. By construction, the truncation at T iterations is exactly a unitary dilation of the variable-runtime computation and so dilation correctness is satisfied. Finally, efficient approximation is follows immediately from [LMS22, Claim 9.5]. As a sketch, let C_γ be the circuit which performs $\gamma^2 t_M$ steps of the computation. C_γ will perfectly simulate the computation on branches which halt before γ, and the amplitude on the steps which have not halted yet is at most $1/\gamma$. □

Now we are prepared to give our revised definition for EQPT_c simulation. At a high level, we require the output of an EQPT_c simulator to be indistinguishable from the real view of the verifier regardless of its dilation choice. In the case of computational zero knowledge, this exact requirement would be too stringent. Instead, we require only that the simulator is correct for all "appropriate" dilations, which we define to be those which are efficiently approximable. In particular, if we are trying to prove the correctness of a computational EQPT_c simulator, the common approach is to reduce to the computational hardness of some cryptographic assumption. For that reduction to go through, the simulator must be efficient or at least have an efficient approximation. In general, as we've seen from the counterexamples, the dilation $\widehat{M}$ (along with $\widehat{M}^\dagger$) has no efficiency guarantees. Thus, the efficiently approximable assumption allows these types of arguments to go through. Perhaps unsurprisingly, this restriction is not necessary for statistical zero knowledge, as their security proofs do not rely on computational hardness.

Definition 3.13. An interactive proof (P, V) for a language L has *statistical (resp. computational) EQPT_c black-box zero knowledge simulation* if there exists a simulator $\mathsf{Sim} = (\{D_n\}, M_1, \ldots, M_k)$ consisting of a uniform family of D-circuits $\{D_n\}_{n\in\mathbb{N}}$ and EQPT_m computations $(M_1, \ldots, M_k)$ such that, for all (resp. efficiently approximable) unitary dilations $\widehat{M}_i$ of M_i, non-uniform QPT malicious verifiers $\tilde{V}$, statements $x \in L$, and unbounded (resp. non-uniform QPT) adversaries Adv,

$$\| \operatorname{View}_{\tilde{V}}(P(1^\lambda, x), \tilde{V}(1^\lambda, x)) - D_n[\widehat{M}_1, \ldots, \widehat{M}_k] \|_\diamond \leq \mathsf{negl}(\lambda).$$

where D_n and $(M_1, \ldots, M_k)$ have black-box access to $\tilde{V}$.

Under this definition of EQPT_c, we can prove the following which is a general version of Claim 3.8.

Lemma 3.14. *Let (P, V) be a zero knowledge proof with EQPT_c simulation. Then (P, V) is also ε-zero knowledge.*

Proof. Let $\mathsf{Sim} = (\{D_n\}, M_1, \cdots, M_k)$ be the EQPT_c simulator. Let $\widehat{M}_1, \ldots, \widehat{M}_k$ be efficiently approximable dilations of $M_1, \ldots, M_k$. By definition, $D_n[\widehat{M}_1, \ldots, \widehat{M}_k]$ δ-approximates the malicious verifier's view $\operatorname{View}_{\tilde{V}}(P(1^\lambda, x), \tilde{V}(1^\lambda, x))$ for some negligible $\delta(n)$ is [3]. Because $\widehat{M}_i$ is efficiently

[3] Note for statistical ZK, the simulator must be correct for all dilations, which includes those which are efficiently approximable.

approximable, there exists an approximate dilation scheme $C^{(i)}$ such that $\widehat{M}_i = C^{(i)}[M_i]$ for all i. In addition, there exists an approximation circuit $C_\gamma^{(i)}$ where $\gamma = ((\varepsilon - \delta)/k)^{-1}$ such that $C_\gamma^{(i)}[M_i]$ is an $((\varepsilon - \delta)/k)$-approximation of $C^{(i)}[M_i]$

Let $\mathsf{Sim}_\varepsilon = D_n[C_\gamma^{(1)}[M_1], \cdots, C_\gamma^{(k)}[M_k]]$ be the ε-simulator. By the triangle inequality, Sim_ε approximates the real view $\mathrm{View}_{\tilde{V}}(P(1^\lambda, x), \tilde{V}(1^\lambda, x))$ with ε accuracy. □

4 Constant-Round ZK for QMA with EQPT_c Simulation

In this section, we prove Theorem 1.5. In Section 4.1, we describe the 4-message protocol following the [GK96] framework. In Section 4.2, we describe the simulator and prove its correctness.

4.1 Protocol Construction & Notation

Let (P_Σ, V_Σ) be a quantum Σ-protocol for the class QMA satisfying SHVZK (as in [BG20]) and let Com be a statistically hiding, computationally binding quantum commitment. Then we construct a four message protocol following the [GK96] framework, described in Fig. 8.

1. The verifier commits to a uniformly sampled challenge r using Com to get commitment and decommitment registers $(\mathcal{C}, \mathcal{D})$. It sends $\mathcal{C}$ to the prover.
2. The prover prepares the initial message $\mathcal{A}$ using $P_\Sigma(x)$ and sends it to the verifier.
3. The verifier sends the decommitment register $\mathcal{D}$ along with r to the prover.
4. The prover aborts if $\mathcal{D}$ is not a valid opening for r. Otherwise, it prepares $\mathcal{Z}$ using $P_\Sigma(x, r)$ and sends it to the verifier.
5. The verifier accepts if V_Σ accepts the transcript $(x, \mathcal{A}, r, \mathcal{Z})$.

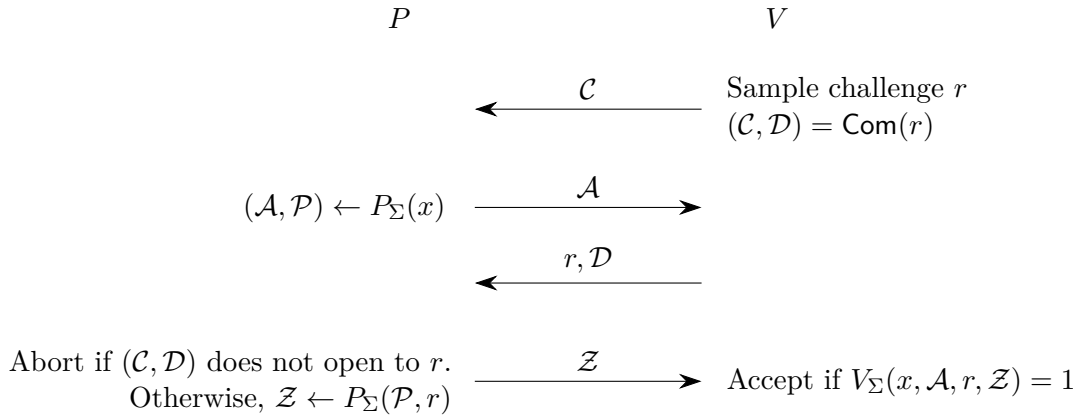

Figure 8: Zero Knowledge Proof for QMA

As it will be helpful later, we present some additional notation concerning the protocol. In addition to the registers for the quantum messages exchanged $(\mathcal{C}, \mathcal{A}, \mathcal{D}, \mathcal{Z})$, we let $\mathcal{R}$ be the register containing the classical challenge and $\mathcal{P}$ be the prover's internal state. When we have a malicious verifier $\tilde{V}$, we denote $\mathcal{V}$ as its internal state and $U_{\tilde{V}}$ to be the unitary on $\mathcal{V} \otimes \mathcal{A} \otimes \mathcal{R} \otimes \mathcal{D}$ which is applied between the prover's first message and the verifier's challenge. In the quantum context, black-box access to $\tilde{V}$ is modeled as having access to $\mathcal{V}$ only by applying the unitary $U_{\tilde{V}}, U_{\tilde{V}}^\dagger$, or their controlled versions.

Let $|P\rangle$ be a state on $\mathcal{P} \otimes \mathcal{A} \otimes \mathcal{Z}$ denoting the honest prover's state upon sending its first message. Note that we can assume $|P\rangle$ to be pure w.l.o.g. where $\mathcal{P}$ is expanded to accommodate any purification. Finally, let U_{P_r} be the unitary applied by P_Σ on $\mathcal{P} \otimes \mathcal{Z}$ to generate the response to challenge r and $|P_r\rangle = (I_\mathcal{A} \otimes U_{P_r}) |P\rangle$ be the final message and state of the honest prover conditioned on challenge r.

Let $\mathsf{SHVZK.Sim}$ be (P_Σ, V_Σ)'s SHVZK simulator as in Definition 2.6. In general $\mathsf{SHVZK.Sim}$ outputs a mixed state on $\mathcal{A} \otimes \mathcal{Z}$, but we denote by $|\mathsf{Sim}_r\rangle$ a purification of that output for a given input r by expanding to a larger Hilbert space. As $\mathsf{SHVZK.Sim}$ is simulating the prover (and to mirror $|P\rangle$), we reuse $\mathcal{P}$ to denote the purification register — i.e., $|\mathsf{Sim}_r\rangle$ is a state on $\mathcal{P} \otimes \mathcal{A} \otimes \mathcal{Z}$. We additionally define the projector

$$\Pi_{\mathsf{Sim}} = \sum_r |r\rangle\langle r|_\mathcal{S} \otimes |\mathsf{Sim}_r\rangle\langle \mathsf{Sim}_r|$$

Recall the projector V_{Com} from Definition A.1 which projects onto the space of valid commitments. Given black-box access to $\tilde{V}$, one can compute the projectors $\Pi_{\mathsf{Com}} = U_{\tilde{V}}^\dagger V_{\mathsf{Com}} U_{\tilde{V}}$ and $\Pi_r = U_{\tilde{V}}^\dagger |r\rangle\langle r|_\mathcal{R} U_{\tilde{V}}$ which act non-trivially on the registers $\mathcal{V} \otimes \mathcal{A} \otimes \mathcal{R} \otimes \mathcal{C} \otimes \mathcal{D}$ and $\mathcal{V} \otimes \mathcal{A} \otimes \mathcal{R} \otimes \mathcal{D}$, respectively.

The protocol is complete by the correctness of the commitment scheme and the completeness of the [BG20] protocol. Moreover, the protocol is sound by the statistical hiding property of the commitment scheme and the soundness of the [BG20] protocol. In the next section, we will prove that the protocol has an EQPT_c simulation.

4.2 Proof of EQPT_c Simulation

Our simulator makes use the following quantum algorithms from [LMS22] (via [GSLW19]): $\mathsf{VarEstimate}$ which estimates singular values, and $\mathsf{Transform}$ which maps left singular vectors to right singular vectors.

Theorem 4.1 ([LMS22])**.** *Let* $\mathsf{A} = (\Pi_\mathsf{A}, I - \Pi_\mathsf{A}), \mathsf{B} = (\Pi_\mathsf{B}, I - \Pi_\mathsf{B})$ *be projective measurements. There is a pair of algorithms* $\mathsf{VarEstimate}[\Pi_\mathsf{A} \rightleftarrows \Pi_\mathsf{B}]$ *and* $\mathsf{Transform}[\Pi_\mathsf{B} \to \Pi_\mathsf{A}]$ *with* $U_\mathsf{A}, U_\mathsf{B}$ *gates with the following properties. Let* $|w_{j,1}\rangle$ *be a left singular vector of* $\Pi_\mathsf{A}\Pi_\mathsf{B}$ *with singular value* $s_j > 0$*, and let* $|v_{j,1}\rangle$ *be the corresponding right singular vector. Then*

- *The subspace* $\mathcal{S}_j$ *is invariant under both* $\mathsf{VarEstimate}$ *and* $\mathsf{Transform}$*.*
- $\mathsf{VarEstimate}(|v_{j,1}\rangle)$ *has running time* $O(\log(1/\delta)/s_j)$ *with probability* $1 - \delta$ *and* $O(\log(1/\delta)/\delta)$ *with probability* 1*.*

$\mathsf{EDT}[\mathsf{A},\mathsf{B},\mathsf{C}]$

1. Apply measurement B obtaining outcome b_1.
2. Apply measurement A obtaining outcome b_2.
3. If $b_1 = 0$ or $b_2 = 0$, stop and output $(0, \perp)$.
4. Run $\mathsf{VarEstimate}_\delta[\mathsf{A} \rightleftarrows \mathsf{B}]$ obtaining output y.
5. Apply measurement C obtaining outcome c.
6. Run $\mathsf{Transform}_y[\mathsf{A} \to \mathsf{B}]$.
7. Output $\mathcal{A}$ and $(1, c)$.

Figure 9: The Estimate-Disturb-Transform Procedure

- *The output* $(q, |\psi\rangle) \leftarrow \mathsf{VarEstimate}(|v_{j,1}\rangle)$ *is such that* $|\psi\rangle = |w_{j,1}\rangle$ *with probability* $1 - \delta$.
- *The running time of* $\mathsf{Transform}(q, |\psi'\rangle)$, *where* $(q, |\psi\rangle) \leftarrow \mathsf{VarEstimate}(|v_{j,1}\rangle)$ *and* $|\psi'\rangle$ *is any state, is* $O(\log(1/\delta)/s_j)$ *with probability* $1 - \delta$ *and at most* $1/\delta$ *with probability* 1.
- *The output state of* $\mathsf{Transform}(\mathsf{VarEstimate}(|v_{j,1}\rangle))$ *is* $|w_{j,1}\rangle$ *with probability* $1 - \delta$.

In particular, we will use the "estimate-disturb-transform" $\mathsf{EDT}[\mathsf{A},\mathsf{B},\mathsf{C}]$ procedure from [LMS22] which is described in Fig. 9. From [LMS22], this procedure is EQPT_m.

Fact 4.2 ([LMS22], Claim 13.4)**.** $\mathsf{EDT}[\mathsf{A},\mathsf{B},\mathsf{C}]$ *is* EQPT_m *for any efficient measurements* $\mathsf{A},\mathsf{B},\mathsf{C}$.

Specifically, for a fixed $\tilde{V}$ we will instantiate EDT with the following three projective measurements

$$\mathsf{M}_{\mathsf{Sim}} = \{\Pi_{\mathsf{Sim}}, I - \Pi_{\mathsf{Sim}}\}, \quad \mathsf{M}_{\mathsf{Com}} = \{\Pi_{\mathsf{Com}}, I - \Pi_{\mathsf{Com}}\}, \quad \mathsf{M}_{\mathcal{R}} = \{\Pi_r\}_r$$

Let $\widehat{\mathsf{EDT}}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_{\mathsf{Sim}}, \mathsf{M}_{\mathcal{R}}]$ be a dilation of $\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_{\mathsf{Sim}}, \mathsf{M}_{\mathcal{R}}]$. Then we define our simulator as follows.

$\mathsf{Sim}^{\tilde{V}}(\mathsf{Hyb}_0)$

1. Run $\tilde{V}$ to initialize the registers $\mathcal{C}, \mathcal{D}, \mathcal{R}$, and $\mathcal{V}$.
2. Prepare the state $|\mathsf{Sim}_0\rangle_{\mathcal{A},\mathcal{Z},\mathcal{P}} |0\rangle_{\mathcal{S}}$.
3. Apply $\widehat{\mathsf{EDT}}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_{\mathsf{Sim}}, \mathsf{M}_{\mathcal{R}}]$ to obtain (b, r) on $\mathcal{B}, \mathcal{R}'$ (in superposition). Measure $\mathcal{B}$ to get b.
4. If $b = 1$, measure $\mathcal{R}'$ to get r and replace the state on $\mathcal{A}, \mathcal{Z}, \mathcal{P}, \mathcal{S}$ with $|\mathsf{Sim}_r\rangle_{\mathcal{A},\mathcal{Z},\mathcal{P}} |r\rangle_{\mathcal{S}}$.
5. Apply $\widehat{\mathsf{EDT}}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_{\mathsf{Sim}}, \mathsf{M}_{\mathcal{R}}]^\dagger$.
6. Apply $U_{\tilde{V}}$.

7. Measure $\mathcal{C} \otimes \mathcal{D} \otimes \mathcal{R}$ with $\{V_{\mathsf{Com}}, I - V_{\mathsf{Com}}\}$. If the outcome is 0, abort and output $(\mathcal{V}, \mathcal{A})$.
8. Otherwise, send $\mathcal{Z}$ to $\tilde{V}$ and output $(\mathcal{V}, \mathcal{A}, \mathcal{Z})$.

Lemma 4.3. $\mathit{Sim}^{\tilde{V}} \approx_c \mathrm{View}_{\tilde{V}}(P, \tilde{V})$ *for all efficiently approximable dilations* $\widehat{\mathsf{EDT}}$ *of* EDT.

Proof. Assume there exists an efficiently approximable dilation $\widehat{\mathsf{EDT}}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_{\mathsf{Sim}}, \mathsf{M}_{\mathcal{R}}]$ of the EQPT_m procedure $\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_{\mathsf{Sim}}, \mathsf{M}_{\mathcal{R}}]$, a QPT Adv, and an inverse polynomial $\varepsilon = \varepsilon(\lambda)$ such that

$$\left|\Pr\left[\mathsf{Adv}(\mathsf{Sim}^{\tilde{V}}) = 1\right] - \Pr\left[\mathsf{Adv}(\mathrm{View}_{\tilde{V}}(P, \tilde{V})) = 1\right]\right| \geq \varepsilon$$

where $\mathsf{Sim}^{\tilde{V}}$ uses $\widehat{\mathsf{EDT}}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_{\mathsf{Sim}}, \mathsf{M}_{\mathcal{R}}]$ in Steps 3 and 5. Then we define hybrids Hyb_0 through Hyb_7, such that $\mathsf{Hyb}_0 = \mathsf{Sim}^{\tilde{V}}$ and $\mathsf{Hyb}_7 = \mathrm{View}_{\tilde{V}}(P, \tilde{V})$ and, by an averaging argument, there must exist Hyb_i such that

$$\left|\Pr[\mathsf{Adv}(\mathsf{Hyb}_i) = 1] - \Pr\left[\mathsf{Adv}(\mathsf{Hyb}_{i+1}) = 1\right]\right| \geq \frac{\varepsilon}{7}$$

Our goal will be to prove that between each hybrid, Adv's distinguishing advantage is $< \varepsilon/7$ and so arrive at a contradiction.

Because $\widehat{\mathsf{EDT}}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_{\mathsf{Sim}}, \mathsf{M}_{\mathcal{R}}]$ is an efficiently approximable dilation, there exists an approximation circuit C_γ where $\gamma = 8/\varepsilon$ such that $|C_\gamma| \leq \mathsf{poly}(8/\varepsilon, \lambda) = \mathsf{poly}(\lambda)$. Using C_γ, we construct Hyb_1.

Hyb_1

1. Run $\tilde{V}$ to initialize the registers $\mathcal{C}, \mathcal{D}, \mathcal{R}$, and $\mathcal{V}$.
2. Prepare the state $|\mathsf{Sim}_0\rangle_{\mathcal{A},\mathcal{Z},\mathcal{P}} |0\rangle_{\mathcal{S}}$.
3. Apply $C_\gamma[\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_{\mathsf{Sim}}, \mathsf{M}_{\mathcal{R}}]]$ to obtain (b, r) on $\mathcal{B}, \mathcal{R}'$ (in superposition). Measure $\mathcal{B}$ to get b.
4. If $b = 1$, measure $\mathcal{R}'$ to get r and replace the state on $\mathcal{A}, \mathcal{Z}, \mathcal{P}, \mathcal{S}$ with $|\mathsf{Sim}_r\rangle_{\mathcal{A},\mathcal{Z},\mathcal{P}} |r\rangle_{\mathcal{S}}$.
5. Apply $C_\gamma[\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_{\mathsf{Sim}}, \mathsf{M}_{\mathcal{R}}]]^\dagger$.
6. Apply $U_{\tilde{V}}$.
7. Measure $\mathcal{C} \otimes \mathcal{D} \otimes \mathcal{R}$ with $\{V_{\mathsf{Com}}, I - V_{\mathsf{Com}}\}$. If the outcome is 0, abort and output $(\mathcal{V}, \mathcal{A})$.
8. Otherwise, send $\mathcal{Z}$ to $\tilde{V}$ and output $(\mathcal{V}, \mathcal{A}, \mathcal{Z})$.

Claim 4.4. $|\Pr[\mathsf{Adv}(\mathsf{Hyb}_0) = 1] - \Pr[\mathsf{Adv}(\mathsf{Hyb}_1) = 1]| < \varepsilon/7$.

Proof. Let K be the quantum channel between $\widehat{\mathsf{EDT}}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_{\mathsf{Sim}}, \mathsf{M}_{\mathcal{R}}]$ and $\widehat{\mathsf{EDT}}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_{\mathsf{Sim}}, \mathsf{M}_{\mathcal{R}}]^\dagger$ in Steps 3 and 4 of $\mathsf{Hyb}_0 = \mathsf{Sim}^{\tilde{V}}$. By Definition 3.10,

$$\|\Phi(\widehat{\mathsf{EDT}}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_{\mathsf{Sim}}, \mathsf{M}_{\mathcal{R}}], K) - \Phi(C_\gamma[\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_{\mathsf{Sim}}, \mathsf{M}_{\mathcal{R}}]], K)\|_\diamond \leq \gamma^{-1} = \varepsilon/8$$

where Φ is defined as in the definition. Therefore the state in Step 6 will be $(\varepsilon/8)$-close between Hyb_0 and Hyb_1 in trace distance. Trace distance contracts under quantum channels, so the state output by Hyb_0 and Hyb_1 are also $(\varepsilon/8)$-close and no measurement can distinguish between them with probability $\geq \varepsilon/8$. □

Let $\mathsf{M}_P = \{|P\rangle\langle P|, I - |P\rangle\langle P|\}$ be the measurement projecting onto the honest prover's initial state $|P\rangle$. Then we construct Hyb_2 as follows.

Hyb_2

1. Run $\tilde{V}$ to initialize the registers $\mathcal{C}, \mathcal{D}, \mathcal{R}$, and $\mathcal{V}$.

2. Prepare the state $|P\rangle_{\mathcal{A},\mathcal{Z},\mathcal{P}}$.
3. Apply $C_\gamma[\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_P, \mathsf{M}_{\mathcal{R}}]]^\dagger$ to obtain (b, r) on $\mathcal{B}, \mathcal{R}'$ (in superposition). Measure $\mathcal{B}$ to get b.

4. If $b = 1$, measure $\mathcal{R}'$ to get r and replace the state on $\mathcal{A}, \mathcal{Z}, \mathcal{P}, \mathcal{S}$ with $|\mathsf{Sim}_r\rangle_{\mathcal{A},\mathcal{Z},\mathcal{P}} |r\rangle_{\mathcal{S}}$.
5. Apply $C_\gamma[\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_{\mathsf{Sim}}, \mathsf{M}_{\mathcal{R}}]]^\dagger$.
6. Apply $U_{\tilde{V}}$.
7. Measure $\mathcal{C} \otimes \mathcal{D} \otimes \mathcal{R}$ with $\{V_{\mathsf{Com}}, I - V_{\mathsf{Com}}\}$. If the outcome is 0, abort and output $(\mathcal{V}, \mathcal{A})$.
8. Otherwise, send $\mathcal{Z}$ to $\tilde{V}$ and output $(\mathcal{V}, \mathcal{A}, \mathcal{Z})$.

Claim 4.5. $|\Pr[\mathsf{Adv}(\mathsf{Hyb}_1) = 1] - \Pr[\mathsf{Adv}(\mathsf{Hyb}_2) = 1]| < \varepsilon/7$.

Proof. Towards contradiction, assume this is not true. Then we can construct an adversary Adv' which can break the SHVZK of (P_Σ, V_Σ). The SHVZK guarantee implies that the states $|P\rangle$ and $|\mathsf{Sim}_0\rangle$ are indistinguishable to any QPT adversary without access to $\mathcal{P} \otimes \mathcal{Z}$[4]. Observe that Hyb_1 and Hyb_2 do access $\mathcal{P} \otimes \mathcal{Z}$ but only in a limited way. In particular, EDT acts on these registers via $\mathsf{M}_{\mathsf{Sim}}$ and M_P. By an oracle indistinguishability result Lemma 2.2 from [Yan23], it suffices to prove that Adv' cannot distinguish $|P\rangle$ and $|\mathsf{Sim}_r\rangle$ without access to $\mathcal{P} \otimes \mathcal{Z}$ even in the presence of a projector onto the entire state.

Let $|\phi\rangle \in \{|P\rangle, |\mathsf{Sim}_0\rangle\}$ denote the state being distinguished and $\mathsf{M}_\phi = \{|\phi\rangle\langle\phi|, I - |\phi\rangle\langle\phi|\}$ be the measurement induced by the oracle projector. Now we construct our SHVZK adversary $\mathsf{Adv}'^{\mathsf{M}_\phi}(|\phi\rangle)$ as follows.

[4]It suffices for this proof that only the initial message on $\mathcal{A}$ is well simulated.

1. Run $\tilde{V}$ to initialize the registers $\mathcal{C}, \mathcal{D}, \mathcal{R}$, and $\mathcal{V}$.
2. Prepare the state $|\phi\rangle |0\rangle_{\mathcal{S}}$ on $\mathcal{A}, \mathcal{Z}, \mathcal{P}, \mathcal{S}$.
3. Apply $C_\gamma[\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_\phi, \mathsf{M}_{\mathcal{R}}]]$ to obtain (b, r) on $\mathcal{B}, \mathcal{R}'$ (in superposition). Measure $\mathcal{B}$ to get b.
4. Continue with Steps 4 – 8 of $\mathsf{Sim}^{\tilde{V}}$.
5. Run Adv and output its decision.

Note that, because both C_γ and EDT are black-box in the parameterized measurements, the approximation circuit $C_\gamma[\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, M_\phi, \mathsf{M}_{\mathcal{R}}]]$ can be instantiated with the oracle projector $|\phi\rangle\langle\phi|$.

We claim that the state provided to Adv is exactly Hyb_1 and Hyb_2 when $|\phi\rangle$ is $|\mathsf{Sim}_0\rangle$ and $|P\rangle$, respectively. In the case where $|\phi\rangle = |\mathsf{Sim}_0\rangle$, M_ϕ will measure according to the projector $|\mathsf{Sim}_0\rangle\langle\mathsf{Sim}_0|$. Observe that $\mathcal{S}$ is only measured with $\mathsf{M}_{\mathsf{Sim}}$, so $|\mathsf{Sim}_0\rangle\langle\mathsf{Sim}_0|$ is equivalent to Π_{Sim} when $\mathcal{S}$ is in the $|0\rangle$ state (as is the case in Hyb_1). Therefore, in this case $C_\gamma[\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_\phi, \mathsf{M}_{\mathcal{R}}]] = C_\gamma[\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_{\mathsf{Sim}}, \mathsf{M}_{\mathcal{R}}]]$ and so Adv' is exactly preparing Hyb_1. In the case where $|\phi\rangle = |P\rangle$, we simply have that $\mathsf{M}_\phi = \mathsf{M}_P$ and so Adv' is preparing Hyb_2.

By assumption, Adv can distinguish these two inputs w.p. $\geq \varepsilon/7$ and so Adv' can distinguish $|\mathsf{Sim}_0\rangle$ and $|P\rangle$ with the same advantage. Because Adv and Adv' are QPT (recall that $|C_\gamma| = \mathsf{poly}(\lambda)$), this is a contradiction and the claim holds. □

Let $\mathsf{M}_{P_r} = \{|P_r\rangle\langle P_r|, I - |P_r\rangle\langle P_r|\}$ be the measurement projecting onto the honest prover's final state $|P_r\rangle$ (conditioned on a particular challenge r). Then we construct Hyb_3 as follows.

Hyb_3

1. Run $\tilde{V}$ to initialize the registers $\mathcal{C}, \mathcal{D}, \mathcal{R}$, and $\mathcal{V}$.
2. Prepare the state $|P\rangle_{\mathcal{A},\mathcal{Z},\mathcal{P}}$.
3. Apply $C_\gamma[\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_P, \mathsf{M}_{\mathcal{R}}]]^\dagger$ to obtain (b, r) on $\mathcal{B}, \mathcal{R}'$ (in superposition). Measure $\mathcal{B}$ to get b.

4. If $b = 1$, measure $\mathcal{R}'$ to get r and replace the state on $\mathcal{A}, \mathcal{Z}, \mathcal{P}$ with $|P_r\rangle$.
5. Apply $C_\gamma[\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_{P_r}, \mathsf{M}_{\mathcal{R}}]]^\dagger$.

6. Apply $U_{\tilde{V}}$.
7. Measure $\mathcal{C} \otimes \mathcal{D} \otimes \mathcal{R}$ with $\{V_{\mathsf{Com}}, I - V_{\mathsf{Com}}\}$. If the outcome is 0, abort and output $(\mathcal{V}, \mathcal{A})$.
8. Otherwise, send $\mathcal{Z}$ to $\tilde{V}$ and output $(\mathcal{V}, \mathcal{A}, \mathcal{Z})$.

Claim 4.6. $|\Pr[\mathsf{Adv}(\mathsf{Hyb}_2) = 1] - \Pr[\mathsf{Adv}(\mathsf{Hyb}_3) = 1]| < \varepsilon/7$.

Proof. An identical argument as in Claim 4.5, but reducing to the indistinguishability of $|\mathsf{Sim}_r\rangle$ and $|P_r\rangle$ (without access to $\mathcal{P}$) for all r. □

Hyb_4

1. Run $\tilde{V}$ to initialize the registers $\mathcal{C}, \mathcal{D}, \mathcal{R}$, and $\mathcal{V}$.
2. Prepare the state $|P\rangle_{\mathcal{A},\mathcal{Z},\mathcal{P}}$.
3. Apply $C_\gamma[\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_P, \mathsf{M}_{\mathcal{R}}]]^\dagger$ to obtain (b, r) on $\mathcal{B}, \mathcal{R}'$ (in superposition). Measure $\mathcal{B}$ to get b.

4. If $b = 1$, measure $\mathcal{R}$ to get r.
5. Apply $C_\gamma[\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_P, \mathsf{M}_{\mathcal{R}}]]^\dagger$ and then U_{P_r}.

6. Apply $U_{\tilde{V}}$.
7. Measure $\mathcal{C} \otimes \mathcal{D} \otimes \mathcal{R}$ with $\{V_{\mathsf{Com}}, I - V_{\mathsf{Com}}\}$. If the outcome is 0, abort and output $(\mathcal{V}, \mathcal{A})$.
8. Otherwise, send $\mathcal{Z}$ to $\tilde{V}$ and output $(\mathcal{V}, \mathcal{A}, \mathcal{Z})$.

Claim 4.7. $|\Pr[\mathsf{Adv}(\mathsf{Hyb}_3) = 1] - \Pr[\mathsf{Adv}(\mathsf{Hyb}_4) = 1]| < \varepsilon/7$.

Proof. If Hyb_3 instead used $\widehat{\mathsf{EDT}}$ in Steps 3 and 5, the state ρ before Step 4 will satisfy $\mathrm{Tr}[|P\rangle\langle P|\,\rho] \geq 1 - \mathsf{negl}(\lambda)$ when $b = 1$ by the correctness of $\mathsf{Transform}$. Recall that $C_\gamma[\mathsf{EDT}]$ is an $(\varepsilon/8)$-approximation of $\widehat{\mathsf{EDT}}$. Therefore the state in Step 4 of Hyb_3 is $(\varepsilon/8 + \mathsf{negl})$-close in trace distance to $|P\rangle\langle P|$. Because $|P_r\rangle = U_{P_r}|P\rangle$, if we apply the unitary U_{P_r} instead of replacing the state with $|P_r\rangle$ in Step 4 of Hyb_3, the resulting state will be $(\varepsilon/8 + \mathsf{negl})$-close as well.

Because both C_γ and EDT are black-box in the parameterized measurements, the approximation circuit $C_\gamma[\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, M_\phi, \mathsf{M}_{\mathcal{R}}]]$ will only act on the registers $\mathcal{A}, \mathcal{Z}, \mathcal{P}$ through the projectors Π_{Com}, $\Pi_{P_r} = |P_r\rangle\langle P_r|$, and $\Pi_{\mathcal{R}}$. Thus, if we apply U_{P_r} instead of the replacement, then Steps 4 and 5 of Hyb_3 are equivalent to the following circuit.

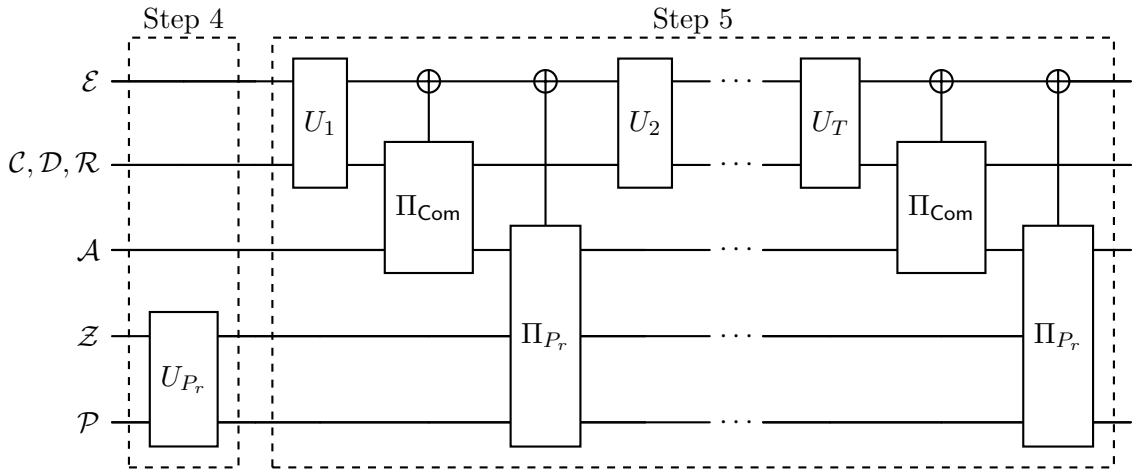

where $(U_1, \cdots, U_T)$ are the other components of $C_\gamma[\mathsf{EDT}]$ which do not act on $\mathcal{P}, \mathcal{A}, \mathcal{Z}$. If we insert $U_{P_r}^\dagger U_{P_r} = I$ between each invocation of Π_{P_r}, we get the equivalent circuit

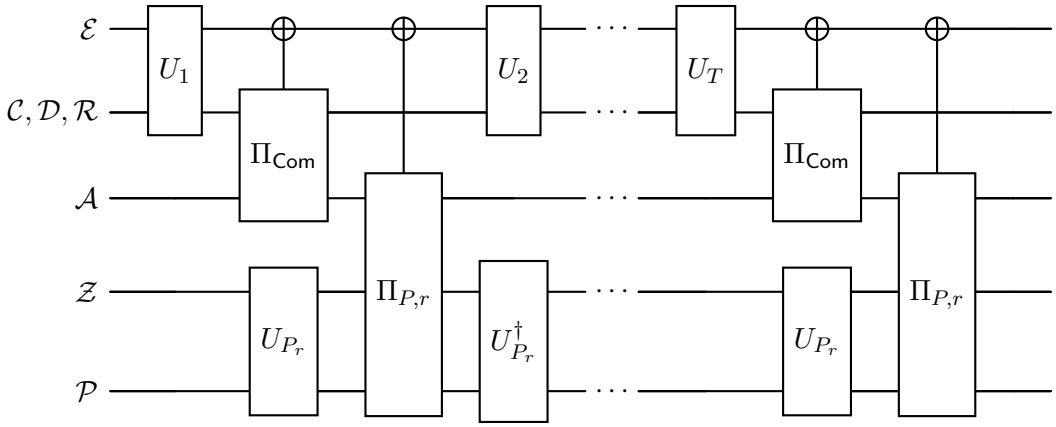

Note that $\Pi_{P_r} = |P_r\rangle\langle P_r| = (I_\mathcal{A} \otimes U_{P_r})\,|P\rangle\langle P|\,(I_\mathcal{A} \otimes U_{P_r})^\dagger = (I_\mathcal{A} \otimes U_{P_r})\Pi_P(I_\mathcal{A} \otimes U_{P_r})^\dagger$. In this way, our previous circuit is equivalent to the following one:

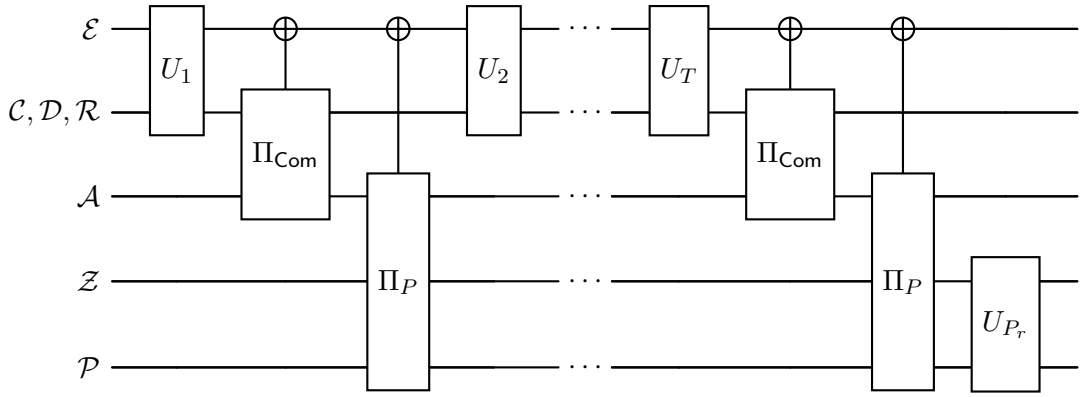

Thus U_{P_r} conjugates Π_P into Π_{P_r} and we have the identity

$$C_\gamma[\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_{P_r}, \mathsf{M}_{\mathcal{R}}]]^\dagger \cdot U_{P_r} = U_{P_r} \cdot C_\gamma[\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_P, \mathsf{M}_{\mathcal{R}}]]^\dagger.$$

Observe that the last circuit is exactly Step 5 in Hyb_4. Because Hyb_3's state after Step 5 is $(\varepsilon/8 + \mathsf{negl})$-close to the state after this circuit, no measurement can distinguish w.p. $\geq \varepsilon/8 + \mathsf{negl}$. □

Hyb_5

1. Run $\tilde{V}$ to initialize the registers $\mathcal{C}, \mathcal{D}, \mathcal{R}$, and $\mathcal{V}$.
2. Prepare the state $|P\rangle_{\mathcal{A},\mathcal{Z},\mathcal{P}}$.
3. Apply $C_\gamma[\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_P, \mathsf{M}_{\mathcal{R}}]]^\dagger$ to obtain (b, r) on $\mathcal{B}, \mathcal{R}'$ (in superposition). Measure $\mathcal{B}$ to get b.
4. If $b = 1$, measure $\mathcal{R}$ to get r'.

5. Apply $C_\gamma[\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_P, \mathsf{M}_{\mathcal{R}}]]^\dagger$.
6. Apply $U_{\tilde{V}}$ and measure $\mathcal{R}$ to get outcome r.

7. Measure $\mathcal{C} \otimes \mathcal{D} \otimes \mathcal{R}$ with $\{V_{\mathsf{Com}}, I - V_{\mathsf{Com}}\}$. If the outcome is 0, abort and output $(\mathcal{V}, \mathcal{A})$.

8. Otherwise, apply U_{P_r}, send $\mathcal{Z}$ to $\tilde{V}$, and output $(\mathcal{V}, \mathcal{A}, \mathcal{Z})$.

Claim 4.8. $|\Pr[\mathsf{Adv}(\mathsf{Hyb}_4) = 1] - \Pr[\mathsf{Adv}(\mathsf{Hyb}_5) = 1]| < \varepsilon/7$.

Proof. Note that U_{P_r} commutes with both $U_{\tilde{V}}$ and the V_{Com} measurement, so Hyb_4 and Hyb_5 will only differ if $r \neq r'$. It suffices then to prove that with all but negligible probability $r = r'$, which

we do through the binding properties of the commitment. Specifically, we reduce to oracle unique message binding which Com satisfies by Lemma A.12.

Towards contradiction, assume the claim is not true and so $r \neq r'$ w.p. $\geq \varepsilon/7$. Then we construct an adversary Adv' which does the following. After preparing $(\mathcal{C}, \mathcal{D}, \mathcal{R})$ as in Steps 1 – 3 of Hyb_4, it sends those registers to the oracle unique message binding referee. Upon receipt of $(\mathcal{D}, \mathcal{R})$ from the referee, Adv applies $C_\gamma[\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_{P_r}, \mathsf{M}_{\mathcal{R}}]]^\dagger$ (using the oracle to instantiate $\mathsf{M}_{\mathsf{Com}}$)[5] and sends $(\mathcal{D}, \mathcal{R})$ back. By assumption, the referee will measure two distinct outcomes w.p. $\geq \varepsilon/7$. Because Adv' is QPT (recall that $|C_\gamma| = \mathsf{poly}(\lambda)$), this is a contradiction and the claim holds. □

Hyb_6

1. Run $\tilde{V}$ to initialize the registers $\mathcal{C}, \mathcal{D}, \mathcal{R}$, and $\mathcal{V}$.
2. Prepare the state $|P\rangle_{\mathcal{A},\mathcal{Z},\mathcal{P}}$.
3. Apply $C_\gamma[\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_P, \mathsf{M}_{\mathcal{R}}]]^\dagger$ to obtain (b, r) on $\mathcal{B}, \mathcal{R}'$ (in superposition). Measure $\mathcal{B}$ to get b.
4. Skip.
5. Apply $C_\gamma[\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_P, \mathsf{M}_{\mathcal{R}}]]^\dagger$.
6. Apply $U_{\tilde{V}}$ and measure $\mathcal{R}$ to get outcome r.
7. Measure $\mathcal{C} \otimes \mathcal{D} \otimes \mathcal{R}$ with $\{V_{\mathsf{Com}}, I - V_{\mathsf{Com}}\}$. If the outcome is 0, abort and output $(\mathcal{V}, \mathcal{A})$.
8. Otherwise, apply U_{P_r}, send $\mathcal{Z}$ to $\tilde{V}$, and output $(\mathcal{V}, \mathcal{A}, \mathcal{Z})$.

Claim 4.9. $|\Pr[\mathsf{Adv}(\mathsf{Hyb}_5) = 1] - \Pr[\mathsf{Adv}(\mathsf{Hyb}_6) = 1]| < \varepsilon/7$.

Proof. Towards contradiction, assume the claim is not true. Then we construct an adversary Adv' that breaks the oracle collapse binding property of Com as follows.

1. Adv' performs Step 1 – 3 as in Hyb_5. It sends $(\mathcal{C}, \mathcal{D}, \mathcal{R})$ to the referee.
2. After receiving back $(\mathcal{D}, \mathcal{R})$, Adv' performs Step 5 – 8 as in Hyb_5.
3. Adv' sends the residual state to Adv and forwards its output.

By construction, if the referee measures $\mathcal{R}$ then the state sent to Adv is exactly Hyb_5. Otherwise, it is Hyb_6. By assumption, Adv distinguishes these w.p. $\geq \varepsilon/7$, and so Adv' wins the game with the same probability. Because Adv' is QPT (recall that $|C_\gamma| = \mathsf{poly}(\lambda)$), this is a contradiction and the claim holds. □

[5] Note this is possible because C_γ and EDT are both black-box in the parameterized measurements.

$\mathsf{Hyb}_7(\mathrm{View}_{\tilde{V}}(P, \tilde{V}))$

1. Run $\tilde{V}$ to initialize the registers $\mathcal{C}, \mathcal{D}, \mathcal{R}$, and $\mathcal{V}$.
2. Prepare the state $|P\rangle_{\mathcal{A},\mathcal{Z},\mathcal{P}}$.
3. Apply the measurement $\mathsf{M}_{\mathsf{Com}}$.
4. Apply $U_{\tilde{V}}$ and measure $\mathcal{R}$ to get outcome r.
5. Measure $\mathcal{C} \otimes \mathcal{D} \otimes \mathcal{R}$ with $\{V_{\mathsf{Com}}, I - V_{\mathsf{Com}}\}$. If the outcome is 0, abort and output $(\mathcal{V}, \mathcal{A})$.
6. Otherwise, apply U_{P_r}, send $\mathcal{Z}$ to $\tilde{V}$, and output $(\mathcal{V}, \mathcal{A}, \mathcal{Z})$.

Claim 4.10. $|\Pr[\mathsf{Adv}(\mathsf{Hyb}_6) = 1] - \Pr[\mathsf{Adv}(\mathsf{Hyb}_7) = 1]| < \varepsilon/7$.

Proof. Let $|\psi\rangle$ be the state after Step 2 of Hyb_5 which we can decompose to $|\psi\rangle = |\psi_0\rangle + |\psi_1\rangle$ where $\Pi_{\mathsf{Com}} |\psi_0\rangle = 0$ and $\Pi_{\mathsf{Com}} |\psi_1\rangle = |\psi_1\rangle$. Consider the modified Hyb_6 where $C_\gamma[\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_P, \mathsf{M}_{\mathcal{R}}]]$ is replaced with $\widehat{\mathsf{EDT}}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_P, \mathsf{M}_{\mathcal{R}}]$. Then in Step 3, by construction of EDT the measurement on $\mathcal{B}$ will yield $b = 1$ if and only if the M_P and $\mathsf{M}_{\mathsf{Com}}$ measurements both give outcome 1 (performed in that order). Because the state on $\mathcal{A}, \mathcal{Z}, \mathcal{P}$ is set to $|P\rangle$, the M_P measurement will always give outcome 1. Thus b is entirely dependent on the measurement $\mathsf{M}_{\mathsf{Com}}$.

In this way, after Step 3 the state would have the form

$$\begin{aligned}\widehat{\mathsf{EDT}}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_P, \mathsf{M}_{\mathcal{R}}] |\psi\rangle |0\rangle_{\mathcal{B}} |0\rangle_{\mathcal{E}} &= \widehat{\mathsf{EDT}}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_P, \mathsf{M}_{\mathcal{R}}] |\psi_0\rangle |0\rangle_{\mathcal{B}} |0\rangle_{\mathcal{E}} \\ &\quad + \widehat{\mathsf{EDT}}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_P, \mathsf{M}_{\mathcal{R}}] |\psi_1\rangle |0\rangle_{\mathcal{B}} |0\rangle_{\mathcal{E}} \\ &= |\perp\rangle |0\rangle_{\mathcal{B}} |\phi_0\rangle_{\mathcal{E}} + |\widetilde{\psi}_1\rangle |1\rangle_{\mathcal{B}} |\phi_1\rangle_{\mathcal{E}} .\end{aligned}$$

where $\mathcal{E}$ is the dilation's ancilla register. Measuring $\mathcal{B}$ collapses the state to $|\widetilde{\psi}_b\rangle |b\rangle |\phi_b\rangle$ with a normalization factor (where $|\widetilde{\psi}_0\rangle = |\perp\rangle$) and Step 5 reverses $\widehat{\mathsf{EDT}}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_P, \mathsf{M}_{\mathcal{R}}]$ leaving

$$\widehat{\mathsf{EDT}}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_P, \mathsf{M}_{\mathcal{R}}]^\dagger |\widetilde{\psi}_b\rangle |b\rangle_{\mathcal{B}} |\phi_b\rangle_{\mathcal{E}} = |\psi_b\rangle |b\rangle_{\mathcal{B}} |0\rangle_{\mathcal{E}}$$

which is exactly the residual state after measuring directly with $\mathsf{M}_{\mathsf{Com}}$ (with a normalization factor). Because $C_\gamma[\mathsf{EDT}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_P, \mathsf{M}_{\mathcal{R}}]]$ is an $(\varepsilon/8)$-approximation of $\widehat{\mathsf{EDT}}[\mathsf{M}_{\mathsf{Com}}, \mathsf{M}_P, \mathsf{M}_{\mathcal{R}}]$, the state after Step 5 of Hyb_6 will be $(\varepsilon/8)$-close to the state after Step 3 of Hyb_7. Thus no measurement can distinguish Hyb_6 and Hyb_7 w.p. $\geq \varepsilon/8$. □

Recall that $\mathsf{M}_{\mathsf{Com}}$ measures according to the projector $\Pi_{\mathsf{Com}} = U_{\tilde{V}}^\dagger V_{\mathsf{Com}} U_{\tilde{V}}$. In this way, Steps 3 – 5 in Hyb_7 are equivalent to applying $U_{\tilde{V}}$ then sequentially measuring V_{Com}, $\mathcal{R}$ in the standard basis, and V_{Com} again. Note that V_{Com} and $\mathcal{R}$ are both diagonal in the standard basis on $\mathcal{R}$ and so commute. In this way, Hyb_7 is exactly equivalent to the real execution $\langle P, \tilde{V}\rangle$.

By Claims 4.3 – 4.9, there is no hybrid Hyb_i such that

$$\left|\Pr[\mathsf{Adv}(\mathsf{Hyb}_i) = 1] - \Pr\left[\mathsf{Adv}(\mathsf{Hyb}_{i+1}) = 1\right]\right| \geq \frac{\varepsilon}{7}$$

Thus Adv cannot distinguish $\mathsf{Hyb}_0 = \mathsf{Sim}^{\tilde{V}}$ from $\mathsf{Hyb}_7 = \mathrm{View}_{\tilde{V}}(P, \tilde{V})$ w.p. ε. This is a contradiction, so the simulator must be correct for all efficiently approximable dilations of EDT. □

5 State-Preserving Extraction

In [LMS22], two notions of post-quantum extraction for proofs/arguments of knowledge are defined: *guaranteed* and *state-preserving* extraction. Guaranteed extraction requires that a protocol have an EQPT_m extractor $\mathsf{Extract}$ with a particular structure such that the probability that $\tilde{P}$ produces a valid transcript but $\mathsf{Extract}$ fails to output a valid witness is negligible.

Definition 5.1. (P_Σ, V_Σ) is a post-quantum proof of knowledge with *guaranteed extraction* if it has an extractor $\mathsf{Extract}^{P^*}$ of the following form:

1. $\mathsf{Extract}^{P^*}$ first runs the cheating prover P^* to generate a (classical) first message a along with an instance x (in a 4-message protocol, this requires first sampling a random vk and running $P^*(\mathsf{vk})$ to obtain x, a).
2. $\mathsf{Extract}^{P^*}$ runs P^* coherently on the superposition $\sum_{r\in\mathcal{R}} |r\rangle$ of all the challenges to obtain a superposition $\sum_{r,z} \alpha_{r,z} |r, z\rangle$ over challenge-response pairs.
3. $\mathsf{Extract}^{P^*}$ then computes (in superposition) the verifier's decision $V(x, a, r, z)$ and measures it. If the measurement outcome is 0, the extractor aborts.
4. If the measurement outcome is 1, run some quantum procedure $\mathsf{FindWitness}^{P^*}$ that outputs a string w.

We require the following two properties to hold:

- **Correctness (guaranteed extraction):** The probability that the first measurement returns 1 but the output witness w is not a valid witness for x is $\mathsf{negl}(\lambda)$.
- **Efficiency:** For any QPT P^*, the procedure $\mathsf{Extract}^{P^*}$ is EQPT_m.

On the other hand, a state-preserving extractor must be an EQPT_c computation (under the previous definition) and output a transcript and residual prover state which are indistinguishable from those generated through an honest interaction with the malicious prover. The authors demonstrated that with sufficient conditions, guaranteed extraction implies state-preserving extraction. In particular, they proved the reduction under the assumption that the protocol is witness binding in the following sense.

Definition 5.2. (P_Σ, V_Σ) is statistically (resp. computationally) *witness binding* with respect to predicate Pred and function f if for any unbounded (resp. non-uniform QPT) quantum adversary Adv,

$$\left|\Pr\left[\mathsf{Exp}^{\mathsf{Adv}}_{\mathsf{wb}}(0, \mathsf{Pred}, f, \lambda) = 1\right] - \Pr\left[\mathsf{Exp}^{\mathsf{Adv}}_{\mathsf{wb}}(1, \mathsf{Pred}, f, \lambda) = 1\right]\right| \leq \mathsf{negl}(\lambda).$$

$\mathsf{Exp}^{\mathsf{Adv}}_{\mathsf{wb}}(b, \mathsf{Pred}, f, \lambda)$

1. Provide Adv a uniformly sampled vk to get an instance x, a prover message a, and a state on $\mathcal{W} \otimes \mathcal{Y}$.
2. Perform the measurement $\{\Pi, I - \Pi\}$ where

$$\Pi = \sum_{(w,y)\,:\,\mathsf{Pred}(x,\mathsf{vk},a,w,y)=1} |w\rangle\langle w|_{\mathcal{W}} \otimes |y\rangle\langle y|_{\mathcal{Y}}$$

   If the outcome is 0, abort and output 0.
3. If $b = 1$, prepare $|0\rangle$ on $\mathcal{K}$, apply the unitary $U_f\,|w\rangle_{\mathcal{W}}\,|k\rangle_{\mathcal{K}} \mapsto |w\rangle_{\mathcal{W}}\,|k \oplus f(w)\rangle_{\mathcal{K}}$, and measure $\mathcal{K}$.
4. Return $\mathcal{W} \otimes \mathcal{Y}$ to Adv who outputs a bit b'.
5. If $b = b'$, output 1. Otherwise output 0.

Under our revised definition, the [LMS22] reduction no longer holds straightforwardly. In particular, the constructed state-preserving extractor cannot be an EQPT_c simulation on its own. Thankfully, the other two properties ("Extraction Correctness" and "State-Preserving") still hold with a minor modification to the assumption. [LMS22] had assumed that any dilation of EQPT_m had an efficient approximation which was key in reducing to the computational witness-binding of the proof of knowledge. Now we must explicitly require that the state-preserving extractor use a dilation which is efficiently approximable. This leads us to the updated statement for this reduction.

Lemma 5.3 ([LMS22], Lemma 10.3 Updated)**.** *Let (P_Σ, V_Σ) be a post-quantum proof/argument of knowledge with guaranteed extraction which we optionally assume outputs some auxiliary information y in addition to the witness w. Additionally, assume that*

- *The protocol (P_Σ, V_Σ) is statistically (resp. computationally) (Pred, f)-witness binding, and*
- *The output of the guaranteed extractor (w, y) satisfies $\mathsf{Pred}(\mathsf{vk}, x, a, w, y) = 1$ with $1 - \mathsf{negl}(\lambda)$ probability.*

Then there exists an EQPT_m computation M and D-circuit D such that, for all quantum algorithms $\tilde{P}$ (with auxiliary state $|\psi\rangle$) and all (resp. efficiently approximable) dilations $\widehat{M}$ of M,

- **Syntax.** *$D[\widehat{M}]$ outputs a protocol transcript τ, a prover state $|\psi'\rangle$, and a witness w.*
- **Extraction Correctness.** *The probability that τ is an accepting transcript but w is an invalid NP witness is negligible.*
- **State-Preserving.** *The pair $(\tau, |\psi'\rangle)$ is statistically (resp. computationally)s indistinguishable from a transcript-state pair $(\tau^*, |\psi^*\rangle)$ obtained through an honest one-time interaction with $P^*(\cdot, |\psi\rangle)$ (where $|\psi^*\rangle$ is the prover's residual state).*

Remark 5.4. *For [LMS22]'s applications, witness extraction was the only variable-time computation done by the zero-knowledge simulators. In this way, the simulation 'inherited' the EQPT_c*

guarantee from the state-preserving extractor. Although slightly more is required with our updated EQPT_c *definition, most of the work is still accomplished by the "state-preserving" property. The subtlety is in dealing with the extracted witness. Because the witness does not appear in the view of the honest protocol, it cannot be indistinguishable from that view. Thus each simulator's handling of the witness must be analyzed case-by-case to prove* EQPT_c *simulation.*

Proof. Let $\mathsf{Extract}$ be (P_Σ, V_Σ)'s guaranteed extractor. Then we define the following variable-runtime computation.

$M^{\tilde{P}}$

1. Apply the measurement $\{|+\rangle\langle+|, I - |+\rangle\langle+|\}$ to $\mathcal{R}$.
2. If the outcome is 0, output $b = 0, w = \perp$ and abort.
3. Otherwise, run $\mathsf{Extract}^{\tilde{P}}$ controlling $\tilde{P}$ on the challenge register $\mathcal{R}$. Output all classical measurement results obtained during the execution — i.e., the instance x, the messages (vk, a), the bit b indicating the verifier's decision, and the extracted output (w, y).

By assumption, $\mathsf{Extract}^{\tilde{P}}$ is EQPT_m for all $\tilde{P}$ so $M^{\tilde{P}}$ is as well. Fix $\tilde{P}$ and let $\widehat{M}$ be an efficiently approximable dilation of $M^{\tilde{P}}$ where the output is written to a register $\mathcal{W}$. Then we define the following D-circuit D.

$D[\widehat{M}]$

1. Prepare $|+\rangle_{\mathcal{R}} |0\rangle_{\mathcal{W}}$.
2. Apply $\widehat{M}$.
3. Measure the appropriate subregisters of $\mathcal{W}$ to get $(x, \mathsf{vk}, a, b, w)$. Note that y is not measured.
4. Apply $\widehat{M}^\dagger$.
5. Run $\tilde{P}(\mathsf{vk})$ to get (x, a). Measure $\mathcal{R}$ to get r.
6. Challenge $\tilde{P}$ with r to get z.
7. Output $(x, \mathsf{vk}, a, r, z, w)$ and $\mathcal{P}$.

Extraction correctness follows immediately from the guaranteed extractor. If $b = 1$, then after $\widehat{M}$ the $\mathcal{R}$ registers will only be on the support of challenges that the prover (with internal state on $\mathcal{P}$) will successfully answer, so the transcript will be accepting. By correctness of $\mathsf{Extract}$, the probability that $b = 1$ and w is not a valid witness is negligible. Thus $D[\widehat{M}]$ satisfies extraction correctness.

Towards contradiction, assume there exists QPT Adv which can distinguish $(x, \mathsf{vk}, a, r, z)$ and $\mathcal{P}$ output by $D[\widehat{M}]$ from those output by $\langle \tilde{P}, V \rangle$. Then we construct Adv' to break the witness binding assumption as follows.

1. Prepare $|+\rangle_{\mathcal{R}} |0\rangle_{\mathcal{W}}$.
2. Apply $\widehat{M}$ and output $\mathcal{W} \otimes \mathcal{Y}$ to the challenger.
3. On return of $\mathcal{W} \otimes \mathcal{Y}$, measure $\mathcal{W}$ to get (x, vk, a, b) but not (w, y).
4. Proceed with Steps 4 – 7 of $D[\widehat{M}]$.
5. Provide $(x, \mathsf{vk}, a, r, z)$ and $\mathcal{P}$ to Adv and forward its output.

Note that the Pred measurement Π made by $\mathsf{Exp}^{\mathsf{Adv}}_{\mathsf{wb}}$ is $\mathsf{negl}(\lambda)$-close to the identity by the additional assumption that $\mathsf{Extract}$ outputs (w, y) consistent with Pred except with negligible probability. Thus if $\mathcal{W}$ is measured, the transcript and state provided to Adv is $\mathsf{negl}(\lambda)$-close to the output of $D[\widehat{M}]$. Otherwise, we must show it is the output of $\langle \tilde{P}, V \rangle$. The idea follows similarly from the proof of Claim 4.10: Steps 2 – 4 together form a measurement that checks whether the verifier accepts when the prover runs on the challenge $\mathcal{R}$. It then commutes with the measurement of r in Step 5, which makes $D[\widehat{M}]$ the same as an honest execution. Let $\varepsilon = \varepsilon(\lambda)$ be the non-negligible advantage of Adv' in the witness binding game.

Lastly, to make Adv' efficient in the computational case, let C be a universal approximation scheme and C_γ be an approximation where $\gamma = 2/\varepsilon$ such that $C_\gamma[M]$ is a poly-sized circuit which approximates $\widehat{M}$ with only an $\varepsilon/2$ cost. Let Adv'' be the modified Adv' that does this approximation. The output of Adv'' is $(\varepsilon/2)$-close to Adv', so Adv'' wins the witness binding game with advantage $\geq \varepsilon/2$. This is a contradiction.

□

5.1 The [GMW86] GNI Protocol has Statistical EQPT_c Simulation

Theorem 5.5. *The [GMW86] 4-message proof system for graph non-isomorphism is post-quantum statistical zero knowledge with EQPT_c simulation.*

In this section, we prove Theorem 5.5 by showing that [LMS22]'s zero knowledge simulator satisfies our updated definition for EQPT_c simulation, i.e., the output of the simulator is statistically close to the output of the real execution regardless of the dilation we use. Although we cannot entirely rely on Lemma 5.3, only a small amount of additional work is required.

Recall that [GMW86] follows the following framework using a non-interactive instance-dependent commitment and a proof of knowledge.

1. The verifier sends an instance-dependent commitment to a bit b.
2. The prover and verifier engage in a proof of knowledge PoK where the verifier proves knowledge of b.
3. If the prover accepts the PoK, it extracts and sends b' from the verifier's commitment. Oth-

erwise, it sends an aborting message $\perp$.

4. The verifier accepts if $b' = b$.

By [LMS22, Corollary 8.23], there is an instance-dependent commitment and PoK such that the above protocol admits guaranteed extraction. Additionally, the commitment scheme is perfectly binding if the graphs are non-isomorphic and perfectly hiding otherwise. Let $L_{G_0,G_1} = \{H : \exists\,(b,\pi)$ such that $\pi G_b \simeq H\}$ be the language being proven in the PoK. Then for any $H \in L_{G_0,G_1}$, the instance has a partial unique witness in b (given that G_0 and G_1 are non-isomorphic). By [LMS22, Corollary 10.7] and Lemma 5.3, there exists a statistical state-preserving extractor for PoK, as the partial witness b is unique. Given all this, we define the simulator as follows:

[GMW86] Sim

1. Given cheating verifier $\tilde{V}$ with classical first message $(\mathsf{com}, \mathsf{pok}_1)$, run the state-preserving PoK extractor on $\tilde{V}$ (which acts as a PoK cheating prover) to get transcript τ, verifier residual state $|\psi'\rangle$, and witness w.
2. If the output transcript is accepting, then output b from partial witness.
3. Otherwise, send an aborting message.

By state-preserving extraction, the partial transcript excluding the final message and the residual verifier state are statistically indistinguishable from an honest execution. It remains only to show that the final message b' preserves this indistinguishability. This can be shown in a hybrid argument through the following.

Hyb

1. Given cheating verifier $\tilde{V}$ with classical first message $(\mathsf{com}, \mathsf{pok}_1)$, run the state-preserving PoK extractor on $\tilde{V}$ (which acts as a PoK cheating prover).
2. If the PoK transcript is accepting, extract and send b' from the verifier's commitment.
3. Otherwise, it sends an aborting message $\perp$.

Sim and Hyb differ only in the source of the final message and, by extraction correctness, with all but negligible probability $b' = b$ when the transcript is accepting. Thus Sim and Hyb are statistically indistinguishable.

Let Real be the view of an honest execution with V^*. Hyb and Real differ only in the PoK transcript and residual verifier state. By the state-preserving property, these will be statistically indistinguishable, and so Hyb and Real are as well. Thus for any dilation of FindWitness, the simulator Sim is statistically indistinguishable from the real view.

5.2 The [FS90] Protocol has Computational EQPT_c Simulation

Theorem 5.6. *Assuming super-polynomially secure non-interactive commitments, [LMS22]'s instantiation of the [FS90] 4-message argument system for NP is post-quantum zero-knowledge with computational EQPT_c simulation.*

Proof Sketch. In [LMS22], the simulator for the [FS90] protocol inherits its EQPT_c structure from AoK's state-preserving extractor. As noted above, under our new definitions the state-preserving extractor may not be EQPT_c itself. Instead, we must show that the [FS90] simulator as a whole is EQPT_c by virtue of the extractor's state-preserving property along with the witness indistinguishability of dAoK.

Thankfully, this proof is identical to [LMS22, Theorem 12.3]'s proof of zero-knowledge. Recall that to prove the existence of a computational EQPT_c simulation, we need to prove indistinguishability for every unitary dilation $\mathsf{Extract}$ which is efficiently approximable. Through careful inspect of the proof, it can be seen that this property is sufficient for proving [LMS22, Theorem 12.3]. Thus, we have that the simulated output is indistinguishable from the real view of the malicious verifier V^* for all efficiently approximable dilations. □

References

[Bar01] Boaz Barak. How to go beyond the black-box simulation barrier. In *42nd Annual Symposium on Foundations of Computer Science, FOCS 2001, Las Vegas, Nevada, USA, October 14-17, 2001*, pages 106–115. IEEE Computer Society, 2001.

[BBBV97] Charles H. Bennett, Ethan Bernstein, Gilles Brassard, and Umesh V. Vazirani. Strengths and weaknesses of quantum computing. *SIAM J. Comput.*, 26(5):1510–1523, 1997.

[BG20] Anne Broadbent and Alex Bredariol Grilo. Qma-hardness of consistency of local density matrices with applications to quantum zero-knowledge. In *IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2020.

[BL02] Boaz Barak and Yehuda Lindell. Strict polynomial-time in simulation and extraction. In John H. Reif, editor, *Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada*, pages 484–493. ACM, 2002.

[BP12] Nir Bitansky and Omer Paneth. From the impossibility of obfuscation to a new non-black-box simulation technique. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 223–232. IEEE Computer Society, 2012.

[BS20] Nir Bitansky and Omri Shmueli. Post-quantum zero knowledge in constant rounds. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy, editors, *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, pages 269–279. ACM, 2020.

[BV97] Ethan Bernstein and Umesh V. Vazirani. Quantum complexity theory. *SIAM J. Comput.*, 26(5):1411–1473, 1997.

[CCLY22] Nai-Hui Chia, Kai-Min Chung, Qipeng Liu, and Takashi Yamakawa. On the impossibility of post-quantum black-box zero-knowledge in constant round. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 59–67. IEEE, 2022.

[Deu85] David Deutsch. Quantum theory, the church–turing principle and the universal quantum computer. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 400(1818):97–117, 1985.

[DNS98] Cynthia Dwork, Moni Naor, and Amit Sahai. Concurrent zero-knowledge. In Jeffrey Scott Vitter, editor, *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*, pages 409–418. ACM, 1998.

[DS22] Marcel Dall'Agnol and Nicholas Spooner. On the necessity of collapsing. *IACR Cryptol. ePrint Arch.*, 2022:786, 2022.

[FS90] Uriel Feige and Adi Shamir. Witness indistinguishable and witness hiding protocols. In Harriet Ortiz, editor, *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing, May 13-17, 1990, Baltimore, Maryland, USA*, pages 416–426. ACM, 1990.

[GK96] Oded Goldreich and Ariel Kahan. How to construct constant-round zero-knowledge proof systems for np. *Journal of Cryptology*, 9(3):167–189, 1996.

[GMW86] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity and a methodology of cryptographic protocol design. In *27th Annual Symposium on Foundations of Computer Science (sfcs 1986)*, pages 174–187. IEEE Computer Society, 1986.

[GSLW19] András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. In Moses Charikar and Edith Cohen, editors, *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*, pages 193–204. ACM, 2019.

[LMS22] Alex Lombardi, Fermi Ma, and Nicholas Spooner. Post-quantum zero knowledge, revisited or: How to do quantum rewinding undetectably. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 851–859. IEEE, 2022.

[LP98] Noah Linden and Sandu Popescu. The halting problem for quantum computers. *arXiv preprint quant-ph/9806054*, 1998.

[Mye97] John M Myers. Can a universal quantum computer be fully quantum? *Physical Review Letters*, 78(9):1823, 1997.

[Oza98a] Masanao Ozawa. Quantum nondemolition monitoring of universal quantum computers. *Physical Review Letters*, 80(3):631, 1998.

[Oza98b] Masanao Ozawa. Quantum turing machines: local transition, preparation, measurement, and halting. In *Quantum Communication, Computing, and Measurement 2*, pages 241–248. Springer, 1998.

[Sti55] W Forrest Stinespring. Positive functions on c*-algebras. *Proceedings of the american mathematical society*, 6(2):211–216, 1955.

[Wat06] John Watrous. Zero-knowledge against quantum attacks. In Jon M. Kleinberg, editor, *Proceedings of the 38th Annual ACM Symposium on Theory of Computing, Seattle, WA, USA, May 21-23, 2006*, pages 296–305. ACM, 2006.

[Yan22] Jun Yan. General properties of quantum bit commitments. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 628–657. Springer, 2022.

[Yan23] Jun Yan. General non-interactive quantum commitments are compatible with quantum rewinding. *Cryptology ePrint Archive*, 2023.

A Quantum Commitments

Definition A.1. A *canonical quantum bit commitment* scheme is represented by an ensemble of polynomial-time uniformly-generated quantum circuit pairs $\{(Q_0(\lambda), Q_1(\lambda))\}_{\lambda\in\mathbb{N}}$ and proceeds as follows:

- In the *commit* stage, to commit a bit $b \in \{0,1\}$, the sender performs the quantum circuit Q_b on the quantum register pair $(\mathcal{C},\mathcal{D})$ initialized in all $|0\rangle$'s state. Now the state on $\mathcal{C},\mathcal{D}$ is $|\mathsf{Com}_b\rangle := Q_b|0\rangle$, and the sender sends the *commitment register* $\mathcal{C}$ to the receiver. Without loss of generality, we assume that $|\mathsf{Com}_0\rangle, |\mathsf{Com}_1\rangle$ are orthogonal.
- In the *reveal* stage, the sender announces the bit b and sends the *decommitment register* $\mathcal{D}$ to the receiver. The receiver will first perform $Q_b^\dagger$ on the quantum register pair $(\mathcal{C},\mathcal{D})$, and measure the result in the computational basis. The receiver accepts if the result is $|0\rangle$. The measurement corresponds to bit b is thus $\Pi_{\mathsf{Com},b} = Q_b|0\rangle\langle 0|Q_b^\dagger$.

The canonical quantum bit commitment can be used in a bitwise fashion to commit to a classical message $m \in \{0,1\}^\ell$ with arbitrary length. In that case, the projectors in the reveal stage are

$$\Pi_{\mathsf{Com},m} = \left(Q_m|0\rangle\langle 0|Q_m^\dagger\right)_{\mathcal{C},\mathcal{D}},$$
$$V_{\mathsf{Com}} = \sum_{m\in\{0,1\}^\ell} |m\rangle\langle m|_{\mathcal{M}} \otimes \Pi_{\mathsf{Com},m}.$$

where $Q_m = Q_{m_1}\otimes Q_{m_2}\otimes\cdots\otimes Q_{m_\ell}$. As with classical commitments, we have the following definition for a hiding quantum commitment:

Definition A.2. A quantum commitment is statistically (resp. computationally) *hiding* if for any $m \neq m'$, the state $\rho_{\mathcal{C}}$ in the commitment phase: $\rho_{\mathcal{C},\mathcal{D}} \leftarrow \mathrm{Commit}(m)$ is statistically (resp. computationally) indistinguishable from the state $\rho'_{\mathcal{C}}$ in the commitment phase: $(\rho'_{\mathcal{C},\mathcal{D}}) \leftarrow \mathrm{Commit}(m')$.

There are many different notions of binding for quantum commitments. Firstly, we have the following definition of "honest binding" for quantum bit commitments, which says that a bit commitment of 0 is hard to convert into a commitment of 1 without access to the commitment register $\mathcal{C}$.

Definition A.3. A quantum bit commitment is statistically (resp. computationally) *honest binding* if for all auxiliary state $|\psi\rangle$ on $\mathcal{E}$ and unbounded (resp. non-uniform QPT) unitary operation U acting on $\mathcal{D}, \mathcal{E}$,

$$\|\Pi_{\mathsf{Com},1} U \left|\mathsf{Com}_0\right\rangle |\psi\rangle\| \leq \mathsf{negl}(\lambda).$$

It turns out that in some applications, an adversary may have additional restricted access to the commitment register $\mathcal{C}$. We have the following definition of unique message binding and collapse binding, with additional access to an oracle O acting on $\mathcal{C}, \mathcal{D}$ and auxiliary registers $\mathcal{B}_0, \ldots, \mathcal{B}_\ell$, defined as follows:

$$O = \sum_{m \in \{0,1\}^\ell} \Pi_{\mathsf{Com},m} \otimes X^m \otimes I + \left(I - \sum_{m \in \{0,1\}^\ell} \Pi_{\mathsf{Com},m} \right) \otimes I \otimes X.$$

Intuitively, O tells the adversary whether the state on $(\mathcal{C}, \mathcal{D})$ is a valid commitment state, and if so, which commitment it is. For simplicity, from now on, we will not explicitly emphasize whether the binding property of a commitment is statistical or computational.

Definition A.4. A quantum bit commitment is *oracle honest binding* if for all auxiliary state $|\psi\rangle$ on $\mathcal{E}, \mathcal{B}_0, \mathcal{B}_1$ and unitary operation $U_0, U_1, \ldots, U_t$ acting on $\mathcal{D}, \mathcal{E}, \mathcal{B}_0, \mathcal{B}_1$,

$$\|\Pi_{\mathsf{Com},1} U_t O \cdots O U_0 \left|\mathsf{Com}_0\right\rangle |\psi\rangle \| \leq \mathsf{negl}(\lambda),$$

where the oracle O on $\mathcal{C}, \mathcal{D}, \mathcal{B}_0, \mathcal{B}_1$ is defined as

$$O = \Pi_{\mathsf{Com},0} \otimes I \otimes I + \Pi_{\mathsf{Com},1} \otimes X \otimes I + (I - \Pi_{\mathsf{Com},0} - \Pi_{\mathsf{Com},1}) \otimes I \otimes X.$$

Lemma A.5. *Let Com be a canonical quantum bit commitment. Then Com is oracle honest binding if and only if it is honest binding.*

Proof. It's easy to see that oracle honest binding implies honest binding. To prove that honest binding implies oracle honest binding, let Adv be an adversary that wins with non-negligible advantage ε in the oracle honest binding game, which makes t queries to O. In that case, by definition, we have that

$$\|\Pi_{\mathsf{Com},1} U_t O U_{t-1} \cdots O U_0 \left|\mathsf{Com}_0\right\rangle |\psi\rangle \| \geq \sqrt{\varepsilon}.$$

We construct a new adversary Adv', which breaks the honest binding game, as follows:

1. Sample $0 \leq i < j \leq t+1$ uniformly at random. If $i = 0$ set $b = 0$. Otherwise sample $b \sim \{0, 1\}$ uniformly at random.
2. On fresh registers $(\mathcal{C}', \mathcal{D}')$ generate the state $\left|\mathsf{Com}_0\right\rangle$.
3. Apply $U_{i-1} O' \cdots U_1 O' U_0$ to registers $(\mathcal{C}', \mathcal{D}', \mathcal{E}, \mathcal{B}_0, \mathcal{B}_1)$, where O' is the oracle with regard to $(\mathcal{C}', \mathcal{D}', \mathcal{E})$.
4. Apply the measurement $(\Pi_{\mathsf{Com},0}, I - \Pi_{\mathsf{Com},0})$ to $\mathcal{C}', \mathcal{D}'$. If the outcome is 0, abort.

5. If $b = 0$, apply X to $\mathcal{B}_0$.
6. Apply $U_{j-1} \cdots U_i$ to $(\mathcal{D}, \mathcal{E})$.
7. Return $\mathcal{D}$.

Note that the success probability of Adv' is unchanged if we replace $(\mathcal{C}', \mathcal{D}', O')$ with $(\mathcal{C}, \mathcal{D}, O)$. Thus the success probability γ of Adv' is

$$\gamma = \frac{1}{t^2 + 2t + 1} \sum_{i,j,b} \|\Pi_{\mathsf{Com},1} U_j \cdots U_i \Pi_{\mathsf{Com},0} X^b_{\mathcal{B}_0} U_{i-1} O \cdots U_1 O U \left|\mathsf{Com}_0\right\rangle \left|\psi\right\rangle \|^2.$$

It remains to show that $\gamma \geq \varepsilon / \mathsf{poly}(t)$. Let $A = \Pi_{\mathsf{Com},1} \otimes X \otimes I + (I - \Pi_{\mathsf{Com},1}) \otimes I \otimes X = O - \Pi_{\mathsf{Com},0} \otimes I \otimes (I - X)$. Then we have

$$U_t O U_{t-1} \cdots O U_0 = \sum_{i=0}^{t} U_t A \cdots U_{i+1} A U_i (\Pi_{\mathsf{Com},0} \otimes I \otimes (I - X)) U_{i-1} O \cdots U_1 O U_0.$$

For each of the i-th summand above, using $A = I + \Pi_{\mathsf{Com},1} \otimes (X \otimes I - I \otimes X)$, we have

$$\begin{aligned} & U_t A \cdots U_{i+1} A U_i (\Pi_{\mathsf{Com},0} \otimes I \otimes (I - X)) U_{i-1} O \cdots U_1 O U_0 \\ = & \sum_{j=i+1}^{t+1} U_t A \cdots U_{j+1} A U_j (\Pi_{\mathsf{Com},1} \otimes (X \otimes I - I \otimes X) U_{j-1} \cdots U_i (\Pi_{\mathsf{Com},0} \otimes I \otimes (I - X)) U_{i-1} O \cdots U_1 O U_0. \end{aligned}$$

By the triangle inequality and unitary invariance of the norm,

$$\begin{aligned} \sqrt{\varepsilon} &\leq 2 \sum_{i<j} \|\Pi_{\mathsf{Com},1} U_j \cdots U_i (\Pi_{\mathsf{Com},0} \otimes (X - I) \otimes I) U_{i-1} O \cdots U_1 O U_0 \left|\mathsf{Com}_0\right\rangle \left|\psi\right\rangle \| \\ &\leq 2 \sum_{i,j,b} \|\Pi_{\mathsf{Com},1} U_j \cdots U_i \Pi_{\mathsf{Com},0} X^b_{\mathcal{B}_0} U_{i-1} O \cdots U_1 O U_0 \left|\mathsf{Com}_0\right\rangle \left|\psi\right\rangle \| \end{aligned}$$

Applying Cauchy-Schwarz,

$$\varepsilon \leq 4(t^2 + 2t + 1) \sum_{i,j,b} \|\Pi_{\mathsf{Com},1} U_j \cdots U_i \Pi_{\mathsf{Com},0} X^b_{\mathcal{B}_0} U_{i-1} O \cdots U_1 O U_0 \left|\mathsf{Com}_0\right\rangle \left|\psi\right\rangle \|^2 \leq 4\gamma(t^2 + 2t + 1)$$

This proves that $\gamma \geq \varepsilon/(4(t^2 + 2t + 1))$, as desired. □

Definition A.6. A quantum commitment is *unique message binding* if for all Adv,

$$\Pr\left[\mathsf{Exp}^{\mathsf{Adv}}_{\mathsf{uniq}}(\lambda) = 1\right] \leq \mathsf{negl}(\lambda),$$

where $\mathsf{Exp}^{\mathsf{Adv}}_{\mathsf{uniq}}$ is defined as follows.

1. Execute Adv to obtain quantum registers $(\mathcal{M}, \mathcal{C}, \mathcal{D})$.
2. Perform binary measurement $\{V_{\mathsf{Com}}, I - V_{\mathsf{Com}}\}$. If the outcome is 0, abort and output 0. Otherwise, measure $\mathcal{M}$ to obtain message m_1.

3. Return registers $(\mathcal{M}, \mathcal{D})$ to Adv, who can perform some quantum algorithm U on $(\mathcal{M}, \mathcal{D})$. Adv returns $(\mathcal{M}, \mathcal{D})$.

4. Perform binary measurement $\{V_{\mathsf{Com}}, I - V_{\mathsf{Com}}\}$. If the outcome is 0, abort and output 0. Otherwise, measure $\mathcal{M}$ to obtain message m_2. Output 1 if $m_1 \neq m_2$, or 0 otherwise.

A quantum commitment is *oracle unique message binding* if any Adv has only non-negligible advantage in $\mathsf{Exp}^{\mathsf{Adv}}_{\mathsf{uniq}}$ even with additional access to O.

Lemma A.7. *Let Com be a canonical quantum bit commitment. Then Com is honest binding if and only if Com is oracle unique message binding.*

Proof. Firstly, we show that oracle unique message binding implies oracle honest binding. Suppose there is a sequence of unitaries $U_t O U_{t-1} \cdots O U_0$ that breaks oracle honest binding with ε advantage, i.e.,

$$\|\Pi_{\mathsf{Com},1} U_t O U_{t-1} \cdots O U_0 \left|\mathsf{Com}_0\right\rangle \left|\psi\right\rangle\| = \varepsilon.$$

To break the oracle unique message binding property, we send $\left|0\right\rangle \left|\mathsf{Com}_0\right\rangle$ to the challenger. On receiving $\mathcal{D}, \mathcal{M}$ back, we perform $U_t O U_{t-1} \cdots O U_0$ on them and sends $\mathcal{D}, \mathcal{M}$. Then $m_1 = 0$ and $\Pr[m_2 = 1] = \varepsilon^2$.

Conversely, suppose we can break the oracle unique message binding game with ε^2 advantage. Let m_1 be the bit obtained in Step 2 of $\mathsf{Exp}^{\mathsf{Adv}}_{\texttt{orac-uniq}}$, and $U_t O U_{t-1} \cdots O U_0$ be the unitary applied in Step 3. Then

- If $m_1 = 0$, the state after Step 2 is exactly $\left|\mathsf{Com}_0\right\rangle$. Thus

$$\|\Pi_{\mathsf{Com},1} U_t O U_{t-1} \cdots O U_0 \left|\mathsf{Com}_0\right\rangle \left|\psi\right\rangle\| \geq \varepsilon.$$

- If $m_1 = 1$, the state after Step 2 is $\left|\mathsf{Com}_1\right\rangle$. This implies that there exists $\left|\phi\right\rangle$ such that

$$\|\Pi_{\mathsf{Com},0} U_t O U_{t-1} \cdots O U_0 \left|\mathsf{Com}_1\right\rangle \left|\phi\right\rangle\| \geq \varepsilon.$$

  Let the internal state of $U_t O U_{t-1} \cdots O U_0 \left|\mathsf{Com}_1\right\rangle \left|\phi\right\rangle$ be $\left|\psi\right\rangle$. Above is equivalent to

$$\|\Pi_{\mathsf{Com},1} U_0^\dagger O^\dagger U_1 \cdots O^\dagger U_t^\dagger \left|\mathsf{Com}_0\right\rangle \left|\psi\right\rangle\| \geq \varepsilon.$$

In either case, we break the oracle honest binding property. □

Definition A.8. A quantum commitment is *collapse binding* if for all Adv, we have

$$\Pr\Big[\mathsf{Exp}^{\mathsf{Adv}}_{\mathsf{coll}}(\lambda) = 1\Big] \leq \frac{1}{2} + \mathsf{negl}(\lambda),$$

where $\mathsf{Exp}^{\mathsf{Adv}}_{\mathsf{coll}}(\lambda)$ is defined as follows.

1. Execute Adv to obtain quantum registers $(\mathcal{M}, \mathcal{C}, \mathcal{D})$.
2. Perform binary measurement $\{V_{\mathsf{Com}}, I - V_{\mathsf{Com}}\}$. If the outcome is 0, abort and output 0. Otherwise, sample $c \in \{0,1\}$ uniformly at random. If $c = 1$, measure $\mathcal{M}$ in the computational basis.
3. Return registers $(\mathcal{M}, \mathcal{D})$ to Adv, who can perform some quantum algorithm U on $(\mathcal{M}, \mathcal{D})$. Adv returns a bit c'.
4. Output 1 if $c = c'$, or 0 otherwise.

A quantum commitment is *oracle collapse binding* if any Adv has only non-negligible advantage in $\mathsf{Exp}^{\mathsf{Adv}}_{\mathsf{coll}}$ even with additional access to O.

Lemma A.9. *Let Com be a canonical quantum bit commitment. Then Com is oracle collapse binding if and only if Com is honest binding.*

Proof. It is proven in [DS22] that collapse binding implies honest binding, so oracle collapse binding implies honest binding. For the other direction, we will prove that oracle honest binding implies oracle collapse binding. To do this, we need to introduce the notion of sum binding.

Definition A.10. A quantum commitment is sum binding if for all Adv, we have $\Pr\left[\mathsf{Exp}^{\mathsf{Adv}}_{\mathsf{sum}}(\lambda) = 1\right] \leq \frac{1}{2} + \mathsf{negl}(\lambda)$, where $\mathsf{Exp}^{\mathsf{Adv}}_{\mathsf{sum}}(\lambda)$ is defined as follows.

1. Execute Adv to obtain quantum registers $\mathcal{C}$.
2. Sample $m \in \{0,1\}$ uniformly at random. Send m to Adv.
3. Obtain registers $(\mathcal{M}, \mathcal{D})$ from Adv, and perform binary measurement $\{V_{\mathsf{Com}}, I - V_{\mathsf{Com}}\}$. If the outcome is 0, abort and output 0. Measure $\mathcal{M}$ to obtain message m'.
4. Output 1 if $m = m'$, or 0 otherwise.

The following proof that sum binding implies collapse binding is adapted from [DS22]. Note that the proof goes through when both are granted access to the oracle.

Sum binding implies collapse binding:

Let Adv be an adversary that achieves ε advantage in the collapse binding game. Let σ be the state of it after Step 1 (i.e., the state of the registers $\mathcal{C} \otimes \mathcal{D} \otimes \mathcal{M}$ and an internal workspace register $\mathcal{E}$) and let $(D, I - D)$ be the distinguishing binary measurement made by Adv in Step 4.

Using Adv, we construct Adv' which breaks sum binding as follows.

1. Prepare $\sigma \leftarrow \mathsf{Adv}$ and send $\mathcal{C}$ to challenger.
2. Receive m and measure $\mathcal{M}$ in the computational basis to get outcome m'.

3. If $m \neq m'$, measure $\mathcal{M} \otimes \mathcal{D} \otimes \mathcal{E}$ using $(D, I - D)$.

4. Send $\mathcal{M}, \mathcal{D}$ to the challenger.

To prove that honest binding implies sum binding, we need the following lemma from [DS22]:

Lemma A.11. *For projectors P, Q and quantum state ρ it holds that*

$$\mathrm{Tr}(PQP\rho) \geq \frac{1}{4}\left(\mathrm{Tr}(P\rho) + \mathrm{Tr}(Q\rho) - 1\right)^2.$$

The following adversary is proposed in [Yan22].

Honest binding implies sum binding:

Let Adv be an adversary that achieves ε advantage in the sum binding game. Let ρ be the state of the sum binding game after Step 1 (i.e., the state of the registers $\mathcal{C} \otimes \mathcal{D}$ and an internal work register $\mathcal{E}$) and let U_0 (resp. U_1) be the unitary operation of Adv in Step 2 (acting on $\mathcal{D} \otimes \mathcal{E}$) when given $m = 0$ (resp. $m = 1$).

Using Adv, we construct Adv', which breaks honest binding as follows.

1. Get $|\mathsf{Com}_0\rangle = Q_0 |0\rangle$ in $\mathcal{C} \otimes \mathcal{D}$. Prepare $\rho \leftarrow \mathsf{Adv}$ on $\mathcal{C}' \otimes \mathcal{D}' \otimes \mathcal{E}'$.

2. Perform U_0 on $\mathcal{D}', \mathcal{E}'$.

3. Perform $\Pi_{\mathsf{Com},0}$ on $\mathcal{C}', \mathcal{D}'$. If the outcome is 0, abort.

4. Perform $U_1 U_0^\dagger$ on $\mathcal{D}, \mathcal{E}'$. Send $\mathcal{D}$ to the challenger.

Proof. Note that after Step 3, conditioned on the adversary not aborting, the state on $\mathcal{C}', \mathcal{D}'$ is exactly $|\mathsf{Com}_0\rangle$. Thus, we may switch to a hybrid where Step 4 is performed on $\mathcal{D}', \mathcal{E}'$ and the challenger has $\mathcal{C}'$ instead. In that case, let $\Pi_0 = U_0^\dagger \Pi_{\mathsf{Com},0} U_0, \Pi_1 = U_1^\dagger \Pi_{\mathsf{Com},1} U_1$. For an adversary that breaks sum binding,

$$\mathrm{Tr}(\Pi_0 \rho) + \mathrm{Tr}(\Pi_1 \rho) - 1 \geq \varepsilon.$$

By Lemma A.11,

$$\frac{\varepsilon^2}{4} \leq \frac{1}{4}(\mathrm{Tr}(\Pi_0 \rho) + \mathrm{Tr}(\Pi_1 \rho) - 1)^2 \leq \mathrm{Tr}(\Pi_0 \Pi_1 \Pi_0 \rho \Pi_0).$$

Finally, $\mathrm{Tr}(\Pi_0 \Pi_1 \Pi_0 \rho \Pi_0 \Pi_1)$ is exactly the probability that the adversary wins against the honest binding game. Thus, we get the desired result. □

Combining the two adversaries together, we may conclude that honest binding is equivalent to oracle collapse binding. □

Importantly, any quantum commitment that satisfies the seemingly "weakest" honest binding property is automatically both oracle unique message binding and oracle collapse binding.

Lemma A.12. *Assuming that honest binding quantum bit commitment exists. Then there exists quantum commitments that are both oracle unique message binding and oracle collapse binding.*

Proof. By Lemma A.7 and Lemma A.9, any quantum bit commitment with honest binding property is automatically oracle unique message binding and oracle collapse binding. We show that the same holds for quantum (string) commitments. Let $\mathsf{Com}(m) = (\mathsf{Com}'(m_1), \ldots, \mathsf{Com}'(m_\ell))$, where Com' is a canonical quantum bit commitment. We prove that if Com' is oracle unique message binding, so is Com.

Indeed, let Adv be an adversary with ε advantage in $\mathsf{Exp}^{\mathsf{Adv}}_{\mathtt{orac\text{-}uniq}}$ for Com. We construct Adv' which breaks $\mathsf{Exp}^{\mathsf{Adv}'}_{\mathtt{orac\text{-}uniq}}$ for Com' as follows.

1. Adv' executes Adv to obtain quantum registers $(\mathcal{M}, \mathcal{C}, \mathcal{D})$.
2. Adv' measures $\mathcal{M}$ to obtain message m_1. It samples a random bit $i \leftarrow \{0, 1, \ldots, \ell - 1\}$, and sends $\mathcal{M}_i, \mathcal{C}_i, \mathcal{D}_i$ to the challenger.
3. On receiving back $(\mathcal{M}_i, \mathcal{D}_i)$ from the challenger, Adv' sends $\mathcal{D} = (\mathcal{D}_0, \ldots, \mathcal{D}_{\ell-1})$ to Adv. When Adv needs to use the oracle O, Adv' simply uses the tensor product of the oracle O_i together with $\{O_j\}_{j \neq i}$ (since Adv' has C_i for $i \neq j$).
4. On receiving $\mathcal{D}$ from Adv, Adv' sends $\mathcal{D}_i$ to the challenger.

Note that Adv' wins the game with probability at least ε/ℓ, which is non-negligible if ε is non-negligible. Similarly, we can show that if Com' is oracle collapse binding, so is Com:

1. Adv' executes Adv to obtain quantum registers $(\mathcal{M}, \mathcal{C}, \mathcal{D})$.
2. Adv' samples a random bit $i \leftarrow \{0, 1, \ldots, \ell - 1\}$, and sends $\mathcal{M}_i, \mathcal{C}_i, \mathcal{D}_i$ to the challenger.
3. On receiving back $(\mathcal{M}_i, \mathcal{D}_i)$ from the challenger, Adv' measures $\mathcal{M}_0, \ldots, \mathcal{M}_{i-1}$ in the computational basis. It sends $\mathcal{D} = (\mathcal{D}_0, \ldots, \mathcal{D}_{\ell-1})$ to Adv. When Adv needs to use the oracle O, Adv' simply uses the tensor product of the oracle O_i together with $\{O_j\}_{j \neq i}$ (since Adv' has C_i for $i \neq j$).
4. On receiving c from Adv, Adv' sends c to the challenger.

We can see that Adv' wins the game with non-negligible probability if Adv does so, by using a hybrid argument: Let H_i be the experiment where the adversary measures $\mathcal{M}_0, \ldots, \mathcal{M}_{i-1}$. Note that H_0 is exactly the collapse binding experiment where no measurement is performed, and H_n is exactly the collapse binding experiment where measurements are performed. Thus, there must be i where H_i and H_{i+1} can be distinguished with non-negligible advantage.

This suggests that Com is oracle unique message binding if and only if it is oracle collapse binding, as desired. □