



An efficient quantum parallel repetition theorem and applications

John Bostanci¹, Luowen Qian², Nicholas Spooner³, and Henry Yuen¹

¹Columbia University

²Boston University

³University of Warwick & NYU

Abstract

We prove a tight parallel repetition theorem for 3-message computationally-secure quantum interactive protocols between an efficient challenger and an efficient adversary. We also prove under plausible assumptions that the security of 4-message computationally secure protocols does not generally decrease under parallel repetition. These mirror the classical results of Bellare, Impagliazzo, and Naor [BIN97]. Finally, we prove that all quantum argument systems can be generically compiled to an equivalent 3-message argument system, mirroring the transformation for quantum proof systems [KW00, KKMV07].

As immediate applications, we show how to derive hardness amplification theorems for quantum bit commitment schemes (answering a question of Yan [Yan22]), EFI pairs (answering a question of Brakerski, Canetti, and Qian [BCQ23]), public-key quantum money schemes (answering a question of Aaronson and Christiano [AC13]), and quantum zero-knowledge argument systems. We also derive an XOR lemma [Yao82] for quantum predicates as a corollary.

Contents

1	Introduction	2
1.1	Applications of 3-message hardness amplification	3
1.2	Barrier for parallel repetition beyond 3-message protocols	5
1.3	Round compression for quantum argument systems	6
1.4	Related works	7
2	Technical overview	8
2.1	2-message non-uniform parallel repetition	8
2.2	Uniform reduction	12
2.3	Round compression	15
3	Preliminaries	16
3.1	Quantum information	16
3.2	Fidelity and Bures distance	16
3.3	Quantum interactive protocols	17
4	Non-uniform parallel repetition of 3-message protocols	18

5	Uniform parallel repetition of 3-message protocols	24
5.1	Jordan’s lemma and alternating projectors	24
5.2	State transformation for almost-projective measurements	26
5.3	Proof of Lemma 5.2	31
6	Barriers to parallel repetition beyond 3-message protocols	32
6.1	Post-quantum bit commitments	32
6.2	Parallel repetition fails for 4-message quantum interactive protocols	34
7	Round compression for quantum argument systems	38
8	Applications	45
8.1	Strong amplification of quantum bit commitment schemes	45
8.2	Quantum XOR lemma	47
8.3	Security amplification for public-key quantum money	49
8.4	3-message quantum zero knowledge	51
8.5	Simpler construction of commitments from undecodable black holes	53

1 Introduction

In this work we study one of the most fundamental questions in theoretical cryptography: can we transform a “weakly” secure construction of a primitive into one that is “truly” secure? A common strategy for such a transformation is *parallel repetition*: if the adversary’s success probability against the original construction is bounded away from 1, then the adversary’s success probability against the repeated construction should tend to zero with the number of repetitions. In classical cryptography this question is well-studied, beginning with the seminal work of Yao [Yao82, Lev87, GNW11] and leading to a long sequence of works [BIN97, CHS05, Hai09, HPWP10, CL10, PV12, CP15, BHT20]. Hardness amplification is also an essential tool for bootstrapping circuit lower bounds (see [SV08] and the references therein).

Our focus in this work is on hardness amplification for *quantum* cryptographic primitives; in particular we focus on the following class of quantum interactive protocols between an *efficient challenger* (specified as part of the protocol) and an *efficient adversary* indexed by a security parameter λ :

- *3-message*: The adversary sends the first message, the challenger the second, and the adversary the third. After the protocol ends, the challenger decides to accept or reject. All messages may be quantum.
- *Weakly computationally secure*: No efficient ($\text{poly}(\lambda)$ -size) adversary can cause the challenger to accept with probability greater than (say) $1 - \frac{1}{\text{poly}(\lambda)}$.

The security of many quantum cryptographic primitives — including quantum (non-interactive) commitments, quantum money and 3-message quantum arguments — can be naturally formulated in terms of the a quantum 3-message protocol associated with the primitive like above. This quantum protocol is often called a *security game*.

Similarly to the classical setting one would like a generic method for amplifying the security of quantum cryptographic primitives. A natural approach is to repeat the construction in parallel; the

security of the repeated construction usually corresponds to the parallel-repeated security game. Ideally, one would want the adversary’s maximum success probability in a repeated security game to decrease exponentially with the number of repetitions.

Our first result is a *tight* parallel repetition theorem for all 3-message quantum protocols.

Theorem 1.1 (3-message efficient parallel repetition, informal). *Let π be a 3-message γ -computationally secure quantum protocol. Then the k -fold parallel repetition $\pi^{\otimes k}$ is $(\gamma^k + \text{negl}(\lambda))$ -computationally secure.*

We prove [Theorem 1.1](#) by identifying the key high level approach used in proving both the classical Yao’s XOR lemma [[Yao82](#), [Lev87](#), [GNW11](#)] and classical tight 3-message parallel repetition theorem of Canetti, Halevi, and Steiner [[CHS05](#)], and then instantiating this high level approach by designing quantum components that work with an arbitrary quantum adversary. As one would expect, handling quantum protocols and adversaries is much more challenging than classical: (1) the classical reduction involves cloning of the adversary’s internal state during protocol execution, which may be computationally infeasible or even information-theoretically impossible (due to entanglement with the challenger); and (2) the classical analysis relies on conditional distributions which breaks down in the quantum setting due to non-commutativity. To resolve these challenges, we combine techniques from recent works on quantum rewinding [[CMSZ22](#)] and quantum algorithmic techniques such as the quantum singular value transform [[GSLW19](#)], as well as additional new ideas to make them compatible with our setting. We explain these in more detail in [Section 2](#).

We stress that our reduction is *uniform* in the strongest possible sense: if an adversary uses quantum advice $|\text{aux}\rangle$ then the reduction uses quantum advice $|\text{aux}\rangle^{\otimes t}$ for some polynomial t . Furthermore, $t = 1$ is possible for any \mathcal{A} as long as $|\text{aux}\rangle$ is an appropriate eigenstate. See [Remark 2.1](#) for details.

On tightness of the reduction. We remark that $\gamma^k + \text{negl}(\lambda)$ is likely the best general bound that one could hope for. The γ^k term is inherent since if the best attack on the original protocol has success probability γ , then simply running this attack on each repetition independently yields an attack achieving success probability γ^k . The negligible term also cannot be eliminated under reasonable assumptions. In particular, the classical 2-message counterexample by Dodis, Jain, Moran, and Wichs [[DJMW12](#)] generalizes to the post-quantum setting, thus the negligible term is inherent assuming existence of exponentially hard post-quantum extended second-preimage resistant hash functions.

1.1 Applications of 3-message hardness amplification

[Theorem 1.1](#) immediately implies hardness amplification for several quantum cryptographic primitives.

Quantum commitments. Bit commitments are a fundamental cryptographic primitive where a sender can commit to a bit b without revealing it at first (this is the *hiding* property), and later can reveal the bit but without the ability to change the bit (this is the *binding* property). Recently our understanding of commitment schemes in the quantum setting has considerably advanced. In particular, there is a robust existential equivalence between commitments and many quantum cryptographic primitives including EFI pairs, which are pairs of efficient mixed states that can only be inefficiently distinguished [[BCQ23](#)]. Therefore, it is likely that commitments and EFI pairs

play a similar “minimal assumption” role (analogous to one-way functions classically) to quantum cryptography.

An important question that has remained open is whether the computational security of quantum commitments (and friends) can be amplified. In other words, given an arbitrary quantum commitment scheme where either the hiding or binding property holds with *weak* (computational) security, can we generically transform it into another quantum commitment scheme where hiding and binding hold with *strong* security? This question was explicitly raised by Yan [Yan22].

Our parallel repetition theorem for computationally secure protocols directly implies hardness amplification for quantum bit commitments, and thus showing robustness of the existence of commitments from a new angle.

Corollary 1.2 (Hardness amplification for commitments). *There is a quantum commitment scheme but only with computational weak hiding (or binding) security, if and only if there is a strong quantum commitment scheme.*

We argue this as follows: without loss of generality it suffices to consider noninteractive commitment schemes using Yan’s compiler [Yan22]. The binding security of the noninteractive scheme can be formulated in terms of the success probability of any efficient adversary in a 2-message security game; correspondingly the security of the repeated scheme can be formulated in terms of any efficient adversary’s success probability in the parallel repeated security game, which by [Theorem 1.1](#) decays to negligible at an exponential rate. Amplification of hiding can be achieved via flavor-switching [Yan22, GJMZ23, HMY23]. We describe this in detail in [Section 8.1](#). We also show how this can be used to drastically simplify constructing commitments from hardness of decoding black hole radiation, originally proven by Brakerski [Bra23], in [Section 8.5](#).

Quantum Yao’s XOR lemma. By the equivalence of quantum commitments and EFI pairs, we also obtain hardness amplification for EFI pairs, answering an open question of Brakerski, Canetti, and Qian [BCQ23]. In fact, we can even use it to show polarization for EFI pairs ([Corollary 8.8](#)).

Corollary 1.3 (XOR lemma for EFI pairs). *If there exists (an ensemble of) weak EFI pairs (ρ_0, ρ_1) that are statistically far but cannot be distinguished with advantage better than ϵ , then the k -fold XOR of (ρ_0, ρ_1) cannot be distinguished with advantage better than $\epsilon^{k/2} + \text{negl}(\lambda)$. In particular, this gives a (strong) EFI pair if ϵ^k is negligible.*

We point out that from this and leveraging an equivalence between quantum state distinguishing and quantum predicates, we can immediately derive a quantum analogue of Yao’s XOR lemma [Yao82], which states that weak computational unpredictability of Boolean predicates (over some distribution of inputs) is amplified when the results of several independent instances are XOR-ed together. A quantum predicate can be defined as two orthogonal average-case inputs ρ_+ (YES), ρ_- (NO) with $\rho_+\rho_- = 0$, and the goal of the predictor is to correctly predict the sign with advantage ϵ . This question was previously asked by Brakerski [Bra23] (private communication) and Colisson [Col19].

Corollary 1.4 (Quantum Yao’s XOR lemma). *The k -fold XOR of an ϵ -unpredictable quantum predicate for ρ_+, ρ_- is $(\epsilon^{k/2} + \text{negl})$ -unpredictable.*

To see a circuit lower bound application of this, we can naturally define “projection complexity classes”, a quantum-input analogue of decision complexity classes. Then we have that for any such class C that is closed under composition with a polynomial fan-in XOR (like the analogue for $PSPACE$), C is strongly hard-on-average against BQP machines if and only if C is weakly hard-on-average against them.

Quantum money. A public-key quantum money scheme consists of quantum states (called *quantum banknotes*) that can be publicly verified by anyone with the public-key, yet remain computationally infeasible to clone. A major goal of quantum cryptography research has been to construct *public-key* quantum money schemes with security based on well-understood assumptions. Aaronson and Christiano [AC13] proved a per-key amplification for a special class of schemes called projective money schemes, and asked whether strong hardness amplification is possible for quantum money schemes. We prove a general amplification that applies to *any* public-key quantum money scheme:

Corollary 1.5 (Hardness amplification for quantum money). *Public-key quantum money schemes satisfying weak uncloneability exist, if and only if there exists a public-key quantum money scheme (satisfying strong unclonability).*

Similar to amplifying commitments, this also follows directly from the observation that the security of a public-key quantum money scheme can be formulated in terms of a 2-message security game, thus it immediately generalizes to e.g. quantum lightning and private-key quantum money. We describe this in detail in Section 8.3.

Amplification of post-quantum security. We remark that, if the original protocol is *classical*, then the repeated protocol is also classical. Hence Theorem 1.1 also implies a parallel repetition theorem for general 3-message *post-quantum* protocols; this was not previously known.

1.2 Barrier for parallel repetition beyond 3-message protocols

We also show that our 3-message parallel repetition theorem (Theorem 1.1) cannot extend to 4-message protocols under reasonable cryptographic assumptions, even if we are restricted to the post-quantum setting. This is a (post-)quantum analogue of the classical result by Bellare, Impagliazzo and Naor [BIN97, Section 3.3]¹.

Theorem 1.6 (Impossibility of parallel repetition, informal). *If there is a post-quantum c -message concurrent-secure many-to-many non-malleable commitment scheme, then for every polynomial k there is a $2c$ -message post-quantum interactive protocol such that the security of a k -fold repetition of the protocol does not decrease compared to the original protocol.*

For the special case of non-interactive commitments ($c = 2$), we would get a 4-message impossibility. We note that while there are no known post-quantum secure non-interactive non-malleable commitments, “pre-quantum” non-interactive non-malleable commitments can be constructed from various subexponential hardness assumptions [KS17, BL18, GKLW21], and so we view this assumption as plausible. Note that a weaker post-quantum *one-to-one* secure constant-round non-malleable commitment scheme is known to exist assuming post-quantum one-way functions [LPY23], and this suffices for a special case of $k = 2$.

We note that classically, stronger impossibilities are known: there is a 4-message protocol whose k -fold computational security cannot be shown to decrease with black-box reductions for any polynomial k [BIN97, Section 3.4], and there is an 8-message protocol whose k -fold computational security is at least constant, regardless of proof techniques [PW12]. These might also generalize to

¹This is essentially the same but one subtle difference is that our counterexample does not require setup. Using setup in a counterexample is arguably problematic as pointed out in [PW12, Section 2.1]. Note that (post-quantum) non-interactive non-malleable commitments with setup can be instantiated from much weaker assumptions, e.g. using a non-malleable encryption or in the (quantum) random oracle model as was done in [BIN97].

the post-quantum setting, assuming strong but reasonable assumptions like post-quantum CCA-secure non-interactive commitments and post-quantum constant-round universal arguments. We consider this sufficient evidence to conjecture that parallel repetition does not amplify 4-message (post-)quantum protocols, but we leave improving the impossibility for future work.

1.3 Round compression for quantum argument systems

An *interactive argument* is a form of interactive proof where the completeness and soundness conditions hold with respect to computationally efficient provers. An important complexity measure of interactive arguments (and interactive proofs in general) is the round complexity. One surprising result in the theory of quantum interactive proofs, due to Kitaev and Watrous [KW00], is that all (single-prover) quantum interactive *proof* systems (where soundness holds against computationally unbounded adversaries) can be compressed to just three messages. We show the analogous statement for quantum interactive arguments via the round compression technique of Kempe, Kobayashi, Matsumoto, and Vidick [KKMV07]. Our technical contribution is to make the reduction efficient.

Theorem 1.7 (Round compression, informal). *Let L be a language with an m -message quantum interactive argument with completeness $1 - c$ and soundness error $s = 1 - \delta$ for $m \geq 3$. Then there exists a 3-message quantum interactive argument for L with completeness $1 - 2c/(m - 1)$ and soundness error $1 - \delta/(m - 1)^4$. The verifier and communication complexity incur only a $\text{poly}(m)$ multiplicative overhead.*

To counteract the worse soundness error, we can again apply [Theorem 1.1](#) to the compressed protocol to obtain a 3-message interactive argument for L with negligible soundness error. Combining these two results, we obtain a general round-preserving soundness amplification theorem for quantum arguments:

Corollary 1.8 (Round-preserving amplification for arguments). *Let L be a language with an m -message quantum interactive argument with completeness $1 - \text{negl}$ (resp., 1) and soundness error $1 - 1/\text{poly}$. Then there exists a $\min\{3, m\}$ -message quantum interactive argument for L with completeness $1 - \text{negl}$ (resp., 1), negligible soundness error, and similar complexity.*

We prove these formally in [Section 7](#). We remark that the crucial aspect of [Theorem 1.7](#) and [Corollary 1.8](#) is that they preserve the communication complexity and the verifier complexity of the original protocol. (Indeed, a trivial round compression for argument systems that is not complexity-preserving can be obtained by having the prover forward its input to the verifier.) We are not aware of any classical analogue of this round compression result.

In [Section 8.4](#), we note that similar techniques allow us to further compile any quantum argument into to a (quantum communication) Σ -protocol [MW05], and thus starting from an honest-verifier zero knowledge protocol, we can get a 3-message malicious-verifier zero knowledge protocol, albeit the soundness becomes worse. We further discuss how to get back negligible soundness at the end of [Section 8.4](#).

Corollary 1.9 (Round compression of zero-knowledge protocols). *For any language L that admits an honest-verifier quantum statistical (resp. computational) zero-knowledge protocol and computational (resp. statistical) soundness, L also admits a malicious-verifier public-coin statistical (resp. computational) zero knowledge protocol with 3 messages, and $1 - 1/\text{poly}$ computational (resp. statistical) soundness, and similar complexity.*

1.4 Related works

Prior works have derived quantum direct product theorems or quantum XOR lemmas in the query-efficient (or communication-efficient) setting [AŠW06, She11, LR13]. Morally these are 2-message (post-quantum) parallel repetitions. However, the query-efficient setting is usually weaker than our time-efficient setting and uses drastically different (non-algorithmic) techniques. For the rest of the discussion we focus on time-efficient hardness amplification.

In [BEM⁺23] a parallel repetition theorem for quantum canonical form commitments was proved, but it only handled classical side information and furthermore only achieved a polynomial rate of decay of the success probability in the repeated protocol. As a consequence, we also improve their Theorem 6.8 such that any inverse polynomial fidelity (or any error that is inverse polynomially bounded away from 1) suffices.

In the classical setting, parallel repetition for three-message arguments (or “weakly-verifiable puzzles”) was studied by [BIN97, CHS05], with the latter showing an optimal exponential soundness amplification. Our three-message parallel repetition result also follows the high-level proof strategy of [CHS05] while borrowing insights from proofs of XOR lemma [Yao82, Lev87, GNW11]. Recent works have observed that in some cases, the [CHS05] amplification result can be applied essentially without modification in the quantum setting. Radian and Sattath [RS19] point out that [CHS05] generalizes to handle 2-message *post-quantum* (classical communication) protocols.

Morimae and Yamakawa [MY22] extend this argument further, adapting [CHS05] to give a parallel repetition theorem for 2-message quantum protocols of the following special form:

0. Both parties a priori agree on a parameter t .
1. The challenger generates a classical verification key k , then uses k to generate t copies of a quantum “puzzle” state $|\text{puz}\rangle$, which it sends to the adversary.
2. The adversary returns a classical answer k' .
3. The challenger accepts or rejects based on k, k' .

They use this result to argue that *weak* one-way state generators (OWSGs) imply OWSGs, analogous to Yao’s amplification of one-way functions. Due to the restriction on the behavior of the challenger — essentially, that its secret state is classical — this result does not suffice for parallel repetition of general 2-message quantum protocols, and does not extend to 3-message protocols even with classical communication. Furthermore, it always requires many copies of the adversary’s auxiliary input whereas our reduction can be advice preserving for eigenstates.

In addition, neither commitments nor quantum money fall within the scope of their result. In the commitment case, this is because both messages in the security game are quantum, and furthermore a general quantum commitment does not have a classical verification key; indeed, the information required to verify the commitment is typically entangled with the state sent to the adversary. For quantum money, the issue is instead that the [MY22] reduction shows only that given an adversary for the parallel repetition of a t -copy protocol, we obtain an adversary for a single repetition of the corresponding t' -copy protocol for some $t' = t \cdot \text{poly}(\lambda)$. This corresponds to giving the adversary multiple copies of the money state, which of course makes the cloning task trivial.

Our reductions share many techniques with prior works in quantum cryptographic reductions, especially in the area of quantum rewinding [Wat09, CCY21, CMSZ22, LMS22]. Like the cited works, we make extensive use of Jordan’s lemma and alternating sequences of projective measurements.

In recent work by Lombardi, Ma, and Spooner [LMS22], they achieved expected polynomial time quantum rewinding, in part by accelerating certain components of [CMSZ22] using the quantum singular value transform (QSVT). In this work, we also make use of the QSVT, but for a quite different purpose: coherent post-selection. Unlike in [LMS22], we crucially rely on the ability of the QSVT to manipulate singular vectors while maintaining coherence between subspaces; see Section 2.1 for more details.

Acknowledgments. We thank Scott Aaronson for bringing to our attention the open question of hardness amplification for quantum money raised by Aaronson and Christiano [AC13]. We thank Ran Canetti for the discussions as well as suggesting the round compression idea. We also thank Xiao Liang and Miranda Christ for the references on post-quantum non-malleability. LQ is supported by DARPA under Agreement No. HR00112020023. JB and HY are supported by AFOSR award FA9550-21-1-0040, NSF CAREER award CCF-2144219, and the Sloan Foundation.

2 Technical overview

2.1 2-message non-uniform parallel repetition

In this section, we give an informal proof sketch for the special case of taking a 2-fold parallel repetition of a 2-message quantum protocol. This special case is easier to understand and cannot be immediately handled by easy changes to [CHS05]. It turns out that the proof for this special case also contains most of the main ideas in the proof for the general non-uniform reduction.

We begin with some notation. A challenger in a 2-message protocol is identified with a pair (V, P) , for V a unitary and P a projector, and an adversary in a 2-message protocol is identified with a pair $(U, |\text{aux}\rangle)$, for U a unitary and $|\text{aux}\rangle$ a quantum input. There are three registers: A, M, C , being the adversary’s register, the message register, and the challenger’s register respectively. We can write the protocol as follows:

- (Challenge) The challenger initializes both M, C to $|0\rangle$, and applies the unitary V to registers MC .
- (Response) The adversary applies some unitary U to registers AM , where A initially contains some “advice” state $|\text{aux}\rangle$.
- (Decision) The challenger applies a projective measurement $\{P, \text{id} - P\}$ to registers MC , and accepts if and only if he gets outcome P .

Without loss of generality, we assume all operations are unitaries or projective measurements since we can expand the private registers C and A appropriately. A 2-fold parallel repetition of (V, P) is simply $(V^{\otimes 2}, P^{\otimes 2})$, acting on registers M_1, M_2, C_1, C_2 . For $i \in \{1, 2\}$, we write V_i to denote the unitary that applies V on registers M_i, C_i ; P_i to denote the projective measurement on registers M_i, C_i .

Suppose (V, P) has computational soundness $\epsilon + \text{negl}$, and we would like to prove that $(V^{\otimes 2}, P^{\otimes 2})$ has computational soundness $\epsilon^2 + \text{negl}$. Assume for the sake of contradiction that there is a 2-fold adversary $(U, |\text{aux}\rangle)$ that achieves an inverse polynomial (for simplicity) advantage over ϵ^2 . That is, the adversary is accepted with probability δ^2 , where $\delta - \epsilon$ is inverse polynomial. Our goal is to construct an 1-fold adversary that is accepted by the original challenger with probability close to δ .

We first give a unified high level approach of the *classical* proof for both tight parallel repetition [CHS05] and the XOR lemma [Yao82] (or Levin’s isolation lemma [Lev87, GNW11]). Later we will extend this high level approach to the quantum setting. The main idea behind all these proofs is similar, we construct a 1-fold adversary by simulating a second challenger with a suitable challenge. Consider the following two cases.

- (i) There exists a fixed challenge c_2 such that running the 2-fold adversary on (c, c_2) outputs a response that is accepted by the first repetition with probability $\geq \delta$.
- (ii) For every challenge in the second repetition, the adversary is accepted by the first repetition with probability $\leq \delta$.

If we are in case (i), then we can construct a non-uniform adversary by giving the 1-fold adversary c_2 as advice. On the other hand, if we are in case (ii), then the 2-fold adversary is accepted by the second repetition with probability $\geq \delta$ whenever it breaks the first repetition. To see why this is the case, let $G_1(c), G_2(c)$ be the events that the adversary is accepted by the first/second repetition on a random challenge c respectively. Then by Bayes’ rule,

$$\delta^2 = \Pr[G_1 \wedge G_2] = \mathbb{E}_{c_1, c_2} [\Pr[G_1] \cdot \Pr[G_2|G_1]] \leq \delta \cdot \mathbb{E}_{c_1, c_2} [\Pr[G_2|G_1]], \quad (1)$$

implying that $\mathbb{E}_{c_2} [\Pr[G_2|G_1]] \geq \delta$. Thus the algorithm for the 1-fold adversary is to simulate the 2-fold protocol, with a real challenger sampling c_1 for the first repetition and the challenge c in the second repetition until the first repetition accepts, and then return the response to the second challenger.

We now attempt to generalize this to the quantum setting. As a first attempt, a natural quantum analogue of case (i) could be the condition

$$\exists |m\rangle, \left\| P_1 U V_1 (|\text{aux}\rangle_A |m\rangle_{M_2} |0\rangle_{M_1 C_1 C_2}) \right\|^2 \geq \delta, \quad (2)$$

which says that there is some message $|m\rangle$ we can insert into the second repetition so the adversary wins the first repetition with probability at least δ . The reduction for this case is straightforward: put the real challenge in M_1 , run the adversary U , then output M_1 ; this succeeds with probability δ by equation (2). We will see soon that case (ii) requires a slightly different condition, but for now we will proceed with equation (2) as stated. Equation (1) suggests the following natural reduction for case (ii):

1. Initialize $|0\rangle_{M_1 C_1}$ and simulate the challenger in the first repetition by running V_1 .
2. Put the real challenge in M_2 .
3. Run the 2-fold adversary U on $AM_1 M_2$.
4. “Post-select” on the event that the challenger accepts in the first repetition (i.e., on P_1).
5. Output M_2 as response.

Before the post-selection step, the state of the system is $U V_1 V_2 |\text{aux}\rangle |0\rangle$. We know that $\|P_1 P_2 U V_1 V_2 |\text{aux}\rangle |0\rangle\|^2 \geq \delta^2$ by assumption (the adversary is accepted with probability $\geq \delta^2$), and that $\|P_1 U V_1 V_2 |\text{aux}\rangle |0\rangle\|^2 < \delta$ by the negation of equation (2). Suppose that we are now able to

post-select on P_1 ; i.e., to prepare the state $\frac{P_1 UV_1 V_2 |\mathbf{aux}\rangle |0\rangle}{\|P_1 UV_1 V_2 |\mathbf{aux}\rangle |0\rangle\|}$. Then as in the classical case, we would be done, since that state achieves success probability

$$\left\| P_2 \cdot \frac{P_1 UV_1 V_2 |\mathbf{aux}\rangle |0\rangle}{\|P_1 UV_1 V_2 |\mathbf{aux}\rangle |0\rangle\|} \right\|^2 = \frac{\|P_1 P_2 UV_1 V_2 |\mathbf{aux}\rangle |0\rangle\|^2}{\|P_1 UV_1 V_2 |\mathbf{aux}\rangle |0\rangle\|^2} \geq \frac{\delta^2}{\delta} = \delta. \quad (3)$$

How do we perform post-selection? Classically, this is achieved by rejection sampling.

As a seasoned reader might expect at this point, naïve rejection sampling does not immediately generalize to the quantum setting. This is because measuring P_1 disturbs M_2 , and it is not in general possible to clone the state on M_2 ; worse, it may be that this state is entangled with the challenger’s private register C_2 . Indeed, for canonical form commitment schemes, M_2 and C_2 are highly entangled, and the challenger will later check for the presence of entanglement.

Attempt: Alternating projectors. Classical rejection sampling can be thought of as a form of *rewinding*. Hence a natural first attempt is to try to apply recent quantum rewinding techniques [Wat09, CCY21, CMSZ22, LMS22]. Following these works, we can implement a form of post-selection without cloning by alternating P_1 (the first repetition accepting) with the projective measurement $Q_1 := (UV_1) |0\rangle\langle 0|_{M_1 C_1} (UV_1)^\dagger$ (the first repetition being initialized correctly) until P_1 accepts.

There are a few issues with this attempt. Alternating projector algorithms can be analyzed via the Jordan (singular value) decomposition of $P_1 Q_1 = \sum_i \varsigma_i |w_i\rangle\langle v_i|$. Before the post-selection step, the state is clearly in Q_1 , and so it can be written as $UV_1 V_2 |\mathbf{aux}\rangle |0\rangle = \sum_i \alpha_i |v_i\rangle$. For simplicity assume for now that we are able to rotate all the singular vectors and the singular values are all non-zero, then the output state of the alternating projectors will be

$$\sum_{i:\varsigma_i>0} \alpha_i |w_i\rangle \otimes |\tilde{\varsigma}_i\rangle,$$

where $|\tilde{\varsigma}_i\rangle$ is the alternating projection history register that only depends on the singular value ς_i (which may be subnormalized). The presence of the history register is problematic since tracing it out amounts to measuring the singular value ς_i . Since this measurement is unlikely to commute with P_2 , we cannot argue that the success probability is at least δ as above. To avoid this problem, we would need to uncompute the history register, which we do not know how to do.

Even if we ignore this issue, and assume we can somehow uncompute the history to obtain the state

$$|\psi\rangle = \sum_{i:\varsigma_i>0} \alpha_i |w_i\rangle,$$

we still would not be able to say that the adversary is accepted with high probability. Recall that our “target” state is

$$\frac{P_1 UV_1 V_2 |\mathbf{aux}\rangle |0\rangle}{\|P_1 UV_1 V_2 |\mathbf{aux}\rangle |0\rangle\|} \approx \sum_i \frac{\varsigma_i}{\sqrt{\delta}} \alpha_i |w_i\rangle. \quad (4)$$

The best bound we can get (via the triangle inequality and equation (3)) is

$$\|P_2 |\psi\rangle\| > \sqrt{\delta} - \left\| P_2 \sum_{i:\varsigma_i>0} \left(1 - \frac{\varsigma_i}{\sqrt{\delta}}\right) \alpha_i |w_i\rangle \right\|,$$

which may be trivial (e.g. if $\alpha_i \approx 1$ for $\varsigma_i \ll \sqrt{\delta}$). Note that this last term can be shown to be non-negative in the classical case, but this could fail quantumly due to the possibility of destructive

interference with respect to P_2 . Therefore, we cannot hope to simply improve the bound on the probability without changing the state $|\psi\rangle$ itself.

Solution: QSVT. To summarize, the alternating projectors approach suffers from two issues: (a) loss of coherence due to explicit computation of ς_i , and (b) incorrect weighting of different singular vectors. To solve both of these issues, we make use of a more sophisticated quantum algorithmic tool, the *quantum singular value transformation* (QSVT) [GSLW19]. Roughly, the QSVT enables efficient, coherent transformations of the form

$$\sum_i \alpha_i |v_i\rangle \rightarrow \sum_i \alpha_i f(\varsigma_i) |w_i\rangle$$

for low-degree real polynomials f with $|f(x)| \leq 1$ when $|x| \leq 1$. We observe that our post-selection task corresponds to $f(\varsigma) = \varsigma/\sqrt{\delta}$. Then Gilyen et al. [GSLW19, Theorem 17] show how to construct a low-degree function g which does satisfy the boundedness conditions, and which approximates $\varsigma/\sqrt{\delta}$ on the range $[0, \sqrt{\delta}]$. Applying the QSVT with respect to this g achieves the necessary post-selection, *provided the spectral norm (maximum singular value) of $P_1 Q_1$ is bounded by $\sqrt{\delta}$* . Furthermore, the reduction goes through as long as the approximation error is $\ll \delta - \epsilon$.

Now we want a promise that all of the singular values of $P_1 Q_1$ are at most $\sqrt{\delta}$ in order to satisfy the necessary boundedness conditions. To achieve this, we simply change the condition for case (ii) to be that the singular values of $P_1 Q_1$ are bounded by $\sqrt{\delta}$, and thus in this case we can safely apply QSVT to approximately post-select. However, we note that the negation of this condition is no longer equation (2), as $P_1 Q_1$ might have a large singular value corresponding to a state that does not come from a state of the form $|\text{aux}\rangle_A \otimes |m\rangle_{M_2}$.

Nevertheless, we can “fix” case 1 by taking advantage of non-uniformity. Suppose that $P_1 Q_1$ has some singular value ς_i larger than $\sqrt{\delta}$, and let $|v_i\rangle_{AM_1 C_1 M_2}$ be a corresponding right singular vector. Since $|v_i\rangle$ is in Q_1 , $U^\dagger |v_i\rangle = (V_1 |0\rangle_{M_1 C_1}) \otimes |\psi\rangle_{AM_2}$ for some advice state $|\psi\rangle$. Then

$$\|P_1 U (V_1 |0\rangle_{M_1 C_1}) |\psi\rangle_{AM_2}\|^2 = \|P_1 |v_i\rangle\|^2 = \varsigma_i^2 \geq \delta.$$

That is, in case (i) the adversary $(U, |\psi\rangle)$ achieves success probability δ , which completes the proof in the non-uniform case. Note that, unlike in the classical case, $|\psi\rangle$ may be entangled across A and M_2 .

Extension to k -fold repetition. In the classical setting for general k , we have k cases as follows. Let G_i be the event that the adversary wins the i -th repetition of the protocol, and suppose that $\Pr[G_k] \geq \delta^k$. It is straightforward to generalize the above to see that there exists some $j \in [k]$ and c_{j+1}, \dots, c_k such that

$$\Pr[G_j | \bigwedge_{i=1}^{j-1} G_i, c_{j+1}, \dots, c_k] \geq \delta,$$

and we can follow the same rejection sampling strategy as above.

In the quantum setting, we similarly generalize the projector Q_1 from the 2-fold case as

$$Q_{\leq j} := UV_1 \dots V_j |0\rangle\langle 0|_{M_{\leq j} C_{\leq j}} (UV_1 \dots V_j)^\dagger,$$

and define $P_{\leq j} := P_1 \dots P_j$. By assumption, $\|P_{\leq k} Q_{\leq k} |\text{aux}\rangle |0\rangle\|^2 \geq \delta^k$, and so in particular the spectral norm of $P_{\leq k} Q_{\leq k}$ is at least $\sqrt{\delta^k}$. It follows that there is some $j \in [k]$ such that

$$\|P_{\leq j} Q_{\leq j}\| \geq \sqrt{\delta^j} \text{ and } \|P_{< j} Q_{< j}\| \leq \sqrt{\delta^{j-1}}. \quad (5)$$

Therefore, given as non-uniform advice a state $|\psi\rangle$ with $\|P_{\leq j} Q_{\leq j} |\psi\rangle\| \geq \sqrt{\delta^j}$, by applying the QSVT with respect to $P_{< j} Q_{< j}$ as in case 2 above we obtain an adversary with success probability δ .

2.2 Uniform reduction

In the previous section, we made crucial use of non-uniformity to provide the adversary with an index j satisfying [equation \(5\)](#) and a vector $|\psi\rangle$ with $\|P_{\leq j}Q_{\leq j}|\psi\rangle\| \geq \sqrt{\delta^j}$. In this section, we will describe how to efficiently prepare $j, |\psi\rangle$ from (polynomially many copies of) the adversary’s initial state $|\text{aux}\rangle$.

We will need to start by relaxing [equation \(5\)](#), as we cannot in general efficiently check the spectral norm of an operator. We address this by observing that our spectral norm condition for post-selection via the QSVT can be substantially weakened: it suffices for the *input state* to have small ($\ll \delta^k$) amplitude on (right) singular vectors $|v_i\rangle$ of $P_{< j}Q_{< j}$ with singular value $\varsigma_i > \sqrt{\delta^{j-1}}$.

Our task then becomes, formally: find an index j and state $|\psi\rangle$ such that (i) $\|P_{\leq j}Q_{\leq j}|\psi\rangle\| \geq \sqrt{\delta^j}$, and (ii) writing $|\psi\rangle = \sum_i \alpha_i |v_i\rangle$ where $\sum_i \varsigma_i |w_i\rangle\langle v_i|$ is the singular value decomposition of $P_{< j}Q_{< j}$, we have $\sum_{i, \varsigma_i > \sqrt{\delta^{j-1}}} |\alpha_i|^2 \ll \delta^k$. This is in fact a quantum analogue of a main algorithmic task in the preprocessing phase of [\[CHS05\]](#). In more detail, the analogous classical task is to find j and a sequence of challenges c_{j+1}, \dots, c_k such that, (i) after fixing challenges c_{j+1}, \dots, c_k in repetitions $j+1, \dots, k$, the residual probability of winning the first j repetitions is at least δ^j , and (ii) with probability $\gg 1 - \delta^k$ over c_j , after fixing c_j, \dots, c_k in repetitions j, \dots, k , the probability of winning the first $j-1$ repetitions is at most δ^{j-1} .

First attempt. Let $|\psi_k\rangle := UV_1 \cdots V_k |\text{aux}\rangle |0\rangle$. Recall that, by assumption, we have that $\|P_{\leq k}Q_{\leq k}|\psi_k\rangle\| \geq \sqrt{\delta^k}$. For each j , let $\sum_i \varsigma_i^{(j)} |w_i^{(j)}\rangle\langle v_i^{(j)}|$ be the singular value decomposition of $P_{\leq j}Q_{\leq j}$.

Let us suppose for now that we have access to the binary projective “singular value threshold” measurement $\Pi^{(j)} = \sum_{i, \varsigma_i^{(j)} > \sqrt{\delta^j}} |v_i^{(j)}\rangle\langle v_i^{(j)}|$, for each j .² We do not know how to realize this measurement efficiently, but it can be approximated in some sense [\[GSLW19, CMSZ22\]](#). This will introduce a number of technical complications that we address later; for now, we assume access to the exact measurement. Observe that we can write condition (ii) equivalently as $\|\Pi^{(j-1)}|\psi\rangle\|^2 \ll \delta^k$.

Our first attempt at a uniform reduction is as follows. We apply $(\Pi^{(k-1)}, I - \Pi^{(k-1)})$ to $t \gg 1/\delta^k$ copies of $|\psi_k\rangle$. If we ever see the outcome $\Pi^{(k-1)}$, the post-measurement state $|\psi_{k-1}\rangle$ is in $\Pi^{(k-1)}$, and so $\|P_{\leq k-1}Q_{\leq k-1}|\psi_{k-1}\rangle\| > \sqrt{\delta^{k-1}}$, and we can then recurse on $|\psi_{k-1}\rangle$. Otherwise, since we never see the outcome $\Pi^{(k-1)}$, with high probability $\|\Pi^{(k-1)}|\psi_k\rangle\|^2 \ll \delta^k$, and so we can output $j = k$ and $|\psi\rangle = |\psi_k\rangle$. Finally, if we get to $|\psi_1\rangle$, we can simply output $j = 1$ and $|\psi\rangle = |\psi_1\rangle$.

Unfortunately, this approach only works for constant k . To see why, notice that to prepare a single copy of $|\psi_{j-1}\rangle$ we may need $1/\delta^k$ copies of $|\psi_j\rangle$. Unlike in the classical setting, we cannot in general clone $|\psi_j\rangle$. Hence the number of copies of $|\psi_k\rangle$ required (and the running time of the algorithm) scales as $\Omega(1/\delta^{k^2})$, which may be superpolynomial for $k = \omega(1)$.

Second attempt. To resolve this issue, we note that in order for the non-uniform reduction to work, it suffices to simply produce j along with *any* state in $\Pi^{(j)}$ with a small enough overlap with $\Pi^{(j-1)}$, therefore in the case we measure $I - \Pi^{(j-1)}$, it suffices to recover a state from $\Pi^{(j)}$ instead of recovering exactly $|\psi_j\rangle$. This is reminiscent of the “state repair” problem encountered in quantum rewinding [\[CMSZ22\]](#); our algorithm will follow that template. In more detail, the reduction works as follows.

²Unlike in the classical setting, it is important here that we do not actually *measure* the singular value, since this would cause too much disturbance.

1. Measure the input state $|\psi_k\rangle_A$ with $\Pi^{(k)}$. If it rejects, start over with a fresh copy of $|\psi_k\rangle$.
2. Repeat for $j = k, \dots, 2$:
 - (a) Measure $\Pi^{(j-1)}$ and $\Pi^{(j)}$ in an alternating fashion for up to $t \gg 1/\delta^k$ iterations.³ If $\Pi^{(j-1)}$ ever accepts, go to the next iteration of the loop ($j - 1$).
 - (b) Otherwise, keep performing alternating projections until $\Pi^{(j)}$ accepts, then output j and A and abort.
3. Output $j = 1$ and register A .

Using Jordan’s lemma, and via similar reasoning to [CMSZ22], it is possible to show that (i) because at the beginning of the j -th loop iteration, the state is in $\Pi^{(j)}$, the number of measurements performed in [step 2b](#) is $O(t)$ in expectation; and (ii) if we never see $\Pi^{(j-1)}$ in [step 2a](#) then with high probability the state $|\psi\rangle$ output by the algorithm on termination satisfies $\|\Pi^{(j-1)}|\psi\rangle\|^2 \ll \delta^k$.

Adapting to approximate POVMs. The algorithm described above is correct assuming access to the projectors $\Pi^{(j)}$. In reality, we can only approximate them using (e.g.) Marriott–Watrous [MW05]. Furthermore, this approximate implementation is not a projection but a POVM; equivalently, it is a projection $\tilde{\Pi}^{(j)}$ acting on the register A and an auxiliary register W_j that is initially set to $|0\rangle$.

Following [CMSZ22], the natural approach to extend the algorithm above to this case is to simply replace $\Pi^{(j-1)}$ and $\Pi^{(j)}$ measurements with their approximate counterparts, $\tilde{\Pi}^{(j-1)} \otimes |0\rangle\langle 0|_{W_j}$ and $\tilde{\Pi}^{(j)} \otimes |0\rangle\langle 0|_{W_{j-1}}$. The projection on the ancilla register for the other measurement aims to ensure its correct initialization.

This approach *almost* works but for a subtle technical issue. Even though W_{j-1} and W_j will be initialized to $|0\rangle$, after applying the first two projections in [step 2a](#), we no longer have any guarantees about the ancilla registers. Therefore, even if we measure that $\tilde{\Pi}$ accepts, it does not imply that we have a state close to Π since it could be that the ancilla registers were malformed.

As a starting point, let us first look at how well the previous algorithm works if we simply plug in $\tilde{\Pi}^{(j)}$ ’s (we omit the zero projector on the ancillas to keep the notations simple). Since the ancillary issue only arises after we perform two projections $\tilde{\Pi}^{(j-1)}$ and $\tilde{\Pi}^{(j)}$, we observe:

1. If $\tilde{\Pi}^{(j-1)}$ accepts in the first iteration, we must still (approximately) have a vector in $\Pi^{(j-1)}$ as the ancilla is initialized to zero at the beginning.
2. Furthermore, the alternating projections can still estimate the singular value. If we, instead of going to $j - 1$ whenever $\Pi^{(j-1)}$ accepts, estimate the singular value and only declare we are in case j when we are below some minuscule threshold, it turns out to still work. This is because as long as the threshold is small enough, when we are below the threshold, by gentle measurement, it must be the case that the auxiliaries are not too far from zero. Thus a small singular vector between $\tilde{\Pi}$ ’s is also a relatively small singular vector between Π ’s.
3. Now it remains to handle the last case where the first measurement rejects but the estimated singular value is still higher than the threshold. The final observation is that in fact the probability that we reach the last case is in fact bounded away from 1 for *any* starting state:

³Technically $\Pi^{(j-1)}$ acts on an additional register C_j . This is a minor point and does not really affect the algorithm nor the analysis. We can simply initialize all C_j ’s to 0 at the beginning and add them to A before aborting.

intuitively if the starting state has a large overlap with $\Pi^{(j-1)}$ then the first clause catches it with noticeable probability, otherwise the second clause catches it with noticeable probability. Therefore, when we reach the last case, we can simply recover *any* state in $\Pi^{(j)}$ again so that we can restart from the beginning. Since the algorithm succeeds for any starting state with some probability, even if in each iteration the starting state is different, we will still eventually reach one of the two good cases with a sufficiently large number of trials.

Leveraging these three observations, we solve this final issue by modifying the loop (step 2) with a more careful algorithm as follows:

- (a) Repeat $t \gg 1/\tau$ times for some inverse polynomial threshold $\tau = \delta^{O(k)}$:
 - (i) Initialize $W_{j,j-1}$ to zero. Measure $(I - \tilde{\Pi}^{(j-1)}) \otimes |0\rangle\langle 0|_{W_j}$. If the measurement rejects, proceed to the next iteration of the outer loop ($j - 1$).
 - (ii) Otherwise, measure $\tilde{\Pi}^{(j)} \otimes |0\rangle\langle 0|_{W_{j-1}}$, $(I - \tilde{\Pi}^{(j-1)}) \otimes |0\rangle\langle 0|_{W_j}$ in an alternating fashion for (say) t^2 iterations. Then, keep alternating until $\tilde{\Pi}^{(j)} \otimes |0\rangle\langle 0|_{W_{j-1}}$ accepts. Let $|\phi\rangle_{A,W_j} |0\rangle_{W_{j-1}}$ be the post-measurement state.
 - (iii) Use the outcomes of the alternating measurements to compute an estimate γ of $\left\| \left((I - \tilde{\Pi}^{(j-1)}) \otimes |0\rangle\langle 0|_{W_j} \right) |\phi\rangle |0\rangle_{W_{j-1}} \right\|^2$. If γ is above $1 - \tau$, terminate the outer loop. Otherwise, proceed to the next iteration of the inner loop.
- (b) Abort without any outputs.

An additional key change is that we are now alternating $\tilde{\Pi}^{(j)} \otimes |0\rangle\langle 0|_{W_{j-1}}$ and $(I - \tilde{\Pi}^{(j-1)}) \otimes |0\rangle\langle 0|_{W_j}$. We also use state repair again to recover a new state for the next iteration. We remark that in order for the algorithm to work we also need to slightly shift the singular value threshold in each iteration, but we refer the readers to the full proof for these technical details.

We now formalize the observations above to analyze this new algorithm. Note that if the *first* application of $(I - \tilde{\Pi}^{(j-1)}) \otimes |0\rangle\langle 0|_{W_j}$ *rejects*, it must be that the post-measurement state is in $\tilde{\Pi}^{(j-1)}$ because W_j is initialized to $|0\rangle$; this is *not true* for subsequent applications because the measurement may have rejected due to a malformed ancilla.

To argue correctness, we consider two cases. The first case is when, in some iteration of the inner loop, the estimate γ is above the threshold $1 - \tau$. In this case we must show that the post-measurement state $\rho = \text{Tr}_{W_j}(|\phi\rangle\langle\phi|)$ on A is (almost completely) in $\Pi^{(j)}$ and has very small overlap with $\Pi^{(j-1)}$. To see this, observe that by gentle measurement the state $|\phi\rangle$ is $\sqrt{\tau}$ -close to a state of the form $|\psi\rangle_A |0\rangle_{W_j}$. The state $|\psi\rangle$ then has the property that $|\psi\rangle |0\rangle_{W_j} |0\rangle_{W_{j-1}}$ is $O(\sqrt{\tau})$ -close to both $\tilde{\Pi}^{(j)}$ and $I - \tilde{\Pi}^{(j-1)}$. Since this latter state has ancillas initialized to zero, it follows that $\tilde{\Pi}^{(j)}$, $\tilde{\Pi}^{(j-1)}$ approximate $\Pi^{(j)}$, $\Pi^{(j-1)}$ on this state, and so $|\psi\rangle$ (which is close to ρ) is close to both $\Pi^{(j)}$ and $I - \Pi^{(j-1)}$.

Otherwise, if γ is always below $1 - \tau$, then in each iteration of the inner loop, we will terminate in step (a)(i) with probability at least τ . It follows that, since $t \gg 1/\tau$, with overwhelming probability the loop will terminate in one of these two cases.

Remark 2.1 (Advice preservation). We note that, while our reduction preserves uniformity, it is not strictly advice-preserving (or constructive [BBK22]), as it requires many copies of the adversary's advice state.

This is inherent for any quantum reduction whose success probability ought to be higher than that of the adversary. Indeed, this is true even classically for randomized advice (and hence also for quantum advice via purification): given an adversary which succeeds with probability δ over the advice distribution, a black-box reduction given only one sample from the advice distribution cannot succeed with probability greater than δ in general.

We remark that the only reason for requiring many copies of the advice is in order to obtain a state in $\Pi^{(k)}$ in [step 1](#). Thus, if the advice state is already in $\Pi^{(k)}$, one copy suffices.

2.3 Round compression

We analyze the soundness of the round compression transformation of Kempe et al. [KKMV07] when applied to argument systems. At a high level, their transformation works by recursively converting an $(r + 1)$ -message protocol into an $(r/2 + 1)$ -message compressed protocol. In an honest execution, the prover begins by simulating the original $(r + 1)$ -message protocol until the $r/2$ -th message, and sends the original uncompressed verifier's private registers to the challenger in the compressed protocol. From there, the verifier flips a coin, deciding whether to continue by running the original protocol forwards or backwards in time.

If the verifier decides to execute the protocol backwards in time, the honest prover and verifier apply the inverse of uncompressed protocol, and at the end the verifier measures whether their private register returns to the state $|0\rangle$. On the other hand, if the verifier decides to execute the protocol forwards, the honest prover and verifier execute the remainder of the uncompressed protocol and the verifier checks the same predicate that the uncompressed verifier does at the end of the uncompressed protocol.

Completeness is straightforward: the honest prover simply simulates the protocol using the original prover and verifier up to the midpoint, and then cooperate with the verifier to compute the protocol either in the forward or the backward direction. To show (computational) soundness, we demonstrate an efficient reduction from an adversary for the compressed protocol to an adversary for the uncompressed protocol. In particular, the adversary for the uncompressed protocol simulates an interaction between the compressed adversary and the compressed verifier, conditioned on the verifier executing the protocol backwards. The adversary can then measure the simulated verifier's register, and conditioned that measurement accepting, the adversary now has a good initial state for the uncompressed protocol, and the state of the simulated verifier's register is $|0\rangle$, so it can be discarded.

From there, the adversary sends their first message and continues by applying the inverse of the compressed adversary until round $r/2$. After round $r/2$, they apply the same unitaries as the compressed adversary, conditioned on the compressed verifier executing the protocol forward in time. Assuming that the compressed adversary was accepted with probability $(1 - \epsilon)$, we show that the state after simulating either the forwards or backwards protocol is $(1 - 4\epsilon)$ -close in squared Bures distance to a state that is accepted by the challenger in both cases. Using the weak triangle inequality for the squared Bures distance, we find that the state of the verifier at the end of the protocol is $(1 - 16\epsilon)$ -close in squared Bures distance to a state that is accepted by the challenger, implying that the challenger accepts with probability $1 - 16\epsilon$. The use of squared Bures distance, instead of the more commonly-used trace distance, avoids a blowup from ϵ to $\sqrt{\epsilon}$ in this step.

This process halves the number of rounds at a cost of mapping $1 - \epsilon$ soundness to $1 - \epsilon/16$. Iterating this protocol $\log m$ times, where m is the number of messages in the original protocol, we arrive at a 3-message protocol with soundness $1 - \epsilon/m^4$.

3 Preliminaries

3.1 Quantum information

A quantum *register* \mathbb{R} is a named finite-dimensional complex Hilbert space. We write $L(\mathbb{R})$ to denote the set of linear transformations on \mathbb{R} and $S(\mathbb{R})$ to denote the set of density matrices on \mathbb{R} (i.e. positive semi-definite and unit-trace operators). For a vector $|\psi\rangle \in \mathbb{R}$, we write ψ to denote the density matrix $|\psi\rangle\langle\psi|$, and for vectors $|\psi\rangle, |\phi\rangle \in \mathbb{R}$, we write $\langle\psi|\phi\rangle$ to denote the inner product. For a vector $|\psi\rangle \in \mathbb{R}$, we write $\| |\psi\rangle \|$ to denote the standard norm over \mathbb{R} , i.e. $\langle\psi|\psi\rangle$. We write $\text{Tr}(\cdot)$ to denote the trace and $\text{Tr}_{\mathbb{R}}$ to denote the partial trace over a register \mathbb{R} .

For a linear operator $X \in L(\mathbb{R})$, let $\|X\|_{\infty}$ be its operator norm and $\|X\|_1 = \text{Tr}(\sqrt{X^\dagger X})$ be its trace norm. For two density matrices $\rho, \sigma \in S(\mathbb{R})$, let $\text{td}(\rho, \sigma) = \frac{1}{2} \|\rho - \sigma\|_1$ be the trace distance between the two. We sometimes write $X_{\mathbb{R}}$ to indicate that X acts on \mathbb{R} . All un-labeled operators act on all registers that do not have an operator acting on them, and if an operator is associated with specific registers, we drop the register subscripts for brevity.

A *binary projective measurement* is a pair $(\Pi, I - \Pi)$, where Π is an orthogonal projector. By convention we refer to the outcome corresponding to Π as 1 and $I - \Pi$ as 0. Since a binary projective measurement is completely specified by Π , we often refer to such a measurement simply as Π .

Definition 3.1 (Eigenspace projectors). *Let $H = \sum_j \lambda_j |j\rangle\langle j|$ be a Hermitian matrix. For $\kappa \in \mathbb{R}$, we denote by $\Pi_{<\kappa}^H := \sum_{j, \lambda_j < \kappa} |j\rangle\langle j|$ the projector on to eigenspaces of H with eigenvalue less than κ . $\Pi_{\geq\kappa}^H$ is defined similarly.*

3.2 Fidelity and Bures distance

An important tool used in the paper will be the quantum fidelity and the related squared Bures distance. Given two quantum states $\rho, \sigma \in S(\mathbb{R})$, the *fidelity* between ρ and σ is given by

$$F(\rho, \sigma) = \text{Tr} \left(\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}} \right)^2.$$

This definition of fidelity might be sometimes be referred to as the “squared” fidelity. The fidelity can be related to the trace distance by a pair of inequalities called the Fuchs-van de Graaf inequalities.

Proposition 3.2 (Fuchs-van de Graaf inequalities). *For all density matrices ρ and σ over the same Hilbert space, we have that*

$$1 - \sqrt{F(\rho, \sigma)} \leq \text{td}(\rho, \sigma) \leq \sqrt{1 - F(\rho, \sigma)}.$$

It is well known that the fidelity is a useful quantity when examining the effects of performing a measurement on a quantum state. Specifically, the gentle measurement lemma gives a bound on the trace distance a state can move after a measurement based on the probability of the measurement accepting.

Proposition 3.3 (Gentle measurement lemma [Win99]). *Given a pure state ρ and a projector Λ , let*

$$\rho' = \frac{\Lambda\rho\Lambda}{\text{Tr}(\Lambda\rho)}$$

be the post-measurement state. Then $F(\rho', \rho) = \text{Tr}(\Lambda\rho)$. It follows from the Fuchs-van de Graaf inequalities that $\text{td}(\rho', \rho) \leq \sqrt{1 - \text{Tr}(\Lambda\rho)}$.

Another way to view the trace of a projector applied to a state ψ is as the fidelity with the closest state in the +1-eigenspace of the projector. Formally we have the following lemma.

Proposition 3.4 (Projector to max fidelity [Wil17, Theorem 9.2.2]). *For any projector Π and state $|\psi\rangle\langle\psi|$,*

$$\text{Tr}(\Pi |\psi\rangle\langle\psi|) = \max_{\text{Tr}(\Pi\sigma)=1} \text{F}(\sigma, |\psi\rangle\langle\psi|)$$

The fidelity corresponds to a squared inner product between states. The *squared Bures distance* is a related distance measure between two states ρ and σ , defined as

$$d_{\text{Bures}}(\rho, \sigma) = 2(1 - \sqrt{\text{F}(\rho, \sigma)}).$$

Being a distance measure, the Bures distance obeys a weak triangle inequality.

Proposition 3.5 (Weak triangle inequality for Bures distance [CS15, Proposition 2.1]). *Let ρ_1, ρ_2, ρ_3 be three quantum states, then*

$$d_{\text{Bures}}(\rho_1, \rho_3) \leq 2(d_{\text{Bures}}(\rho_1, \rho_2) + d_{\text{Bures}}(\rho_2, \rho_3)).$$

3.3 Quantum interactive protocols

A $(2r + 1)$ -message quantum interactive protocol π is specified by a quantum interactive algorithm C (the “challenger”), which interacts with an arbitrary quantum interactive algorithm A (the “adversary”). An execution of π consists of r *interactions*, each one consisting of a (quantum) message from the adversary followed by a (quantum) message from the challenger; and then a *decision* after the adversary sends their final message to the challenger, wherein the challenger either accepts or rejects.

In the following we give a detailed description of an interactive protocol and introduce the notation for registers we use throughout the paper. A visual representation of a quantum interactive protocol can be found in [Figure 1](#).

The adversary in an interactive protocol starts with an initial private register A_0 , and an initial (0-dimensional) message register R_{-1} , and the challenger starts with an initial private register W_0 .⁴ Each interaction in an interactive protocol proceeds (without loss of generality) as follows:

1. The adversary applies a unitary A_i on $A_i R_i$ to obtain a state on registers $A_{i+1} M_i$ and sends the message register M_i to the challenger.
2. The challenger then performs a unitary C_i to registers $M_i W_i$, to obtain a state over registers $R_i W_{i+1}$, here R_i represents the response register. The challenger sends R_i to the adversary.

We use M_i and R_i to distinguish the messages sent by the adversary and challenger respectively.

In the special case where $r = 1$ (i.e., a 3-message interactive protocol), once the adversary prepares their initial state, both the adversary and challenger have a single unitary to apply. As a result, we drop the round index and refer to the adversary and challenger unitaries as A and C , and the adversary’s initial state as $|\text{aux}\rangle$.

⁴In the case of a non-uniform adversary, A_0 may be initialized to an advice state $|\text{aux}\rangle$, in which case A_0 may be taken to be the identity. For uniform adversaries we can assume that A_0 is in the all-zero state.




Figure 1: A $2r + 1$ -message quantum interactive protocol.

Without loss of generality, we assume that after the adversary has sent their final message, the challenger performs some 2-outcome measurement, described by the POVM $\{D, id - D\}$, on $M_r W_r$. If the measurement outputs D , then the challenger accepts and otherwise the challenger rejects. In this work, we focus on the *soundness* of an interactive protocol. We say that an interactive protocol has soundness s if for every polynomial time (in the security parameter λ) adversary, the adversary is accepted by the challenger with probability at most $s(\lambda) + \text{negl}(\lambda)$.

Given an interactive protocol, π , we define the k -fold *parallel repetition* of the protocol, $\pi^{\otimes k}$ to be the protocol where the challenger and adversary execute π k -times in parallel. In $\pi^{\otimes k}$, the challenger performs the unitary $C_i^{\otimes k}$ on round i and measures the two outcome POVM $\{D^{\otimes k}, id - D^{\otimes k}\}$. This means that the challenger in the k -fold parallel repetition only accepts if every decision POVM accepts. We note that the adversary in the k -fold parallel repetition of the protocol may play a correlated, or even entangled strategy across the k copies of the protocol.

We use superscripts to denote the registers and unitaries applied by the individual repetitions of the protocol, so the first message register M_0 of the k -fold parallel repetition of π consists of k many registers M_0^i where each M_0^i is sent to the i 'th repetition of the protocol, and similarly for the registers W_j , and R_j . We assume that there is only one adversary private register.

4 Non-uniform parallel repetition of 3-message protocols

In this section, we prove the following non-uniform version of our main theorem.

Theorem 4.1. *For any 3-message quantum interactive protocol π with soundness ϵ against adversaries*

of size s , $\pi^{\otimes k}$ has soundness δ^k against adversaries of size $O\left(\frac{(\epsilon-\delta)\sqrt{\delta^{k-1}}}{\log \delta} \cdot \frac{s}{k}\right)$ for any k and $\delta > \epsilon$.

In particular, if a family of π has soundness ϵ against non-uniform polynomial-time adversaries, then $\pi^{\otimes k}$ has soundness ϵ^k against non-uniform polynomial-time adversaries for any polynomial k .

Let $(A, |\mathbf{aux}\rangle)$ be an adversary for $\pi^{\otimes k}$ that achieves success probability δ^k . For a collection of registers $\{\mathbf{C}^i\}_i$, let \mathbf{C} be the concatenation of all \mathbf{C}^i , and $\mathbf{C}^{(\leq j)}$ be the concatenation of the first j many \mathbf{C}^i . Let the following operators be the challenger unitaries C and decisions D that are only concerned about the first i repetitions of the protocol:

$$\begin{aligned} \left(C^{(\leq i)}\right)_{M_0^{(\leq i)}W_0^{(\leq i)}} &= \left(\bigotimes_{j=1}^i C_{M_0^j W_0^j}\right) \otimes \text{id}, \\ \left(D^{(\leq i)}\right)_{M_1^{(\leq i)}W_1^{(\leq i)}} &= \left(\bigotimes_{j=1}^i D_{M_1^j W_1^j}\right) \otimes \text{id}. \end{aligned}$$

Note that $C^{(\leq 0)} = D^{(\leq 0)} = \text{id}$. $D^{(<i)}$, $D^{(\neq i)}$, $D^{(>i)}$, etc. are defined similarly to be the restriction of $D^{\otimes k}$ to registers that satisfy the condition in the superscript.

Imagine a prefix of the k -fold game, where the k -fold adversary plays their strategy against all k challengers but only the first i repetitions perform C and measure D . Similar to [CHS05], our strategy for constructing an adversary for the 1-fold game will be to find an index i such that the i 'th repetition has a high probability of accepting, conditioned on the first $i - 1$ repetitions accepting, and then have the adversary “post-select” on the first $i - 1$ repetitions accepting. To that end, define the following projectors

$$G_i = \left(C^{(\leq k)}\right)^\dagger A_{A_1 R^{(\leq k)}}^\dagger \left(D^{(\leq i)}\right) A_{A_1 R^{(\leq k)}} \left(C^{(\leq k)}\right). \quad (6)$$

This further gives rise to the following POVM where we further enforce that the private registers of the verifiers is correctly initialized to $|0\rangle$.

$$\tilde{G}_i = \left(\text{id} \otimes \langle 0|_{W_0^{(\leq i)}}\right) G_i \left(\text{id} \otimes |0\rangle_{W_0^{(\leq i)}}\right).$$

Crucially note that this operator only enforces the initialization of the first i folds but permits arbitrary initialization of the other $(k - i)$ repetitions. In other words, its input registers are $A_1, R^{(\leq k)}, W_0^{(>i)}$.

Observe that $\left\|\tilde{G}_i |\mathbf{aux}\rangle\right\|^2$ exactly captures the probability of the success probability of winning the first i repetitions when initialized with $|\mathbf{aux}\rangle$, and thus $\left\|\tilde{G}_i\right\|_\infty^2$ captures the maximum success probability for the first i repetitions over all possible initializations. The following corollary shows that there does always exist a “good” index to put the real challenger, for a particular definition of “good” that would suffice later for the reduction.

Fact 4.2 (Discrete intermediate value theorem). *Let (p_0, \dots, p_k) be a sequence of reals such that $p_0 \leq 0 \leq p_k$. Then there exists an integer $1 \leq i \leq k$ such that $p_{i-1} \leq 0 \leq p_i$.*

Proof. Suppose this is not the case then we have $p_{i-1} \leq 0 \implies p_i < 0$, and thus by induction, $p_k < 0$, a contradiction. \square

Corollary 4.3. *There exists some $1 \leq i \leq k$ such that $\|\tilde{G}_{i-1}\|_\infty^2 \leq \delta^{i-1}$ and $\|\tilde{G}_i\|_\infty^2 \geq \delta^i$.*

Proof. Apply [Fact 4.2](#) to the sequence $\left\{ \left\| \tilde{G}_i \right\|_\infty^2 - \delta^i \right\}_i$. □

Since our reduction is non-uniform for now, we can assume that the adversary knows a good index and starts with an advice state that certifies the largeness of $\|\tilde{G}_i\|_\infty$. In order to leverage the smallness of $\|\tilde{G}_{i-1}\|_\infty$, we need to invoke the Quantum Singular Value Transform.

Theorem 4.4 (Uniform singular value amplification [[GSLW19](#), Theorem 17 (rephrased)]). *Let $\Pi, \tilde{\Pi}$ be projectors and $\gamma > 1$ and $\mu, \nu \in (0, \frac{1}{2})$. Let $\tilde{\Pi}\Pi = \sum_i \varsigma_i |w_i\rangle\langle v_i|$ be a singular value decomposition. Then there is an $m = O(\frac{\gamma}{\mu} \log(\frac{\gamma}{\nu}))$ and efficiently computable $\Phi \in \mathbb{R}^m$ such that*

$$\left(\langle + | \otimes \tilde{\Pi}_{\leq \frac{1-\mu}{\gamma}} \right) U_\Phi \left(| + \rangle \otimes \Pi_{\leq \frac{1-\mu}{\gamma}} \right) = \sum_{i: \varsigma_i \leq \frac{1-\mu}{\gamma}} \tilde{\varsigma}_i |w_i\rangle\langle v_i|, \text{ where } \left| \frac{\tilde{\varsigma}_i}{\gamma \varsigma_i} - 1 \right| \leq \nu$$

and where, for $\kappa \in [0, 1]$ (using notation from [Definition 3.1](#)),

$$\Pi_{\leq \kappa} := \sum_{i: \varsigma_i \leq \kappa} |v_i\rangle\langle v_i| = \Pi_{\leq \kappa^2}^{\tilde{\Pi}\Pi} \quad \text{and} \quad \tilde{\Pi}_{\leq \kappa} := \sum_{i: \varsigma_i \leq \kappa} |w_i\rangle\langle w_i| = \tilde{\Pi}_{\leq \kappa^2}^{\tilde{\Pi}\Pi}.$$

Moreover U_Φ can be implemented using a single ancilla qubit with m uses of $C_{\Pi}NOT$, m uses of $C_{\tilde{\Pi}}NOT$ and m single qubit gates.

Note that in the theorem statement, ς_i are the *singular values* of $\tilde{\Pi}\Pi$; the eigenvalues of $\tilde{\Pi}\Pi\tilde{\Pi}$, used later in the Jordan decomposition, are obtained as $p_i = \varsigma_i^2$. Applying U_Φ to any pure state “simulates” the boosted singular value transform up to a small error ν . The following corollary makes the error more convenient later in our proof of the main theorem.

Corollary 4.5. *Let $\varsigma_i |w_i\rangle\langle v_i|$, $\tilde{\varsigma}_i$, μ and ν be as defined in [Theorem 4.4](#). For any pure state $|\psi\rangle$, let*

$$\begin{aligned} |\tilde{\phi}\rangle &= \left(\sum_{i: \varsigma_i \leq \frac{1-\mu}{\gamma}} \tilde{\varsigma}_i |w_i\rangle\langle v_i| \right) |\psi\rangle, \\ |\phi\rangle &= \left(\sum_{i: \varsigma_i \leq \frac{1-\mu}{\gamma}} \gamma \varsigma_i |w_i\rangle\langle v_i| \right) |\psi\rangle. \end{aligned}$$

Then $|\tilde{\phi}\rangle = |\phi\rangle + |\nu\rangle$ where $\|\nu\|_2 \leq \nu$.

Proof. We simply verify by calculating

$$\begin{aligned}
\| |\nu\rangle \|_2 &= \left\| |\tilde{\phi}\rangle - |\phi\rangle \right\|_2 \\
&= \left\| \sum_{i: \zeta_i \leq \frac{1-\mu}{\gamma}} (\tilde{\zeta}_i - \gamma \zeta_i) |w_i\rangle \langle v_i | \psi \rangle \right\|_2 \\
&= \sqrt{\sum_{i: \zeta_i \leq \frac{1-\mu}{\gamma}} |(\tilde{\zeta}_i - \gamma \zeta_i) \langle v_i | \psi \rangle|^2} \\
&\leq \max_{i: \zeta_i \leq \frac{1-\mu}{\gamma}} |\tilde{\zeta}_i - \gamma \zeta_i| \\
&\leq \max_{i: \zeta_i \leq \frac{1-\mu}{\gamma}} \gamma \zeta_i \nu \\
&\leq \nu,
\end{aligned}$$

where the second inequality is by the guarantee of the algorithm. \square

We are now ready to state and prove the main result of this section. In order to use this lemma in [Section 5](#), we introduce a parameter τ to account for a loss in the uniform reduction. For the non-uniform case, it suffices to set $\tau = 1$ since we can prepare the advice state without any loss.

Lemma 4.6. *Let i and $|\text{aux}\rangle$ be an index and state such that $\left\| \Pi_{>\delta^{i-1}}^{\tilde{G}_{i-1}} |\text{aux}\rangle \right\|^2 = 0$ and $\text{Tr}(\tilde{G}_i |\text{aux}\rangle \langle \text{aux}|) \geq \tau \delta^i$ for some $0 < \tau \leq 1$. Then for any $0 < \mu < \frac{1}{2}$, there exists an adversary that takes as input i, μ, δ, τ and $|\text{aux}\rangle$ and is accepted by the 1-fold verifier with probability at least $(1 - 2\mu)^2 \tau \delta$. The adversary's running time is dominated by running the original k -fold prover $O\left(\frac{i}{\mu \sqrt{\delta^{i-1}}} \log \frac{1}{\mu \delta \tau}\right)$ times.*

Proof. At a high level, the goal of the single-fold adversary will be to use [Theorem 4.4](#) to coherently do post selection such that the first $(i - 1)$ repetitions are accepted, in which case the i 'th repetition should also be accepted with decent probability by the theorem assumption.

Let A be the unitary that the adversary performs on registers $\text{AR}^{(\leq k)}$ in the k -fold game. Apply [Theorem 4.4](#) with the following projectors and parameters.

$$\Pi = \left(C^{(<i)} \otimes \text{id} \right) \left(\text{id} \otimes |0\rangle \langle 0|_{W_0^{(<i)}} \right) \left(\left(C^{(<i)} \right)^\dagger \otimes \text{id} \right),$$

$$\tilde{\Pi} = \left(A_{\text{A}_1 \text{R}^{(\leq k)}}^\dagger \otimes \text{id} \right) \left(D^{(<i)} \otimes \text{id} \right) \left(A_{\text{A}_1 \text{R}^{(\leq k)}} \otimes \text{id} \right), \quad (7)$$

$$\gamma = \frac{1 - \mu}{\sqrt{\delta^{i-1}}}, \quad (8)$$

$$\nu = \mu \sqrt{\tau \delta}.$$

Let W be the unitary satisfying the conclusions [Theorem 4.4](#) for that choice of parameters. Note that W acts on A , the k -fold adversaries private register, $W_1^{(<i)}$, the $(i - 1)$ many simulated challengers' private workspaces after sending the challenges, $\text{R}^{(\leq k)}$, the response registers for all k challengers

(which A expects to act on), and an ancilla register P , which will be projected onto the $|+\rangle\langle+|$ state. Now consider the following prover for the 1-fold game (the challenger's actions are included in monospace font to aid understanding).

Algorithm 1. *Non-uniform adversary Amp for the 1-fold protocol*

Input: Quantum registers $AM_0^{(\leq k)}W_0^{>i}$, index i , and slackness parameter μ , and black-box oracle access to A, C, D .

1. Initialize registers $W_0^{(<i)}$ to $|0\rangle$, the private workspace registers for the first $(i-1)$ simulated challengers.
2. Send M_0^i to the challenger as the first message.
(Challenger performs $C_{W_0^i M_0^i}$ and sends R^i back.)
3. Perform $C^{(\neq i)}$ on registers $W_0 M_0^{(\neq i)}$ to get a state on $W_1 R^{(\neq i)}$.
4. Create ancilla register P initialized in $|+\rangle_P$.
5. Perform $W_{AW_1^{(<i)} R^{(\leq k)} P}$ as defined above.
6. Perform $A_{AR^{(\leq k)}}$.
7. Measure $D^{(<i)} \otimes |+\rangle\langle+|_P$. If the measurement rejects, abort.
8. Send M_1^i to the challenger.
(Challenger measures $D_{W_1 M_1^i}$ and accepts or rejects.)

We analyze the algorithm by describing the state of the combined prover-verifier system after every step.

After step 1: We assume that the registers $AM_0^{(\leq k)}$ are initialized in the state $|\text{aux}\rangle_{AM_0^{(\leq k)}W_0^{>i}}$ satisfying the theorem statement. Thus, state of the adversary and challenger after **step 1** is

$$|\text{init}\rangle := |\text{aux}\rangle_{AM_0^{(\leq k)}W_0^{(>i)}} \otimes |0\rangle_{W_0^{(\leq i)}}.$$

After step 4: In **step 3**, together with the verifiers action, the verifier and prover perform $C^{(\leq k)}$ on registers $W_0^{(\leq k)}M_0^{(\leq k)}$, so the state of the system after **step 4** is given by

$$C^{(\leq k)} |\text{init}\rangle \otimes |+\rangle_P.$$

This state is in the +1 eigenstate of $(\Pi \otimes |+\rangle\langle+|_P)$ as the first $(i-1)$ repetitions are initialized correctly, so we can write the state as

$$(\Pi \otimes |+\rangle_P) C^{(\leq k)} |\text{init}\rangle.$$

After step 7: After measuring $D^{(<i)} \otimes |+\rangle\langle+|_P$, we get the following state.

$$(D^{(<i)} \otimes \langle+|_P) AW (\Pi \otimes |+\rangle_P) C^{(\leq k)} |\text{init}\rangle.$$

The state might be subnormalized since we might abort in the event that the measurement rejects; aborting also ensures that the rejection part of the amplitude would not interfere with the rest of the algorithm and the analysis. Now recall that $\tilde{\Pi} = A^\dagger D^{(<i)} A$, thus $D^{(<i)} = A\tilde{\Pi}A^\dagger$. Performing this substitution gives us the following expression for the state above.

$$A(\tilde{\Pi} \otimes \langle + |_{\mathbf{P}})W(\Pi \otimes | + \rangle_{\mathbf{P}})C^{(\leq k)} |\text{init}\rangle.$$

Note that we choose the parameter at (8) so that $\frac{1-\mu}{\gamma} = \sqrt{\delta^{i-1}}$ and thus we can now apply the guarantee of [Corollary 4.5](#) and get that the state can be written as

$$\gamma A\tilde{\Pi}\Pi C^{(\leq k)} |\text{init}\rangle + |\nu\rangle$$

for some $\| |\nu\rangle \|_2 \leq \nu$. We now use again the fact that $C^{(\leq k)} |\text{init}\rangle$ is invariant under Π to remove Π . We further plug in the definition of $\tilde{\Pi}$, yielding

$$\gamma D^{(<i)} A C^{(\leq k)} |\text{init}\rangle + |\nu\rangle,$$

which is exactly the “post-selection” state we would like to prepare up to a small error.

After the final verifier decision: After the verifier measures $D_{W_1 M_1}$, we obtain the state

$$\gamma D^{(\leq i)} A C^{(\leq k)} |\text{init}\rangle + |\nu'\rangle,$$

where $|\nu'\rangle := D_{W_1 M_1} |\nu\rangle$ which still has 2-norm at most ν as $D \preceq \text{id}$.

By the theorem’s assumption, the state $|\text{aux}\rangle$ satisfies $\text{Tr}(\tilde{G}_i |\text{aux}\rangle \langle \text{aux}|) \geq \tau \delta^i$, thus $\| D^{(\leq i)} A C^{(\leq k)} |\text{init}\rangle \| = \| G_i |\text{init}\rangle \| = \| \tilde{G}_i |\text{aux}\rangle \| \geq \sqrt{\tau \delta^i}$. Therefore, the above state has 2-norm at least

$$\gamma \sqrt{\tau \delta^i} - \nu = (1 - \mu) \sqrt{\tau \delta} - \mu \sqrt{\tau \delta} = (1 - 2\mu) \sqrt{\tau \delta}.$$

So the prover is accepted with probability at least $(1 - 2\mu)^2 \tau \delta$. Finally the running time can be verified by plugging in the appropriate parameters and noting that A is only used in [step 6](#) and $\tilde{\Pi}$ as defined in [equation \(7\)](#) in [step 5](#). \square

Note that since the theorem holds for any pure state $|\text{aux}\rangle$ and since the algorithm is linear, it immediately extends to mixed state inputs satisfying the same condition as well.

Finally, combining [Lemma 4.6](#) (instantiating $\mu = \frac{\delta - \epsilon}{8}$ and $\tau = 1$) with [Corollary 4.3](#), we obtain the following non-uniform reduction.

Proof of [Theorem 4.1](#). The first part of the theorem immediately follows by picking the parameters above.

For the second part, assume the adversary’s success probability is non-negligibly larger than ϵ^k . Then there exists some function $\delta = \delta(\lambda)$ such that $\delta^k - \epsilon^k$ is some inverse polynomial and the adversary’s success probability is at least δ^k infinitely often. Whenever the adversary in the k -fold parallel repetition of the original protocol achieves δ^k , the success probability of [Algorithm 1](#) is at least $(1 - 2\mu)^2 \delta \geq (1 - 4\mu) \delta \geq \frac{\delta + \epsilon}{2}$, which is larger than ϵ by a non-negligible function in λ , since $\delta - \epsilon \geq \frac{\delta^k - \epsilon^k}{k}$ for any real $\epsilon \leq \delta \leq 1 \leq k$. Finally, the running time of this adversary has a multiplicative overhead of $\tilde{O}\left(\frac{k}{\mu \sqrt{\delta^{k-1}}}\right) = \tilde{O}\left(\frac{k}{(\delta - \epsilon) \sqrt{\delta^{k-1}}}\right) = \tilde{O}\left(\frac{k^2}{(\delta^k - \epsilon^k) \sqrt{\delta^{k-1}}}\right) = \tilde{O}(k^2 \delta^{-3k/2})$; therefore, this new adversary is efficient as $\delta^k - \epsilon^k$ is inverse polynomial and k is polynomial. \square

5 Uniform parallel repetition of 3-message protocols

In this section, we prove the uniform version of our main theorem.

Theorem 5.1. *Let π be a 3-message quantum interactive protocol with soundness s against polynomial-time (resp. polynomial-size) quantum adversaries. Then for any polynomial $k = k(\lambda)$, $\pi^{\otimes k}$ has soundness s^k against polynomial-time (resp. polynomial-size) quantum adversaries.*

We do this by giving an efficient algorithm which prepares such a state from (polynomially many copies of) the adversary’s initial state $|\text{aux}\rangle$. Formally, we show the following, from which the theorem is immediate.

Lemma 5.2. *There is a polynomial-time quantum oracle algorithm Amp-U with the following guarantee. Let $\pi = \{C, D\}$ be a 3-message quantum interactive protocol. For $k \in \mathbb{N}$, $\delta \in [0, 1]$, let $(A, |\text{aux}\rangle)$ be an adversary against $\pi^{\otimes k}$ which causes the challenger to accept with probability δ^k . Then for every ϵ there is a $t = (\epsilon\delta^k)^{-O(1)}$ such that $(\text{Amp-U}^{A,C,D}(1^{1/\delta^k}, 1^{1/\epsilon}), |\text{aux}\rangle^{\otimes t})$ is an adversary against π (i.e., a single repetition) which causes the challenger to accept with probability $\delta - \epsilon$.*

5.1 Jordan’s lemma and alternating projectors

Lemma 5.3 (Jordan’s lemma [Jor75]). *For any two Hermitian projectors Π_A and Π_B on a Hilbert space \mathbb{H} , there exists an orthogonal decomposition of $\mathbb{H} = \bigoplus_j \mathcal{S}_j$ (the Jordan decomposition with respect to Π_A, Π_B) into one-dimensional and two-dimensional subspaces $\{\mathcal{S}_j\}_j$ (the Jordan subspaces), where each \mathcal{S}_j is invariant under both Π_A and Π_B . Moreover:*

- in each one-dimensional space, Π_A and Π_B act as identity or rank-zero projectors; and
- in each two-dimensional subspace \mathcal{S}_j , Π_A and Π_B are rank-one projectors. In particular, there exist distinct orthogonal bases $\{|v_j^1\rangle, |v_j^0\rangle\}$ and $\{|w_j^1\rangle, |w_j^0\rangle\}$ for \mathcal{S}_j such that Π_A projects onto $|v_j^1\rangle$ and Π_B projects onto $|w_j^1\rangle$.

In order to unify the treatment of one- and two-dimensional subspaces, for a one-dimensional subspace $\mathcal{S}_j = \text{span}(|v\rangle)$ we denote $|v\rangle$ both by $|v_j^{\lambda_0}\rangle$ for $\Pi_A |v\rangle = \lambda_0 |v\rangle$, and by $|w_j^{\lambda_1}\rangle$ for $\Pi_B |v\rangle = \lambda_1 |v\rangle$. We define $|v_j^{1-\lambda_0}\rangle$ and $|w_j^{1-\lambda_1}\rangle$ to be the zero vector.

Definition 5.4. *For two Hermitian projectors Π_A, Π_B we define the Jordan measurement to be the projective measurement $\mathcal{P}_{\text{Jor}}[\Pi_A, \Pi_B] := (\Pi_j^{\text{Jor}})_j$ with outcomes j , where Π_j^{Jor} projects on to \mathcal{S}_j .*

We define the value of the subspace \mathcal{S}_j to be $p_j := |\langle v_j^1 | w_j^1 \rangle|^2$.

The following straightforward but useful claim relates the Jordan decomposition with respect to Π_A, Π_B to the spectral decompositions of $\Pi_A \Pi_B \Pi_A$ and $\Pi_B \Pi_A \Pi_B$.

Claim 5.5. $\Pi_A \Pi_B \Pi_A = \sum_j p_j |v_j^1\rangle\langle v_j^1|$, and $\Pi_B \Pi_A \Pi_B = \sum_j p_j |w_j^1\rangle\langle w_j^1|$.

Alternating projectors. Jordan’s lemma allows us to characterize the behavior of an alternating sequence of binary projective measurements. Define the following (classical) probability distribution $\text{MWDist}(p, T)$ (for “Marriott–Watrous distribution”), parameterized by $p \in [0, 1]$ and $T \in \mathbb{Z}$:

MWDist(p, T):

1. For each $i \in [T]$, set $a_i := 0$ with probability p and $a_i := 1$ otherwise.
2. Let $b_0 := 1$. For each $i \in [T]$, define $b_i := b_{i-1} \oplus a_i$.
3. Output b_1, b_2, \dots, b_T .

The following is a straightforward consequence of Jordan's lemma; see e.g. [CMSZ22] for a proof.

Lemma 5.6. *The measurement outcomes that result from applying T alternating binary projective measurements $\Pi_A, \Pi_B, \Pi_A, \Pi_B, \dots$ to the state $\sum_j \alpha_j |w_j^1\rangle$ have the following distribution:*

1. sample p_j with probability $|\alpha_j|^2$;
2. output MWDist(p_j, T).

Via a Chernoff bound, we then obtain the following very useful result.

Claim 5.7. *For $b_0, b_1, \dots, b_n \in \{0, 1\}$, define*

$$\text{NumReps}(b_0, b_1, \dots, b_n) := \frac{|\{j \in \{1, \dots, n\} : b_{j-1} = b_j\}|}{n}.$$

Fix $p \in [0, 1], T \in \mathbb{Z}$. Let $X := \text{NumReps}(1, b_1, \dots, b_T)$ for $b_1, \dots, b_T \leftarrow \text{MWDist}(T, p)$. Then $\mathbb{E}[X] = p$, and for any $\varepsilon, \delta \in [0, 1]$, if $T \geq \lceil \ln(\frac{1}{2\delta}) / (2\varepsilon^2) \rceil$,

$$\Pr[|X - p| \leq \varepsilon] \geq 1 - \delta.$$

An important consequence of the above is the existence of an efficient (ε, δ) -almost projective measurement related to the Jordan decomposition.

Definition 5.8 ([Zha20]). *A real-valued measurement \mathcal{M} is (ε, δ) -almost-projective if applying \mathcal{M} twice in sequence to any state ρ produces measurement outcomes p, p' where*

$$\Pr[|p - p'| \leq \varepsilon] \geq 1 - \delta.$$

Lemma 5.9. *For any $\varepsilon, \delta > 0$, and binary projective measurements (Π_A, Π_B) , there is an (ε, δ) -almost projective measurement $\text{EffJor}_{\varepsilon, \delta}$ which applies Π_A, Π_B a total of $O(\frac{1}{\varepsilon^2} \log \frac{1}{\delta})$ times, with the following properties:*

- if $\text{Tr}(\Pi_A \rho) = 1$ then $\mathbb{E}_{p \leftarrow \text{EffJor}(\rho)}[p] = \text{Tr}(\Pi_B \rho) = \text{Tr}(\Pi_A \Pi_B \Pi_A \rho)$;
- if $\Pr[\text{EffJor}(\rho) \geq p] \geq \gamma$ then $\sum_{j, p_j \geq p - \varepsilon} \langle v_j^1 | \rho | v_j^1 \rangle \geq \gamma - \delta$;
- similarly, if $\Pr[\text{EffJor}(\rho) \leq p] \geq \gamma$ then $\sum_{j, p_j \leq p + \varepsilon} \langle v_j^1 | \rho | v_j^1 \rangle \geq \gamma - \delta$;

Proof sketch. Let $\Pi'_A := \Pi_A \otimes |00\rangle\langle 00|$, $\Pi'_B = \Pi_B \otimes |++\rangle\langle ++| + I \otimes |--\rangle\langle --|$. The algorithm works as follows:

1. Measure Π'_A on $\rho \otimes |00\rangle\langle 00|$; if the outcome is 0, abort.

2. Alternate Π'_B, Π'_A $T = 2 \lceil \ln(\frac{1}{\delta}) / \varepsilon^2 \rceil$ times, obtaining outcomes b_1, \dots, b_T .
3. Continue alternating until $\Pi'_A \rightarrow 1$, or at most $k = 3 \log \frac{2}{\delta}$ times.
4. Output $4(\text{NumReps}(1, b_1, \dots, b_T) - \frac{1}{4})$.

If $|v_j^1\rangle$ is an eigenvector of $\Pi_A \Pi_B \Pi_A$ with eigenvalue p_j , then $|v_j^1\rangle \otimes |00\rangle$ is an eigenvector of $\Pi'_A \Pi'_B \Pi'_A$ with eigenvalue $\frac{1}{4}p_j + \frac{1}{4} \in [\frac{1}{4}, \frac{1}{2}]$. Hence the transition probability when alternating Π'_A, Π'_B is between $\frac{1}{4}$ and $\frac{1}{2}$, so the probability that after k applications we have not reached Π'_A is at most $(\frac{3}{4})^k \leq \delta/2$. Combined with the usual analysis of alternating projectors [Zha20, CMSZ22], the lemma statement follows. \square

5.2 State transformation for almost-projective measurements

In this section, we describe an algorithm for the following problem. Let $\mathcal{M}_0, \mathcal{M}_1$ be (ε, δ) -almost projective measurements (Definition 5.8). Given a state ρ such that $\mathcal{M}_0(\rho) \geq \alpha$ with high probability, and a target $\beta \in [0, 1]$, efficiently prepare a state σ so that either:

- (i) $\mathcal{M}_1(\sigma) \gtrsim \beta$ with high probability, or
- (ii) both $\mathcal{M}_0(\sigma) \gtrsim \alpha$ and $\mathcal{M}_1(\sigma) < \beta$ with high probability.

That is, the algorithm either converts a “good” state with respect to \mathcal{M}_0 into a “good” state with respect to \mathcal{M}_1 , or produces a “good” state with respect to \mathcal{M}_0 which has small overlap with *any* state that is “good” with respect to \mathcal{M}_1 .

Before giving the algorithm, we set up some preliminaries. By Naimark dilation, any measurement $\mathcal{M} = (M_q)_{q \in \{0,1\}^n}$, can be implemented as a unitary $U_{\mathcal{M}}$ on $\mathbb{H} \otimes \mathbb{W}$ for some ancilla register \mathbb{W} initialized to zero, followed by some projective measurement $(\Pi_q)_{q \in \{0,1\}^n}$ on \mathbb{W} , where the Π_q are independent of \mathcal{M} . Formally, for each $q \in \{0,1\}^n$, the unitary $U_{\mathcal{M}}$ and projector Π_q satisfy $M_q \rho M_q^\dagger = \text{Tr}_{\mathbb{W}}(\Pi_q U_{\mathcal{M}}(\rho \otimes |0\rangle\langle 0|_{\mathbb{W}}) U_{\mathcal{M}}^\dagger)$ for all states ρ . By “black-box unitary access to \mathcal{M} ”, we mean access to $U_{\mathcal{M}}, U_{\mathcal{M}}^\dagger$, and access to the unitary $\sum_{q, q' \in \{0,1\}^n} \Pi_q \otimes |q' \oplus q\rangle\langle q'|$.

The main result of this section is the following lemma.

Lemma 5.10. *Let $\mathcal{M}_0, \mathcal{M}_1$ be (ε, δ) -almost projective measurements on the same system \mathbb{H} for some $0 < \delta < 1$. There is an algorithm `StateTrans` such that for every state ρ , and real numbers $\alpha, \beta, \gamma \in [0, 1]$, $\tau \in [2\varepsilon, 1 - \delta]$ satisfying $\Pr[\mathcal{M}(\rho) \geq \alpha] \geq 1 - \gamma$, letting $(\sigma, c) \leftarrow \text{StateTrans}_{\varepsilon, \delta, \tau}[\mathcal{M}_0, \mathcal{M}_1, \beta](\rho)$ and $q_b \leftarrow \mathcal{M}_b(\sigma)$, the following hold, for $K = \lceil \frac{2}{\tau} \ln \frac{1}{\delta} \rceil$:*

1. $\Pr[c = \perp] \leq 4K\sqrt{\delta}$,
2. $\Pr[c = 0 \wedge q_1 < \beta - \varepsilon] \leq \delta$,
3. $\Pr[c = 1 \wedge q_0 < \alpha - 2K\varepsilon] \leq \gamma + \sqrt{\tau + \varepsilon + \delta}$, and
4. $\Pr[c = 1 \wedge q_1 > \beta] \leq \tau + \varepsilon + \delta$.

Furthermore `StateTrans` _{$\varepsilon, \delta, \tau$} runs in expected time $O(K)$, given black-box unitary access to $\mathcal{M}_0, \mathcal{M}_1$.

That is, `StateTrans` takes as input a state ρ and outputs a state σ and a flag c such that if $c = 0$ then σ satisfies condition (i) and if $c = 1$ then σ satisfies condition (ii). ($c = \perp$ indicates failure.)

Let $\mathcal{M}_0, \mathcal{M}_1$ be almost projective measurements on the same register \mathbf{H} , and let $\mathbf{W}_0, \mathbf{W}_1$ be corresponding ancilla registers for the Naimark dilations of \mathcal{M}_0 and \mathcal{M}_1 respectively. We define the projectors:

$$\Pi_{\geq p}^b := \sum_{q \geq p} U_{\mathcal{M}_b}^\dagger \Pi_q \Pi_q U_{\mathcal{M}_b} \quad \text{and} \quad \Pi_{< p}^b := \text{id}_{\mathbf{H} \otimes \mathbf{W}_b} - \Pi_{\geq p}^b.$$

We define $\tilde{\Pi}_{\geq p}^b := \Pi_{\geq p}^b \otimes |0\rangle\langle 0|_{\mathbf{W}_{1-b}}$, and define $\tilde{\Pi}_{< p}^b := \Pi_{< p}^b \otimes |0\rangle\langle 0|_{\mathbf{W}_{1-b}}$.

Algorithm 2. `StateTrans` $_{\varepsilon, \delta, \tau}[\mathcal{M}_0, \mathcal{M}_1, \beta]$.

Input: Quantum register \mathbf{H} .

1. Let $K := \lceil \frac{2}{\tau} \ln \frac{1}{\delta} \rceil$, $T := \lceil 1/\sqrt{\delta} \rceil$. For $i = 1, \dots, K$:
 - (a) Apply the (ε, δ) -almost-projective measurement \mathcal{M}_0 , obtaining outcome α_i .
 - (b) Initialize ancilla registers $\mathbf{W}_0, \mathbf{W}_1$ to $|0\rangle$.
 - (c) Apply the measurement $\tilde{\Pi}_{< \beta}^1$, obtaining outcome b_1 . If $b_1 = 0$, apply $U_{\mathcal{M}_1}$ to $(\mathbf{H}, \mathbf{W}_1)$, discard the \mathbf{W} registers, and output $c = 0$ along with the \mathbf{H} register.
 - (d) Apply the measurements $\tilde{\Pi}_{\geq \alpha_i - \varepsilon}^0, \tilde{\Pi}_{< \beta}^1$ in an alternating fashion $K-1$ times, obtaining outcomes b_2, \dots, b_{2K-1} .
 - (e) Apply the measurements $\tilde{\Pi}_{\geq \alpha_i - \varepsilon}^0, \tilde{\Pi}_{< \beta}^1$ in an alternating fashion until $\tilde{\Pi}_{\geq \alpha_i - \varepsilon}^0 \rightarrow 1$, or until $2TK + 1$ measurements have been applied in this step ($2(T+1)K$ overall). In the latter case, abort (output \perp).
 - (f) If $\text{NumReps}(b_1, \dots, b_{2K-1}) \geq 1 - \tau$, discard the \mathbf{W} registers and output $c = 1$ along with the \mathbf{H} register.
 - (g) Apply $U_{\mathcal{M}_0}$ to $(\mathbf{H}, \mathbf{W}_0)$, and discard the \mathbf{W} registers.
2. If the algorithm does not terminate with output above, abort (output \perp).

In the proof we will make use of the following lemma, due to [LMS22]:

Lemma 5.11 (Pseudoinverse lemma). *Let $\Pi_{\mathbf{A}}, \Pi_{\mathbf{B}}$ be projectors, and $(\Pi_j^{\text{Jor}})_j := \mathcal{P}_{\text{Jor}}[\Pi_{\mathbf{A}}, \Pi_{\mathbf{B}}]$ the corresponding Jordan measurement. Let ρ be a state such that $\Pi_{\mathbf{A}}\rho = \rho\Pi_{\mathbf{A}}$ and $\text{Tr}(\Pi_{\mathbf{A}}\rho) \geq 1 - \gamma$, and let $\Pi_0 := \sum_{j, p_j=0} \Pi_j^{\text{Jor}}$. Let $E := \sum_{j, p_j>0} \frac{1}{p_j} \Pi_j^{\text{Jor}}$. There exists a “pseudoinverse” state σ with $\text{Tr}(\Pi_{\mathbf{B}}\sigma) = 1$ such that all of the following are true:*

1. $\text{Tr}(\Pi_{\mathbf{A}}\sigma) = \frac{1 - \text{Tr}(\Pi_0\rho)}{\text{Tr}(E\rho)}$,
2. $\text{td}\left(\rho, \frac{\Pi_{\mathbf{A}}\sigma\Pi_{\mathbf{A}}}{\text{Tr}(\Pi_{\mathbf{A}}\sigma)}\right) \leq \sqrt{\text{Tr}(\Pi_0\rho)}$,
3. for all j such that $p_j > 0$ it holds that $\text{Tr}(\Pi_j^{\text{Jor}}\sigma) = \frac{\text{Tr}(\Pi_j^{\text{Jor}}\rho)}{p_j \cdot \text{Tr}(E\rho)}$, and
4. $\text{Tr}(\Pi_0\sigma) = 0$.

Claim 5.12. *Let ρ be such that $\Pi_A \rho = \rho \Pi_A$ and $\text{Tr}(\Pi_0 \rho) = 0$. Then there exists a state ρ' such that $\text{Tr}(E \rho') = \text{Tr}(E \rho)$, $\text{Tr}(\Pi_A \rho') = 1$, and $\text{td}(\rho, \rho') \leq 1 - \text{Tr}(\Pi_A \rho)$.*

Proof. This is a variant of [LMS22, Claim 7.2]. Let

$$U := \sum_{j, p_j \notin \{0,1\}} (|v_j^1\rangle \langle v_j^0| + |v_j^0\rangle \langle v_j^1|) + \sum_{j, p_j \in \{0,1\}} |v_j^{p_j}\rangle \langle v_j^{p_j}|$$

and let $\rho' := \Pi_A \rho + U(I - \Pi_A)\rho U^\dagger$. Then $\text{td}(\rho, \rho') = \frac{1}{2} \|(I - \Pi_A)\rho - U(I - \Pi_A)\rho U^\dagger\|_1 \leq \text{Tr}((I - \Pi_A)\rho)$. \square

Claim 5.13. *For all ϵ, δ, τ and all initial states in H at [step 1a](#), every time the [StateTrans](#) $_{\epsilon, \delta, \tau}$ reaches [step 1e](#), it aborts with probability at most $2\sqrt{\delta}$ at that step and makes at most $9K$ measurements in expectation. Moreover, for all $1 < i \leq K$, $\Pr[b_1 = 1 \text{ in iteration } i - 1 \wedge \alpha_i < \alpha_{i-1} - 2\epsilon] \leq 3\sqrt{\delta}$.*

Proof. This follows via an analysis similar to [CMSZ22, Lemma 4.10], modified to account for the additional measurements in [step 1d](#).

First note that, because \mathcal{M}_0 is (ϵ, δ) -almost projective, the probability that a measurement of $\tilde{\Pi}_{\geq \alpha_i - \epsilon}^0$ accepts immediately after [step 1b](#) is $1 - \delta$. Hence by [Proposition 3.3](#), post-selecting on $\tilde{\Pi}_{\geq \alpha_i - \epsilon}^0 \rightarrow 1$ at this point disturbs the state by at most $\sqrt{\delta}$ in trace distance. Henceforth we therefore assume that the state immediately after [step 1b](#) is some $|\phi\rangle \in \text{img}(\tilde{\Pi}_{\geq \alpha_i - \epsilon}^0)$. (We without loss of generality argue this for all pure states $|\phi\rangle \in \text{img}(\tilde{\Pi}_{\geq \alpha_i - \epsilon}^0)$ and by linearity the argument also generalizes to mixed states.) This changes the probability of any event by at most an additive $\sqrt{\delta}$ and the expected running time by at most an additive $\sqrt{\delta}(2TK + 1) \leq \sqrt{\delta}(4K/\sqrt{\delta} + 1) \leq 5K$.

We first analyse the distribution of measurement outcomes in [step 1e](#). Indeed, let $b_{2K}, b_{2K+1}, \dots, b_{2TK}$ be the outcomes of those measurements. By [Lemma 5.6](#), we have that $b_1, \dots, b_{2TK} \sim \sum_j a_j \text{MWDist}(p_j, 2TK)$ for some $a_j \in [0, 1]$, $\sum_j a_j = 1$.

Fix any $p \in [0, 1]$, and consider $c_1, \dots, c_{2TK} \leftarrow \text{MWDist}(p, 2TK)$. Let $r := \Pr[c_{2K} = 0]$. Note that the variables c_2, c_4, \dots, c_{2TK} form a Markov chain with symmetric transition probabilities and initial state $c_0 = 1$. Hence, we have that $\Pr[c_{2iK} = 1 \mid c_{2(i-1)K} = 0] = r$ for all $2 \leq i \leq T$. It follows that

$$\Pr \left[\bigwedge_{i=1}^T c_{2iK} = 0 \right] = \Pr[c_{2K} = 0] \prod_{i=2}^T \Pr[c_{2iK} = 0 \mid c_{2(i-1)K} = 0] = r(1-r)^{T-1} \leq \frac{1}{T}, \quad (9)$$

where the last inequality follows by finding that LHS is maximized when $r = \frac{1}{T}$, giving maximum $\frac{1}{T} \cdot (1 - \frac{1}{T})^{T-1} \leq \frac{1}{T}$. Next, let D' be a random variable corresponding to the smallest $i \geq 1$ such that $c_{2iK} = 1$. Then

$$\mathbb{E}[D'] = \sum_{t=1}^T \Pr[D' \geq t] \leq 1 + r \sum_{t=0}^{\infty} (1-r)^t = 2. \quad (10)$$

Since the distribution of b_1, \dots, b_{2TK} is a convex combination of $\text{MWDist}(p_j, 2TK)$, it holds by convexity and (9) that $\Pr \left[\bigwedge_{i=1}^T b_{2iK} = 0 \right] \leq 1/T$. If (in particular) $b_{2iK} = 1$ for any $1 \leq i \leq T$, then [StateTrans](#) does not abort; hence A aborts with probability at most $1/T \leq \sqrt{\delta}$. Next, let D be the number of measurements applied in [step 1e](#). By convexity, linearity of expectation and [Equation \(10\)](#), $\mathbb{E}[D] \leq 4K$.

Finally we bound the probability that $\alpha_i < \alpha_{i-1} - 2\varepsilon$. The post-measurement state after [step 1c](#), conditioned on $b_1 = 1$, is $\tilde{\Pi}_{<\beta}^1 |\phi\rangle / \sqrt{q}$, where $q = \Pr[b_1 = 1] = \|\tilde{\Pi}_{<\beta}^1 |\phi\rangle\|^2 = \sum_j q_j p_j$ for $q_j \in [0, 1]$ such that $\sum_j q_j = 1$. $q > 0$ since otherwise we would reach [step 1c](#) with probability 0. Observe that $\|\Pi_j^{\text{Jor}} |\phi\rangle / \sqrt{q}\|^2 = q_j p_j / q$ for each j . Let ρ be the reduced density matrix on register $\mathsf{H}, \mathsf{W}_0, \mathsf{W}_1$ after [step 1e](#) (again conditioned on $b_1 = 1$). Let ρ' be the state guaranteed by [Claim 5.12](#); note that $\text{td}(\rho, \rho') \leq \sqrt{\delta}$ by (9). Since Π_j^{Jor} commutes with both $\tilde{\Pi}_{\geq \alpha_i - \varepsilon}^0$ and $\tilde{\Pi}_{<\beta}^1$, $\text{Tr}(\Pi_j^{\text{Jor}} \rho) = q_j p_j / q$, so

$$\text{Tr}(E\rho') = \text{Tr}(E\rho) = \sum_{j: p_j > 0} \frac{q_j}{q} \leq \frac{1}{q} \quad (11)$$

and $\text{Tr}(\Pi_0 \rho') = \text{Tr}(\Pi_0 \rho) = 0$ for E, Π_0 as defined in [Lemma 5.11](#).

By [Lemma 5.11](#) there is a pseudoinverse state $\sigma \in \text{img}(\tilde{\Pi}_{<\beta}^1)$ with $\rho' = \tilde{\Pi}_{\geq \alpha_i - \varepsilon}^0 \sigma \tilde{\Pi}_{\geq \alpha_i - \varepsilon}^0 / q'$ for $q' = \frac{1}{\text{Tr}(E\rho)} \geq q$ by (11). Since $\text{img}(\tilde{\Pi}_{<\beta}^1) \subseteq \text{img}(I_{\mathsf{H}} \otimes |0\rangle\langle 0|_{\mathsf{W}_0} \otimes I_{\mathsf{W}_1})$, $\sigma = \sigma'_{\mathsf{H}, \mathsf{W}_1} \otimes |0\rangle\langle 0|_{\mathsf{W}_0}$ for some σ' . Hence $\text{Tr}_{\mathsf{W}}(U_{\mathcal{M}_0} \rho' U_{\mathcal{M}_0}^\dagger)$ is precisely the post-measurement state after applying \mathcal{M}_0 to σ' and post-selecting on obtaining an answer greater than $\alpha_i - \varepsilon$. It follows that $\Pr[\alpha_i < \alpha_{i-1} - 2\varepsilon \mid b_1 = 1] = \Pr[p' < \alpha_{i-1} - 2\varepsilon \mid p \geq \alpha_i - \varepsilon]$ where p, p' are the results of applying \mathcal{M}_0 twice in sequence to σ . By the definition of conditional probability,

$$\Pr[p' < \alpha_{i-1} - 2\varepsilon \mid p \geq \alpha_i - \varepsilon] = \frac{\Pr[p' < \alpha_{i-1} - 2\varepsilon \wedge p \geq \alpha_i - \varepsilon]}{\Pr[p \geq \alpha_i - \varepsilon]} \leq \frac{\delta}{q}$$

since $\Pr[p \geq \alpha_i - \varepsilon] = q' \geq q$ and \mathcal{M}_0 is (ε, δ) -almost projective. The claim follows since $\Pr[b_1 = 1] = q$. \square

Claim 5.14. [StateTrans](#) aborts with probability at most $4K\sqrt{\delta}$.

Proof. We show that for each i , the probability that the algorithm stops (outputting either 0 or 1) in the i -th iteration is at least τ . Consider the following experiment:

[StateTrans](#) i -th round abort experiment

1. Initialize ancilla registers $\mathsf{W}_0, \mathsf{W}_1$ to $|0\rangle$.
- ★. Apply the Jordan subspace measurement, obtaining a subspace label j .
2. Apply the measurement $\tilde{\Pi}_{<\beta}^1$, obtaining outcome b_1 . If $b_1 = 0$, output YES.
3. Apply the measurements $\tilde{\Pi}_{\geq \alpha_i - \varepsilon}^0, \tilde{\Pi}_{<\beta}^1$ in an alternating fashion $K - 1$ times, obtaining outcomes b_2, \dots, b_K . If $\text{NumReps}(b_1, \dots, b_K) \geq 1 - \tau$, output YES.
4. Otherwise, output NO.

Observe that [StateTrans](#) stops at the i -th iteration without aborting if and only if this experiment, without step ★, outputs YES. Since step ★ commutes with $\tilde{\Pi}_{\geq \alpha_i - \varepsilon}^0$ and $\tilde{\Pi}_{<\beta}^1$, inserting ★ does not change the outcome probabilities. Note that given outcome j from step ★, the probability that step 2 outputs YES is $1 - p_j$, and that $\text{NumReps}(b_1, \dots, b_K) \sim \text{Bin}(K, p_j) / K$. Hence there are two cases:

- if $p_j < 1 - \tau + \varepsilon$, the probability that step 2 outputs YES is at least $\tau - \varepsilon$;
- if $p_j \geq 1 - \tau + \varepsilon$, the probability that step 3 outputs YES is at least $1 - \delta \geq \tau - \varepsilon$.

Since this lower bound holds regardless of the initial state at iteration i , it follows that the probability that the procedure fails to terminate with output within K steps is at most $(1 - (\tau - \varepsilon))^K \leq (1 - \tau/2)^K \leq e^{-\tau K/2} \leq \delta$.

Then the overall probability, using union bound and [Claim 5.13](#), is at most $3K\sqrt{\delta} + \delta \leq 4K\sqrt{\delta}$. \square

Next we show that, for any $\alpha, \beta \in [0, 1]$, a state with large overlap with both $\tilde{\Pi}_{\geq \alpha}^0$ and $\tilde{\Pi}_{< \beta}^1$ is likely to return an outcome $\geq \alpha$ when measured with \mathcal{M}_0 and $< \beta$ when measured with \mathcal{M}_1 . (Note this is not trivial due to the presence of ancillas.)

Claim 5.15. *For any state ρ satisfying $\text{Tr}(\tilde{\Pi}_{\geq \alpha}^0 \rho) = 1$ and $\text{Tr}(\tilde{\Pi}_{< \beta}^1 \rho) \geq 1 - \gamma$, it holds that*

$$\Pr[\mathcal{M}_0(\rho') \geq \alpha] \geq 1 - \sqrt{\gamma} \quad \text{and} \quad \Pr[\mathcal{M}_1(\rho') < \beta] \geq 1 - \gamma,$$

where $\rho' := \text{Tr}_{\mathcal{W}_0, \mathcal{W}_1}(\rho)$.

Proof. We have that $\Pr[\mathcal{M}_1(\rho') < \beta] = \text{Tr}(\Pi_{< \beta}^1(\rho' \otimes |0\rangle\langle 0|_{\mathcal{W}_1}))$. Since $\text{Tr}(\tilde{\Pi}_{\geq \alpha}^0 \rho) = 1$, $\rho' \otimes |0\rangle\langle 0|_{\mathcal{W}_1} = \text{Tr}_{\mathcal{W}_0}(\rho)$, and so $\Pr[\mathcal{M}_1(\rho') < \beta] = \text{Tr}(\tilde{\Pi}_{< \beta}^1 \rho) = 1 - \gamma$.

Similarly, $\Pr[\mathcal{M}_0(\rho') \geq \alpha] = \text{Tr}(\tilde{\Pi}_{\geq \alpha}^0(\rho' \otimes |0\rangle\langle 0|_{\mathcal{W}_0}))$. Since $\text{Tr}(\tilde{\Pi}_{< \beta}^1 \rho) \geq 1 - \gamma$, the states $\text{Tr}_{\mathcal{W}_1}(\rho)$ and $\rho' \otimes |0\rangle\langle 0|_{\mathcal{W}_0}$ are $\sqrt{\gamma}$ -close in trace distance by [Proposition 3.3](#). Hence $\Pr[\mathcal{M}_0(\rho') \geq \alpha] \geq \text{Tr}(\tilde{\Pi}_{\geq \alpha}^0 \rho) - \sqrt{\gamma} = 1 - \sqrt{\gamma}$. \square

We use this claim to prove that the algorithm outputs the correct state when it does not abort.

Claim 5.16. *For all quantum states σ and $\varepsilon, \delta, \tau$, let $(c, \rho) \leftarrow \text{StateTrans}_{\varepsilon, \delta, \tau}[\mathcal{M}_0, \mathcal{M}_1, \beta](\sigma)$ be the output of the algorithm. Then the following hold, where i is (a random variable corresponding to) the last iteration of the algorithm:*

1. $\Pr[c = 0 \wedge \mathcal{M}_1(\rho) < \beta - \varepsilon] \leq \delta$,
2. $\Pr[c = 1 \wedge \mathcal{M}_0(\rho) < \alpha_i - \varepsilon] \leq \sqrt{\tau + \varepsilon + \delta}$, and
3. $\Pr[c = 1 \wedge \mathcal{M}_1(\rho) \geq \beta] \leq \tau + \varepsilon + \delta$.

Proof. By the description of [StateTrans](#), $c = 0$ when the measurement $\tilde{\Pi}_{< \beta}^1$ yields outcome 0 in [step 1c](#). By the definition of an almost projective measurement, the probability that applying \mathcal{M}_1 subsequently yields outcome $< \beta - \varepsilon$ is less than δ , which proves the first inequality.

Consider inserting the Jordan subspace measurement, \mathcal{P}_{Jor} , after [step 1e](#), obtaining outcome j . Since \mathcal{P}_{Jor} commutes with $\tilde{\Pi}_{\geq \alpha - \varepsilon}^0$ and $\tilde{\Pi}_{< \beta}^1$, we can equivalently insert the Jordan subspace measurement before [Step 1d](#). Suppose that the outcome j satisfies $p_j < 1 - \tau - \varepsilon$; conditioned on receiving this outcome, by [Claim 5.7](#), $\Pr[\text{NumReps}(b_1, \dots, b_{2K-1}) \geq 1 - \tau] \leq \delta$.

By the definition of conditional probability we have that

$$\Pr[\text{NumReps}(b_1, \dots, b_{2K-1}) \geq 1 - \tau \wedge p_j < 1 - \tau - \varepsilon] \leq \delta. \quad (12)$$

Let ρ be the output state of [StateTrans](#), conditioned on halting with output $c = 1$; let $q := \Pr[c = 1]$. Then by construction, $\text{Tr}(\tilde{\Pi}_{\geq \alpha_i - \varepsilon}^0 \rho) = 1$, and by [equation \(12\)](#), $\text{Tr}(\sum_{j, p_j < 1 - \tau - \varepsilon} \Pi_j^{\text{Jor}} \rho) \leq \delta/q$. Then $\text{Tr}(\tilde{\Pi}_{< \beta}^1 \rho) \geq (1 - \tau - \varepsilon)(1 - \delta/q) \geq 1 - \tau - \varepsilon - \delta/q$. It follows from [Claim 5.15](#) that

$$\Pr[\mathcal{M}_0(\rho) \geq \alpha_i - \varepsilon] \geq 1 - \sqrt{\tau + \varepsilon + \delta/q} \quad \text{and} \quad \Pr[\mathcal{M}_1(\rho) < \beta] \geq 1 - \tau - \varepsilon - \delta/q.$$

The second two inequalities in the claim then follow by definition of the conditional probability and since $q \leq 1$. \square

We are now ready to put together the proof of Lemma 5.10.

Proof of Lemma 5.10. Claim 5.14 shows that $\Pr[c = \perp] \leq 4K\sqrt{\delta}$. Claim 5.16 shows the correctness of the algorithm when the algorithm does not abort. Claim 5.13 shows that **StateTrans** makes $O(K)$ measurements in expectation. \square

5.3 Proof of Lemma 5.2

We are now ready to prove Lemma 5.2. First recall the definition of G_i from equation (6):

$$G_i = (C^{\otimes k})^\dagger A^\dagger \left(D^{(\leq i)} \right) A \cdot C^{\otimes k}.$$

Let $\epsilon_0 := \epsilon\delta^k/4$, $\tau := \epsilon_0^2/100$, $\hat{\delta} := \frac{\tau^6}{8k^3}$, $\hat{\epsilon} := \frac{\epsilon_0^2\tau}{10\log 1/\delta}$. For all $i \in [k]$, define $\mathcal{M}_i := \text{EffJor}_{\hat{\epsilon}, \hat{\delta}}[|0\rangle\langle 0|_{W_0^{\leq i}}, G_i]$. The algorithm **Amp-U** is defined as follows.

Algorithm 3. *Uniform adversary Amp-U for the 1-fold protocol.*

Input: registers $(A_j)_{1 \leq j \leq t}$, unitary oracles A, C, D , $\delta, \epsilon \in [0, 1]$, $k \in \mathbb{N}$.

1. For $j = 1, \dots, t$, initialise a fresh register $W_{0,j}, M_{0,j}$ to $|0\rangle$ and measure $A_j, W_{0,j}, M_{0,j}$ with \mathcal{M}_k , obtaining outcome γ_j .
2. Let j be such that $\gamma_j \geq \delta^k - \hat{\epsilon}$; if no such j exists, abort. Set $A := A_j, W_0 := W_{0,j}$.
3. For $i = k, \dots, 2$, apply $\text{StateTrans}_{\hat{\epsilon}, \hat{\delta}, \tau}[\mathcal{M}_i, \mathcal{M}_{i-1}, \delta^{i-1} - \hat{\epsilon}]$ to (A, W_0) , obtaining outcome c_i . If $c_i = \perp$, abort. If $c_i = 1$, stop and proceed to the next step.
4. If $c_i = 0$ for all $i = 2, \dots, k$, set $c_1 := 1$. Let i^* be the (unique) index for which $c_{i^*} = 1$.
5. Run $\text{Amp}_{\tau, \mu}$ (Algorithm 1) on input i^*, δ and registers $AW_0^{\geq i^*+1}M_0^{\leq k}$.

We first bound the probability that Algorithm 3 aborts. By Lemma 5.9 and assumption, for all j , $\mathbb{E}[p_j] = \text{Tr}(\tilde{G}_k \rho_j) \geq \delta^k$. Since $p_j \leq 1$, by an averaging argument we get $\Pr[p_j \geq \delta^k - \hat{\epsilon}] \geq \frac{\hat{\epsilon}}{1 - (\delta^k - \hat{\epsilon})} \geq \hat{\epsilon}$, and so the probability that Algorithm 3 aborts in step 2 is at most $(1 - \hat{\epsilon})^t \leq \epsilon_0/10$ if $t = \frac{1}{\hat{\epsilon}} \ln \frac{10}{\epsilon_0}$.

By Lemma 5.10, the probability that $c_i = \perp$ for any i is at most $4kK\sqrt{\hat{\delta}} = 4k\sqrt{\hat{\delta}} \cdot \lceil \frac{2}{\tau} \ln \frac{1}{\hat{\delta}} \rceil \leq \frac{9k}{\tau} \cdot \sqrt{\hat{\delta}} \ln \frac{1}{\hat{\delta}} \leq \frac{20k}{\tau} \cdot \sqrt[3]{\hat{\delta}} = \epsilon_0/10$.

By Lemma 5.10, the state ρ at the beginning of iteration i^* of Step 3 has $\Pr[\mathcal{M}_{i^*}(\rho) < \delta^{i^*} - 2\hat{\epsilon}] = O(K\delta) \leq \epsilon_0/10$. Then, again by Lemma 5.10, the state ρ' at the end of iteration i^* of Step 3 has $\mathcal{M}_{i^*}(\rho') \geq \delta^{i^*} - 2K\hat{\epsilon}$ with probability $1 - \epsilon_0/10 - \sqrt{\tau + \hat{\epsilon} + \hat{\delta}} \geq 1 - \epsilon_0/5$. It then follows by Lemma 5.9 that $\text{Tr}(\tilde{G}_{i^*} \rho') \geq \delta^{i^*} - 2K\hat{\epsilon} - \epsilon_0/5 \geq \delta^{i^*} - \epsilon_0/4$. We also have that $\mathcal{M}_{i^*-1}(\rho') \leq \delta^{i^*-1} - \hat{\epsilon}$ with probability $1 - \tau - \hat{\epsilon} - \hat{\delta}$, from which it follows that $\text{Tr}(\Pi_{>\delta^{i^*-1}}^{\tilde{G}_{i^*-1}} \rho') \leq \tau + \hat{\epsilon} + 2\hat{\delta} \leq \epsilon_0^2/16$.

Let $\sigma := \Pi_{>\delta^{i^*-1}}^{\tilde{G}_{i^*-1}} \rho' \Pi_{>\delta^{i^*-1}}^{\tilde{G}_{i^*-1}} / \text{Tr}(\Pi_{>\delta^{i^*-1}}^{\tilde{G}_{i^*-1}})$. By gentle measurement, $\text{td}(\rho', \sigma) \leq \epsilon_0/4$. By Lemma 4.6, applying **Amp** (Algorithm 1) to σ with parameters $\tau = 1 - \frac{\epsilon_0}{\delta^{i^*}} \geq 1 - \epsilon/4$, $\mu = \epsilon/4$ yields an adversary that succeeds with probability at least $(1 - \epsilon/4)^3 \delta \geq (1 - \epsilon)\delta \geq \delta - \epsilon$, and runs in time $\text{poly}(k, \frac{1}{\delta^k}, \frac{1}{\epsilon})$.

6 Barriers to parallel repetition beyond 3-message protocols

In this section we show that, under a cryptographic assumption, for every k there exists a constant round quantum interactive protocol such that the k -fold parallel repetition of the protocol has the *same* soundness as the original protocol. Before introducing the interactive quantum protocol, we need to define post-quantum bit commitment schemes and non-malleability.

6.1 Post-quantum bit commitments

We begin by describing post-quantum bit commitments similar to how they are described in [CCY21]. At a high level, interactive quantum bit commitments are quantum interactive protocols executed between a sender and receiver. Post-quantum bit commitments describe a special form of general quantum bit commitments where the protocol is entirely classical but security holds against quantum adversaries. Since the messages are entirely classical we will not use the quantum register notation used in the rest of the paper for the sake of simplicity. To talk about the security of post-quantum bit commitments, we must first define computational indistinguishability. We say that two families of random variables, indexed by λ , $\mathcal{X} = \{\mathcal{X}_\lambda\}$ and $\mathcal{Y} = \{\mathcal{Y}_\lambda\}$ are *computationally indistinguishable*, denoted $\mathcal{X} \approx_c \mathcal{Y}$ if for all polynomial-time quantum adversaries $(A_\lambda)_\lambda$,

$$\left| \Pr_{x \sim \mathcal{X}_\lambda} [A_\lambda(x) \text{ accepts}] - \Pr_{y \sim \mathcal{Y}_\lambda} [A_\lambda(y) \text{ accepts}] \right| \leq \text{negl}(\lambda).$$

For a post-quantum commitment scheme between a sender $A = (A_\lambda)_\lambda$ and receiver $B = (B_\lambda)_\lambda$, denote by $\text{OUT}_S(m, A, B)$ and $\text{OUT}_R(m, A, B)$ the random variables corresponding to the classical strings held in the private workspace registers at the end of the commit phase when the sender is committing to the message m , where the randomness is over the internal randomness of the algorithm. Let $\tau(m, A, B)$ be the public transcript (i.e. an ordered list of messages sent). We now define the hiding and binding properties of post-quantum commitments.

Definition 6.1 (Hiding property of post-quantum bit commitments). *Let A be the algorithm that performs an honest execution of the sender in a post-quantum bit commitment. The post-quantum bit commitment is computationally hiding if for all polynomial-time quantum adversaries $B = (B_\lambda)_\lambda$,*

$$\text{OUT}_R(0, A, B) \approx_c \text{OUT}_R(1, A, B).$$

At a high level, the hiding property means that any bounded receiver can not tell if they are receiving a commitment to 0 or 1.

Definition 6.2 (Binding property of post-quantum commitments). *Let OUT_R^b and τ^b be the private receiver workspace and public transcript after an execution of a post-quantum commitment scheme where the sender committed to b . Let $\text{ACCEPT}_\lambda(\text{OUT}_R^b, \tau^b, A)$ be the bit corresponding to whether or an honest receiver accepts in the reveal phase when the sender acts according to the algorithm A , conditioned on the receiver having private workspace distributed as OUT_R^b , and public transcript τ^b in the commit phase, and $\text{REVEAL}(\text{OUT}_R^b, \tau^b, A)$ be the bit that is revealed to the receiver in the same context. The post-quantum commitment scheme is binding if for all polynomial time non-uniform quantum adversaries A ,*

$$\Pr[\text{REVEAL}(\text{OUT}_R^b, \tau^b, A) = 1 - b \text{ and } \text{ACCEPT}(\text{OUT}_R^b, \tau^b, A) = 1] \leq \text{negl}(\lambda). \quad (13)$$




Figure 2: An c, r -message post-quantum bit commitment scheme. In a post-quantum bit commitment, all messages are classical, but the sender might be quantum.

At a high level, the binding property means that a bounded adversary can not switch their commitment from b to $1 - b$ after the commit phase.

Now we turn our attention to post-quantum interactive protocols with a specific type of security. The notion of security that we will care about is called *non-malleability* [DDN91]. At a high level, a post-quantum commitment scheme is non-malleable if any bounded man-in-the-middle adversary can not modify a commitment into another valid commitment in a systematic way (i.e. they may be able to output commitments to random bits, or forward commitments they receive). We say that the man-in-the-middle adversary receives commitments on the left and sends their commitments on the right. When discussing man-in-the-middle attacks against interactive protocols, there are two additional concepts we need to introduce: tags and schedules.

A *tag-based* commitment scheme is one where in addition to λ , the sender receives a classical tag $\mathbf{tg}_\lambda \in \{0, 1\}^{t(\lambda)}$, which we assume (w.l.o.g.) is contained in the sender’s initial workspace register (i.e. the sender starts with $|b\rangle \otimes |\mathbf{tg}_\lambda\rangle \otimes |0\rangle$), and the receiver’s private workspace register after the commit phase. For a fixed sequence of tags $\mathbf{tg} = \{\mathbf{tg}_\lambda\}_\lambda$, we require that the corresponding family of commitment schemes satisfy hiding and binding. At a high level, a tag is meant to prevent an adversary from forwarding communication from the left to the right, in the sense that we will make the definitions such that in order for an adversary to break the non-malleable property of a commitment scheme, we will require that the man in the middle uses different tags in the left and right commitments.

For man-in-the-middle adversaries, a schedule refers to the sequence of messages sent in the left and right that the adversary accepts. In this paper we define non-malleability with respect to a specific schedule, which we call **schedule**, defined by Figure 3. Certain post-quantum bit commitments might be secure against more general schedules, or even all schedules, but any scheme that is secure against this schedule will suffice for our result.

Definition 6.3 (Many-to-many synchronous non-malleable property of commitment scheme). *Let*

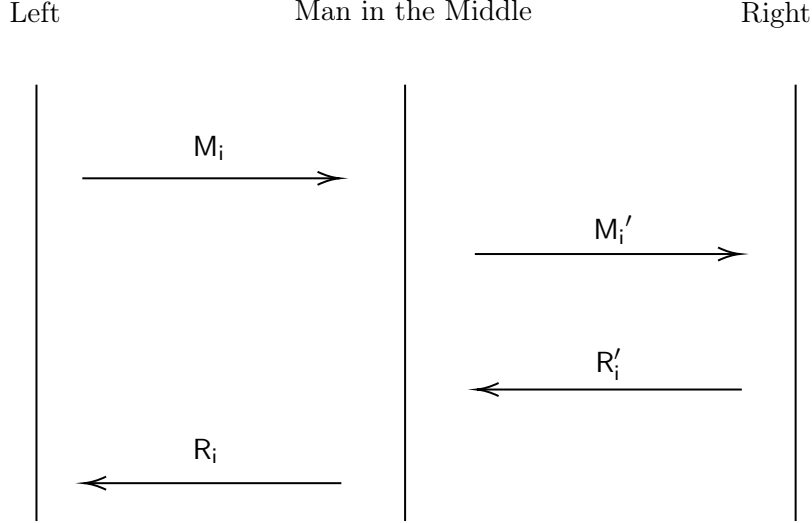


Figure 3: Schedule of messages sent in schedule. The schedule is chosen specifically so that an adversary might forward commitments to other challengers, in particular the left and right always send their messages *first*, before expecting a response. In [Algorithm 4](#), both the left and right are executed by the challenger, and the adversary executes the man-in-the-middle strategy.

\mathcal{C} be a post-quantum, tagged, commitment scheme. For a man-in-the-middle quantum adversary $A = (A_\lambda)_\lambda$, that receives l commitments on the left and outputs commitments to the l' bits on the right according to the schedule `schedule`, let $\text{mim}_{\mathcal{C}}^A(\lambda, x)$ denote the random variable over the private workspace of the adversary and the message bits it outputs, $(\text{OUT}_S(m_i, A, R), m)_{i \leq l'}$, where m is the l' bits that the adversary commits to on the right, where R is an honest execution of the receiver, whenever the l bits being committed to on the left are x . Let $\text{mim}_{\mathcal{C}}^A(\lambda, x) = \perp$ whenever the adversary uses the tag that the original commitment used. We assume that the adversary commits to each m_i individually using the bit commitment scheme. A post-quantum bit commitment scheme is one-to-many non-malleable if

$$\text{mim}_{\mathcal{C}}^A(\lambda, x_1) \approx_c \text{mim}_{\mathcal{C}}^A(\lambda, x_2)$$

for all message on the left, x_1 and x_2 .

The use of the word many-to-many indicates that the adversary can receive any number of commitment on the left, and commits to many bits on the right. Being non-malleable means no matter what joint distribution over bits the adversary commits to on the right, the state of the adversary and bits committed to on the right is indistinguishable when the left commits to 0 versus 1. We note that we define mim to take the value \perp whenever the adversary re-uses the tag to allow for the “un-interesting” case when the adversary does not try to tamper with the commitment.

6.2 Parallel repetition fails for 4-message quantum interactive protocols

We show that, if there exists a $(c+r)$ -message post-quantum interactive bit commitment scheme that satisfies one-to-many non-malleability, then the following protocol is a $2(c+r)$ -message quantum interactive protocol such that the soundness does not change after k repetitions. The protocol involves the challenger and adversary simultaneously making commitments to each other, and

then revealing their commitments, with the adversary winning if they can “flip” the challenger’s commitment. At a high level, the challenger and adversary will do the following: The challenger will begin by sending a message corresponding to the next message of commitment to $k - 1$ copies of a bit b from the challenger to the adversary. Before the adversary sends to response to this message, the adversary and challenger will compute and send (in parallel) the next message of $k - 1$ commitments from the adversary to the challenger. After doing this, the adversary will send their response to the challenger’s initial commitment. Formally, assume that a $(c + r)$ -message commitment scheme exists and consider the following protocol:

Algorithm 4. *k*-fold unrepeatable challenger

Input: Security parameter λ , tag tg .

1. Sample a bit b uniformly at random.
2. Repeat $\lceil c/2 \rceil$ times:
 - (a) The challenger sends the next message in $k - 1$ commitments to b .
 - i. The adversary sends the next message in the commitments to $k - 1$ many bits.
 - ii. The challenger sends their response to the $k - 1$ commitments from the adversary.
 - (b) The adversary sends the responses to the challengers commitments to b .
3. Check that the adversary does not use the tag tg . If not, reject.
4. Repeat $\lceil r/2 \rceil$ times:
 - (a) The challenger sends the next message in $k - 1$ reveals to b .
 - i. The adversary sends the next message in the reveal to $k - 1$ many bits.
 - ii. The challenger sends their response to the $k - 1$ reveals from the adversary.
 - (b) The adversary sends the response to the challengers reveals to b .
5. Let $\{c_i\}_{i=0}^{k-1}$ be the bits that the adversary revealed. Accept if
 - (a) $\oplus_{i=0}^{k-1} c_i \neq b$ and
 - (b) Every commitment sent by the adversary is accepted.

The use of the doubly indented bullet points is meant to highlight that steps 2a and 2b correspond to an execution of the commitment scheme, and steps 2(a)i and 2(a)ii correspond to another execution of the commitment scheme, interwoven with the first. Note that in the final iteration of each loop, the response might be an empty message (for example, if the final message the commitment or reveal stages has already been sent). With this, it is clear that Algorithm 4 is a $2(c + r)$ -message protocol if the commitment scheme involves exchanging c messages in the commit stage and r messages in the reveal stage. We will show that both Algorithm 4 and the k -fold parallel repetition of Algorithm 4 have soundness $1/2$, assuming that the commitment scheme satisfies certain properties.

Lemma 6.4. *Assume that \mathcal{C} is a post-quantum commitment scheme satisfying one-to-many non-*

malleability. Then for every k , and every polynomial-time quantum adversary A , there exists a negligible function ϵ such that the probability that A is accepted by the challenger executing [Algorithm 4](#) is at most

$$\frac{1}{2} + \epsilon(\lambda). \quad (14)$$

Proof. Assume for the sake of contradiction that there is an adversary $(A_\lambda)_\lambda$ that is accepted by the challenger with advantage $1/p(\lambda)$ for some polynomial $p(\cdot)$. We are going to construct an adversary for the non-malleability property of \mathcal{C} as follows: The challenger on the left commits to the string b^{k-1} using $k - 1$ copies of the commitment scheme (here b can be either 0 or 1). The adversary then runs A on the commitments from the left to receive a commitment to a string of length $k - 1$, denoted \tilde{m} . By step 3, we can assume that all of the commitments in used to commit to \tilde{m} use a different tag than the challenger's.

By the assumption about A , \tilde{m} has parity $1 - b$ with probability $1/2 + 1/p(\lambda)$. Thus, with probability $1/2 + 1/p(\lambda)$, the message string in $\text{mim}_{\mathcal{C}}^A(0^{k-1}, \lambda)$ has XOR 1, and the message string in $\text{mim}_{\mathcal{C}}^A(1^{k-1}, \lambda)$ has XOR 0. Thus, the two distributions are distinguishable with non-negligible advantage by a polynomial-time adversary that accepts if the messages sent to the right have even parity. Thus, if there is an adversary that succeeds in [Algorithm 4](#) with non-negligible advantage, \mathcal{C} is not many-to-many non-malleable, a contradiction. \square

Next we show that there is a adversary that is accepted with probability $1/2$ in the k -fold parallel repetition of [Algorithm 4](#). At a high level, for each challenger in the k -fold parallel repetition, the adversary will forward commitments from all $k - 1$ other challengers to them. If the XOR of all of the committed bits is equal to 1, every bit that the adversary commits to will be $b \oplus 1$, and the adversary will win every game in the k -fold repetition.

Lemma 6.5. *For every commitment scheme \mathcal{C} , there exists an adversary that accepted with probability $1/2$ in the k -fold parallel repetition of [Algorithm 4](#), when all k repetitions are given unique tags tg .*

Proof. In step 2a of [Algorithm 4](#), the challenger sends $k - 1$ messages to the adversary. In the k -fold parallel repetition, we will denote by $m_i^{i,\alpha}$ the α^{th} commitment sent from the i^{th} repetition of the protocol, and similarly for responses. Consider the following adversary for the k -fold parallel repetition of [Algorithm 4](#).

Algorithm 5. Adversary for k -fold repetition of Algorithm 4.**Input:** Security parameter λ .

1. For l from 0 to $\lfloor c/2 \rfloor$:
 - (a) Receive k messages, $\{m_l^{0,\alpha}\}_{\alpha \in [k-1]}, \dots, \{m_l^{k-1,\alpha}\}_{\alpha \in [k-1]}$ from the challengers.
 - i. To every challenger i , send the $k-1$ messages $\{m_l^{j+i \bmod k, j+1 \bmod k}\}_{j \in [1:k]}$.
 - ii. From challenger i , receive $k-1$ responses, $\{r_l^{j+i \bmod k, j+1 \bmod k}\}_{j \in [1:k]}$ from the challengers.
 - (b) To every challenger i , send the $k-1$ responses $\{r_l^{i,\alpha}\}_{\alpha \in [k-1]}$.

(The challengers reveal bits b_0, \dots, b_k . To each challenger i the adversary reveal $\{b_j\}_{j \neq i}$.)
2. For l from 0 to $\lfloor r/2 \rfloor$:
 - (a) Receive k messages, $\{m_l^{0,\alpha}\}_{\alpha \in [k-1]}, \dots, \{m_l^{k-1,\alpha}\}_{\alpha \in [k-1]}$ from the challengers.
 - i. To every challenger i , send the $k-1$ messages $\{m_l^{j+i \bmod k, j+1 \bmod k}\}_{j \in [1:k]}$.
 - ii. From challenger i , receive $k-1$ responses, $\{r_l^{j+i \bmod k, j+1 \bmod k}\}_{j \in [1:k]}$ from the challengers.
 - (b) To every challenger i , send the $k-1$ responses $\{r_l^{i,\alpha}\}_{\alpha \in [k-1]}$.

We first note that Algorithm 5 runs in polynomial time and always produces valid commitments to the bits that it announces, and assuming that all of the challengers used unique tags, every challenger sees commitments that use unique tags.

Let $B = \bigoplus_{j=0}^{k-1} b_j$ be the XOR of all of the bits that the challengers committed to. Because the adversary announces to every challenger, j , the bits of the other challengers (not including j), the adversary announces bits that XOR to $\left(\bigoplus_{l \neq j} b_l\right) = b_j \oplus B$, to the j^{th} repetition of the protocol. If $B = 1$, Algorithm 5 wins all of the repetitions of Algorithm 4. Since every challenger samples a bit uniformly at random, this happens with probability $1/2$. Thus the adversary is accepted in the k -fold repetition of Algorithm 4 with probability $1/2$. \square

If $c = r = 1$, then we say that a quantum bit commitment is *non-interactive*. Classical non-interactive, non-malleable bit commitments are known to exist in a number of settings and from a wide variety of assumptions [LPS20], and there is no formal evidence ruling out the existence of post-quantum (or fully quantum) non-interactive non-malleable bit commitments. Under this assumption, combined with the previous claims, we have the following corollary.

Corollary 6.6. *If there exists a non-interactive post-quantum bit commitment scheme satisfying one-to-many non-malleability, then there is a 4-message quantum interactive protocol for which parallel repetition fails.*

Thus, it can not be the case that both non-interactive one-to-many non-malleable commitments exist, and that strong parallel repetition, where the soundness for the k -fold parallel repetition

goes down as ϵ^k , holds 4-message quantum interactive protocols. If one believes that these kinds of commitments do exist, then the 3-message parallel repetition theorem proved here is tight. However this argument only shows that for every fixed k , there exists a protocol for which soundness of the (exactly) k -fold parallel repetition stays high. It could still be the case that for every 4-message protocol, there exists a large k for which the soundness of the k -fold parallel repetition of the protocol has negligible soundness. We leave the task of ruling out this (weaker) form of parallel repetition as an open question.

7 Round compression for quantum argument systems

A quantum interactive *argument* is a special type of quantum interactive proof where the completeness and soundness conditions hold with respect to computationally efficient provers. Formally, a promise decision problem $L = (L_{yes}, L_{no})$ admits a quantum interactive argument with completeness c and soundness s if there exists a polynomial-time quantum verifier V and a polynomial-time quantum “honest” prover P such that

1. For all $x \in L_{yes}$ there exists a quantum advice state $|\psi_x\rangle$ such that P given $|\psi_x\rangle$ interacting with V on input x gets accepted with probability at least c ;
2. For all polynomial-time “malicious” provers P^* , for all $x \in L_{no}$, for all quantum advice states $|\psi_x\rangle$, the prover P^* interacting with V on input x gets accepted with probability at most s .

In other words, the quantum protocol that the verifier V engages in is s -computationally secure when V is given “no” instances from L_{no} .

In this section we prove a *round compression* result for quantum interactive arguments. It was proved by Kitaev and Watrous in [KW00] that every m -round quantum interactive *proof* (i.e. where the prover has unbounded computational power) for a (promise) language L can be *compressed* into another interactive proof for L that only has 3 messages and the soundness is worsened by a factor $\text{poly}(m)$. By parallel repetition, the soundness can be improved back to a constant [KW00]. This stands in contrast to the classical case, where an analogous round compression technique for interactive proofs is unknown and considered unlikely.

We prove an analogue of the Kitaev–Watrous compression procedure for interactive arguments:

Theorem 7.1 (Round compression for quantum interactive arguments). *Let L be a promise language with an $m(n)$ -message quantum interactive argument with completeness c and soundness s for $m \geq 3$. Then there exists a 3-message quantum interactive argument for L with completeness $1 - \frac{2(1-c)}{m-1}$ and soundness $1 - \frac{1-s}{(m-1)^4}$. Furthermore, if the original prover has complexity t_P and verifier has complexity t_V , then the compressed protocol has prover complexity $m^{O(1)}(t_P + t_V)$ and verifier complexity $m^{O(1)}t_V$.*

We follow the proof approach of Kempe et al. [KKMV07] (see also [VW16]) who gave an iterative procedure to round-compress quantum proof systems (whereas the original Kitaev-Watrous proposal achieved the compression in one step). The idea is as follows: given a $(2r + 1)$ -message “original” protocol, one can obtain an $(r + 1)$ -message “compressed” protocol where the prover first sends the intermediate state of the verifier at round $r + 1$ of the original protocol (i.e., the midpoint). The verifier in the compressed protocol randomly decides whether to play the original protocol forwards or backwards to check respectively whether the original verifier would have accepted, or

was initialized properly. Our protocol is similar in spirit to this reduction, except that we must slightly tweak the reduction and analysis to ensure that the reduction is efficient.

Remark 7.2. In the quantum interactive proof setting, the assumption of perfect completeness on the original protocol is not necessary, because every quantum interactive proof can be transformed to have perfect completeness (see [KW00] and [VW16]). However we do not know whether quantum arguments can be generically transformed to have perfect completeness.

Remark 7.3. To maintain consistency with existing work regarding interactive proofs and arguments, in this section we adopt the terminology “verifier”, “honest prover”, and “adversary” (or malicious prover). The verifier corresponds to the challenger in an interactive argument. When talking about completeness, the honest prover corresponds to the adversary, and when talking about soundness we still use the term adversary.

Remark 7.4. When we refer to the *total run-time* of an efficient adversary in an interactive argument, as the total time it takes to execute the entire interaction between the adversary and the challenger. If every action the adversary takes is polynomial time, then the total run-time will be a polynomial in n as well.

We first describe a “round-halving” compression procedure that transforms a verifier for a $(2r + 1)$ -message protocol to a verifier for an equivalent $(r + 1)$ -message protocol (in the completeness and soundness sense). Our reduction adds at most a constant overhead to the gate complexity of the verifier. At the end, we will iterate this procedure logarithmically many times to obtain a 3-message protocol.

Let $\{C_i\}_{i \leq r}$ be a verifier in a $(2r + 1)$ -message quantum interactive protocol. We describe a $(r + 1)$ -message protocol and analyse its completeness and soundness. We briefly recall the notation used to describe quantum interactive arguments. A_i denotes the adversary’s private workspace in round i , M_i denotes the message register sent from the adversary to the challenger, and R_i denotes the response register sent from the challenger to the adversary. The adversary applies the unitary A_i on $A_i R_{i-1}$ to get a state on $A_{i+1} M_i$, and the verifier applies a unitary C_i on $W_i M_i$ to get a state on $W_{i+1} R_i$. In the final round, the challenger applies a POVM $\{D, \text{id} - D\}$ on $W_r M_r$ to decide whether to accept or reject.

Algorithm 6. Verifier for $(r + 1)$ -message compressed protocol

1. Receive registers $M_{r/2}W_{r/2}$ from the adversary. Flip a unbiased coin.
2. If the outcome of the coin flip is heads:
 - (a) Run the original $2r + 1$ -message protocol starting from interaction $r/2$, i.e. apply unitary $C_{r/2}$ to registers $M_{r/2}W_{r/2}$, send the result of the coin flip and $R_{r/2}$ to the adversary, and continue for interactions $r/2 + 1$ through r .
 - (b) After $r/2$ many interactions, receive register M_r , measure Π on registers M_rW_r . Accept if the measurement accepts.
3. If outcome of the coin flip is tails:
 - (a) Run the original $2r + 1$ -message protocol in reverse, i.e. send the result of the coin flip and $M_{r/2}$ back to the adversary, receive register $R_{r/2-1}$ from the adversary, apply $(C_{r/2-1})^\dagger$ to $R_{r/2-1}W_{r/2}$, and continue for interactions $r/2 - 1$ through 1.
 - (b) After $r/2$ many interactions, receive register R_1 . Apply $(C_1)^\dagger$ on registers R_1W_2 to get registers M_0W_0 , and measure $|0\rangle\langle 0|_{W_0}$. Accept if the measurement accepts.

Claim 7.5 (Completeness). *If there is an non-uniform (resp. uniform) honest prover (denoted the original honest prover) that succeeds in the original $(2r + 1)$ -message protocol with probability $1 - \epsilon$, then there is a non-uniform (resp. uniform) honest prover that succeeds in the $(r + 1)$ -message compressed protocol with probability $1 - \epsilon/2$. Furthermore, if the original protocol has running time t_P, t_V for the prover and the verifier, then the new protocol has running time $O(t_P + t_V)$ for the prover and $O(t_V)$ for the verifier.*

Proof. Let A_i be unitary operations that the honest prover implements in the i 'th interaction. Specifically, the honest prover applies A_i to registers $R_{i-1}A_i$ to get a message register M_i and new private register A_{i+1} , and then sends M_i to the verifier. Initially, R_0 is an empty register, and A_0 is initialized to some state $|\text{aux}\rangle\langle \text{aux}|_{A_0}$ (in the non-uniform case this might be an advice state, otherwise $|0\rangle\langle 0|$).

The honest prover for the compressed protocol will essentially implement the honest strategy, applying the unitaries of the un-compressed honest prover forward or backwards depending on the outcome of the verifier's coin flip. The honest prover begins by performing preparing a register W_0 in the state $|0\rangle_{W_0}$, and then performs $A_0, C_0, A_1, C_1, \dots, A_{r/2}$ to the appropriate registers to get a state over registers $A_{r/2+1}M_{r/2}W_{r/2}$. The honest prover then sends $M_{r/2}W_{r/2}$ to the challenger as their initial message.

If the outcome of the verifier's coin flip is heads, the honest prover applies $A_{r/2+1}$ to the verifier's response, and continues as the un-compressed honest prover would starting from round $(r/2) + 1$. After all $r/2$ interactions, the verifier and honest prover hold the state after implementing the unitaries A_0, C_0, \dots, A_r on the initial state $|\text{aux}\rangle\langle \text{aux}|_{A_0} \otimes |0\rangle\langle 0|_{W_0}$ and then measuring D on registers M_rW_r . Since the original protocol accepted with probability $1 - \epsilon$, we have that

$$\text{Tr}(D(A_r C_{r-1} \dots C_0 A_0)(|\text{aux}\rangle\langle \text{aux}|_{A_0} \otimes |0\rangle\langle 0|_{W_0})(A_0^\dagger C_0^\dagger \dots C_{r-1}^\dagger A_r^\dagger)) = 1 - \epsilon.$$

Thus, the probability that the honest prover is accepted is also $1 - \epsilon$.

If the outcome of the verifier's coin flip is tails, when the honest prover receives register M_i , the honest prover applies $(A_i)^\dagger$ to $M_i A_{i+1}$ and sends registers R_{i-1} to the verifier, and continues by applying the inverse of the un-compressed honest prover unitary, going backwards in round number. When the protocol reaches the final message, the honest prover and verifier will have performed the inverse of $C_0, A_1, C_1, \dots, A_{r/2}$ to yield some state $|\psi\rangle_{M_0 A_1} \otimes |0\rangle_{W_0}$. When the verifier measures $|0\rangle\langle 0|$, the measurement will accept with probability 1.

So with probability $1/2$, the honest prover is accepted with probability $1 - \epsilon$, and with probability $1/2$ the honest prover is accepted with probability 1. Thus, the overall probability that the honest prover is accepted is $1 - \epsilon/2$.

To analyze the runtime of the honest prover, note that the honest prover for the compressed protocol applies every unitary A_i at most twice, and runs the first verifier's unitaries C_i at most once. The verifier runs every C_i at most once, so together they run every A_i and C_i at most twice, so its total running time of the protocol is at most twice the total running time of the original honest prover. We note that the compressed honest prover can be implemented by controlling the A_i and C_i on the message register (so that they act correct controlled on the outcome of the verifier's coin flip). If the honest prover can only use 2-qubit gates, this might result in a constant multiplicative overhead to the gate complexity of the honest prover. \square

Claim 7.6 (Soundness). *If there is an adversary that has total run-time t that succeeds in the $(r + 1)$ -message compressed protocol with probability $1 - \epsilon$, then there is an adversary that has total run-time $O(t + t_V)$ and succeeds in the original $(2r + 1)$ -message protocol with probability $1 - 16\epsilon$.*

Proof. We can assume that $\epsilon \leq 1/16$, otherwise the claim is trivially true. Assume that the adversary for the $(r + 1)$ -message compressed protocol implements an initial unitary A_0 on register A_0 , which is initialized in state $|\text{aux}\rangle_{A_0}$. Call the output registers of this unitary $R_{r/2+1} M_{r/2} W_{r/2}$, of which $M_{r/2} W_{r/2}$ are sent to the challenger as the adversary's first message. After seeing the result of the coin flip, $b \in \{H, T\}$, from the challenger, we denote by A_i^b the unitaries that the adversary implements in interactions $i = 1$ through $r/2$. We say that the unitary A_i^H acts on registers $R_{r/2+i-1} A_{r/2+i}$ (so that the first one acts on registers $R_{r/2} A_{r/2+1}$ and counts up from there) and the unitary A_i^T acts on registers $M_{r/2-i+1} A_{r/2-i+2}$ (so that the first one acts on registers $M_{r/2} A_{r/2+1}$). The challenger's actions in the game are specified in [Algorithm 6](#). Define the following projectors.

$$\begin{aligned} \Pi_H &= (A_0^\dagger C_{r/2}^\dagger A_1^{H\dagger} C_{r/2+1}^\dagger \dots A_{r/2}^{H\dagger}) D_{M_r W_r} (A_{r/2}^H \dots C_{r/2+1}^H A_1^H C_{r/2}^H A_0), \\ \Pi_T &= (A_0^\dagger A_1^{T\dagger} C_{r/2-1}^\dagger \dots A_{r/2}^{T\dagger} C_0^\dagger) (\text{id} \otimes |0\rangle\langle 0|_{W_0}) (C_0^\dagger A_{r/2}^T \dots C_{r/2-1}^\dagger A_1^T A_0). \end{aligned}$$

It is clear that $\text{Tr}(\Pi_H |\text{aux}\rangle\langle \text{aux}|)$ and $\text{Tr}(\Pi_T |\text{aux}\rangle\langle \text{aux}|)$ are the probabilities that the adversary is accepted when the result of the challenger's coin flip is H and T respectively. By assumption, the adversary is accepted in the $r + 1$ -message compressed protocol with probability $1 - \epsilon$, so the adversary must be accepted in the protocol if the coin flip is fixed to being either heads or tails with probability at least $1 - 2\epsilon$. Therefore we have the inequalities

$$\begin{aligned} \text{Tr}(\Pi_H |\text{aux}\rangle\langle \text{aux}|) &\geq 1 - 2\epsilon, \\ \text{Tr}(\Pi_T |\text{aux}\rangle\langle \text{aux}|) &\geq 1 - 2\epsilon. \end{aligned}$$

Now we construct an adversary for the original $(2r + 1)$ -message protocol as follows:

Algorithm 7. Adversary for the $(2r + 1)$ -message protocol

Input: Initial state $|\text{aux}\rangle$.

1. Run A_0 to receive a state on registers $A_{r/2+1}M_{r/2}W_{r/2}$. Implement unitaries $A_1^T, C_{r/2-1}^\dagger, \dots, A_{r/2}^T, C_0^\dagger$ to receive a state on registers M_0W_0 . Measure $\text{id} \otimes |0\rangle\langle 0|_{W_0}$, if the measurement fails, abort. If it accepts, send register M_0 to the verifier.
2. For steps $i = 1$ through $r/2$:
 - (a) Run $(A_{r/2-i}^T)^\dagger$ to get the next message register and send it to the verifier.
3. For steps $i = r/2 + 1$ through r :
 - (a) Run $A_{i-r/2}^H$ to get the next message register and send it to the verifier.

Let ρ_T be the state of the adversary after measuring Π_T on $|\text{aux}\rangle\langle\text{aux}|$ and accepting. Let ρ_H be the state after measuring Π_H on $|\text{aux}\rangle\langle\text{aux}|$ and accepting. By the gentle measurement lemma (Proposition 3.3) and the definition of the squared Bures distance, we have that the the following inequalities

$$\begin{aligned} d_{\text{Bures}}(\rho_T, |\text{aux}\rangle\langle\text{aux}|) &\leq 2(1 - \sqrt{1 - 2\epsilon}) \leq 3\epsilon, \\ d_{\text{Bures}}(\rho_H, |\text{aux}\rangle\langle\text{aux}|) &\leq 3\epsilon. \end{aligned}$$

Here the first inequality holds for any $0 \leq \epsilon \leq \frac{4}{9}$. Using the weak triangle inequality for the squared Bures distance (Proposition 3.5), we have that

$$d_{\text{Bures}}(\rho_H, \rho_T) \leq 12\epsilon.$$

By the definition of the squared Bures distance, we have that

$$F(\rho_H, \rho_T) \geq (1 - 6\epsilon)^2 \geq 1 - 12\epsilon, \tag{15}$$

where we applied the Bernoulli inequality here. Now we examine the state of the adversary in the un-compressed game. If the un-compressed adversary does not abort in step 1, after step 2, the adversary is left with exactly ρ_T , because they condition on $|0\rangle\langle 0|_{W_0}$ accepting. At the end of the protocol, the challenger will have implemented the measurement Π_H on the adversary's state after step 2, ρ_T . Since Π_H is a projector and ρ_H is the post-measurement state of some initial state after measuring Π_H , we have that $\text{Tr}(\Pi_H \rho_H) = 1$. Applying Proposition 3.4 and Equation (15), we have that $\text{Tr}(\Pi_H \rho_T) \geq F(\rho_H, \rho_T) \geq 1 - 16\epsilon$. Thus, the adversary is accepted with probability at least $1 - 16\epsilon$. By assumption, the first measurement in the step succeeds with probability $(1 - 2\epsilon)$, thus the total probability of failure is given by

$$2\epsilon + (1 - 2\epsilon) \cdot 12\epsilon \leq 16\epsilon.$$

The un-compressed adversary applies every the unitaries A_i^T (and $(A_i^T)^\dagger$) two times, the unitaries C_i^\dagger once, and the unitaries A_i^H once. Thus the total run time of the adversary is $O(t + t_V)$ accounting for constant multiplicative overheads in implementing controlled unitaries. \square

Proof of Theorem 7.1. To prove the theorem, we iterate the round-halving procedure. Let $L = (L_{yes}, L_{no})$ be a (promise) decision language with a quantum interactive proof π where for inputs of length n , the protocol has $m(n)$ messages. We assume without loss of generality that $m(n)$ is of the form $2^{k(n)} + 1$; the adversary and challenger can send empty messages at the beginning of the protocol. In other words, $k = \lceil \log(m - 1) \rceil$ and $2^{k-1} < m - 1 \leq 2^k$.

Let π' denote protocol that applies the round-halving procedure described above $k - 1$ times. In other words, when the challenger receives input x , it first computes the description of the challenger for the $(2^{k-1} + 1)$ -message protocol (where $n = |x|$), and then based on this computes the description of the challenger for the $(2^{k-2} + 1)$ -message protocol, and so forth, until it obtains a challenger for a 3-message protocol. Since $m(n)$ is polynomial in n , $k = O(\log n)$. The final honest prover runs in time $m^{O(1)}(t_P + kt_V)$ and the verifier runs in time $m^{O(1)}t_V$, where the $O(1)$ factors here are exactly the constant multiplicative overhead above.

We now analyze the completeness and soundness of the 3-message protocol. Given a “yes” instance $x \in L_{yes}$, by definition there is an adversary that is accepted with probability $c = 1 - \delta$. By Claim 7.5, there is an adversary for the $2^{k(n)-1} + 1$ -message protocol that is accepted with probability $1 - \delta/2$. Repeating the compression $k - 1$ times, there exists an adversary that is accepted by the 3-message protocol with probability

$$1 - \frac{\delta}{2^{k-1}} \geq 1 - \frac{2\delta}{m(n) - 1}.$$

Given a “no” instance $x \in L_{no}$, the original protocol accepts with probability at most $s = 1 - \epsilon$ for all efficient adversaries. Then by Claim 7.6, the soundness of the $2^{k(n)-1} + 1$ -message protocol is at most $1 - \epsilon/16$. Similarly the soundness of the $2^{k(n)-1} + 1$ -message protocol is at most $1 - \epsilon/16^2$. Iterating this we see that the 3-message protocol has soundness at most

$$1 - \frac{\epsilon}{16^{k-1}} \leq 1 - \frac{\epsilon}{(m - 1)^4}.$$

We note that all of the reductions are efficient, even after $k - 1$ iterations of round collapse. Specifically, for “no” instances, if there is an adversary for the compressed protocol that has total run-time $t(n)$, then there is an adversary for the original protocol that has total run-time $m^{O(1)}(t + kt_V)$, which is efficient. \square

We can slightly extend our round compression to compile any 3-message quantum interactive argument into a public coin argument using the same strategy as [MW05, Theorem 5.4]. A *public coin* quantum interactive argument is a quantum interactive argument where all of the challenger messages are uniformly random coin flips. At a high level, if we were to apply round compression starting with a 3-message protocol, the challenger only ever needs to receive one register (either R_0 to go backwards, or M_1 to go forwards) from the adversary, so they do not need to send back a quantum register as in round compression for $m > 3$. Formally, we have the following.

Theorem 7.7 (Compilation to public coin). *Let L be a promise language with a 3-message quantum interactive argument with completeness c and soundness s . Then there exists a 3-message public coin quantum interactive argument for L with completeness $1 - \frac{1-c}{2}$ and soundness $1 - \frac{1-s}{16}$.*

Proof. Let (C_0, Π) be a challenger for a 3-message quantum interactive protocol. Consider the following challenger for a 3-message protocol.

Algorithm 8. Challenger for 3-message public coin protocol

1. Receive registers W_1 from the adversary. Flip a unbiased coin and send it to the adversary.
2. If the outcome of the coin flip is heads:
 - (a) The adversary sends register M_1 .
 - (b) Measure Π on M_1W_1 , accept if the measurement accepts.
3. If outcome of the coin flip is tails:
 - (a) The adversary sends register R_0 .
 - (b) Perform C_0^\dagger , and measure $|0\rangle\langle 0|$ on registers W_0 , accept if the measurement accepts.

This protocol is clearly public coin, as the challenger only sends the outcome of a single coin flip to the adversary. An honest adversary runs the original 3-message protocol up to the first two messages and sends the challenger’s private workspace to the challenger. Upon seeing heads they perform the honest adversary unitary and send M_1 , and upon seeing tails they send back R_0 to the challenger. A similar argument as [Claim 7.5](#) shows that the completeness behaves the same way as applying a single round of round compression.

Similarly, given an adversary for the public coin protocol, a malicious adversary for the original protocol can simulate the interaction between the public coin adversary and the challenger, and measure the challenger’s private workspace to get $|0\rangle\langle 0|_{W_0}$, and then get a suitable state on M_0 to send to the challenger. From there they run the public coin adversary, conditioned on seeing tails, in reverse for one step, and then the public coin adversary, conditioned on seeing heads, forward for one step. The same argument as in [Claim 7.6](#) shows that the soundness of the public coin protocol is equivalent to having applied round compression one additional step.

The run-time of the adversary in both directions is multiplied by a constant factor in the reduction, so the reduction is efficient. \square

Finally, we combine [Theorems 7.1](#) and [7.7](#) with our hardness amplification result to show that all polynomial-message quantum interactive arguments can be compressed into 3-message quantum interactive arguments with negligible soundness.

Corollary 7.8 (Parallelization and amplification for quantum interactive arguments). *Let L be a (promise) language that has a polynomial-message quantum interactive argument with completeness $c \geq 1 - \text{negl}(n)$ (or 1, respectively) and soundness $s \leq 1 - \frac{1}{\text{poly}(n)}$. Then there is a 3-message public coin quantum interactive argument for L with completeness $1 - \text{negl}(n)$ (or 1, respectively) and soundness $\text{negl}(n)$.*

We remark that this parallel repetition is not as efficient as what is possible classically. Classically, to amplify any interactive argument (say of soundness $\frac{1}{2}$ to soundness $\frac{1}{4}$) while preserving round complexity r , the state of the art incurs a multiplicative cost of order either r [[BHT20](#)] or λ [[CL10](#)], where λ is a security parameter. Turning back to quantum interactive arguments, we note that we can first round collapse the argument into three messages (if it has more than three messages),

which only incurs a constant multiplicative overhead, and then apply the three-message soundness amplification. In the end, the cost of the overall compiled protocol is $\Omega(r^4)$, since we need to make up for the loss in soundness in the round collapse theorem. A more careful analysis of [Claim 7.6](#) could improve the exponent from $\log 16 = 4$ down to $\log 10 \approx 3.32$ but we leave as future work to further improve the amplification efficiency.

8 Applications

The parallel repetition theorem has many immediate applications in cryptography, from boosting the security of commitments, to round reduction for zero-knowledge proofs.

8.1 Strong amplification of quantum bit commitment schemes

Here we show that the security of canonical quantum bit commitment schemes can be amplified through parallel repetition. While we have already defined bit commitments and notions of post-quantum security, here we review *canonical* quantum bit commitment schemes, and the standard notions of security they are described by, honest hiding and binding

Canonical quantum bit commitments. A canonical quantum bit commitment is a kind of non-interactive bit commitment scheme. The scheme consists of two quantum circuits C_0, C_1 , to commit to be a bit b , the sender generates the bipartite pure state $|\psi_b\rangle_{\text{CR}} = C_b |0 \dots 0\rangle$ and sends the C register of $|\psi_b\rangle$. To reveal, the sender sends R to the receiver, along with the bit b . The receiver can then verify the original bit b by applying the circuit C_b^\dagger and measuring $|0 \dots 0\rangle$. Importantly, Yan [\[Yan22\]](#) showed that all quantum commitment schemes can be compiled to this canonical form while preserving honest hiding and binding security, and honest security for canonical form commitments is equivalent to the stronger standard security notions. Throughout this section we focus on canonical quantum commitment schemes.

The honest hiding property of a commitment scheme guarantees that an adversarial receiver can not reveal the bit committed to before the reveal phase. Formally,

Definition 8.1 (Honest hiding property of commitment scheme). *Let $\epsilon(\lambda)$ denote a function. We say that a commitment scheme $(\pi_\lambda)_\lambda$ satisfies ϵ -computational (resp. ϵ -statistical) honest hiding if for all non-uniform polynomial-time algorithms (resp. for non-uniform algorithms) $A = (A_\lambda)_\lambda$ that take as input the receiver's register immediately after an honest execution of the commit stage of π_λ , the following holds for sufficiently large λ :*

$$\left| \Pr[A_\lambda(\rho_{\lambda,0}) = 1] - \Pr[A_\lambda(\rho_{\lambda,1}) = 1] \right| \leq \epsilon(\lambda).$$

Here, $\rho_{\lambda,b}$ denotes the reduced density matrix of the receiver's register after a honest execution of the commit stage of π_λ when the sender is committing to b . If ϵ is a negligible function of λ then we simply say that the scheme satisfies strong computational (resp. statistical) hiding. If $\epsilon(\lambda) \leq 1 - \frac{1}{p(\lambda)}$ for some polynomial $p(\lambda)$ we say it satisfies weak computational (resp. statistical) hiding.

The honest binding property, at a high level, says that after the commit phase of a commitment scheme, an adversarial sender can only reveal to the bit that they committed to. Formally,

Definition 8.2 (Honest binding property of commitment scheme). *Let $\epsilon(\lambda)$ denote a function. We say that a commitment scheme $(\pi_\lambda)_\lambda$ satisfies ϵ -computational (resp. ϵ -statistical) honest binding if for all non-uniform polynomial-time algorithms (resp. for all non-uniform algorithms) $A = (A_\lambda)_\lambda$ that take as input the sender's register immediately after an honest execution of the commit stage of π_λ , the following holds for sufficiently large λ :*

$$F\left(\left(A_\lambda \otimes \text{id}_C\right)(\psi_{\lambda,0}, \psi_{\lambda,1})\right) \leq \epsilon(\lambda),$$

where $\psi_{\lambda,b}$ denotes the state of the joint sender-receiver system after an honest execution of the commit stage of π_λ when the committed bit is b .

If ϵ is a negligible function of λ then we simply say that the scheme satisfies strong computational (resp. statistical) honest binding. Otherwise if $\epsilon(\lambda) \leq 1 - \frac{1}{p(\lambda)}$ for some polynomial $p(\cdot)$ we say that it satisfies weak computational (resp. statistical) honest binding.

Parallel repetition of quantum bit commitments. We show that parallel repetition can be used to amplify weak computationally binding (and strong statistically hiding) commitments into commitments where both hiding and binding are strong.

Corollary 8.3. *Let $\{C_{\lambda,b}\}_{\lambda,b}$ be a canonical commitment scheme satisfying the $\left(1 - \frac{1}{p(\lambda)}\right)$ -computational honest binding property for some polynomial $p(\cdot)$, and strong statistical hiding. Then there exists a polynomial $q(\cdot)$ such that the commitment scheme $\{C_{\lambda,b}^{\otimes q(\lambda)}\}_{\lambda,b}$ satisfies the strong computational honest binding property and strong statistical hiding property.*

Proof. The computational honest binding property of (canonical) quantum commitments can be equivalently viewed as the maximum winning probability of the following 2-message game (over efficient adversaries):

1. The challenger commits to 0, and sends the *reveal* register to the adversary.
2. The adversary sends back a reveal register.
3. The challenger accepts if applying $C_{\lambda,1}$ and measuring in the computational basis yields $|0^\lambda\rangle$.

The maximum winning probability of this game is exactly the fidelity present in [Definition 8.2](#). By [Lemma 4.6](#), no non-uniform adversary can win the parallel repeated game with probability non-negligibly greater than $(1 - \frac{1}{p(\lambda)})^{q(\lambda)} = \text{negl}(\lambda)$. Thus, the repeated commitment satisfies strong computational honest binding. Since q is polynomial, the game still satisfies strong statistical hiding. \square

We can also amplify commitments that satisfy weak computational hiding (instead of binding) using flavor switching [[GJMZ23](#), [HMY23](#)], which we state below.

Proposition 8.4 ([[HMY23](#), Theorem 7]). *Let $\epsilon(n), \delta(n)$ be functions. If $\{C_{\lambda,b}\}_{\lambda,b}$ is an ϵ -computationally (resp. statistical) hiding and δ -statistical (resp. computational) binding commitment scheme, then there exists a $\sqrt{\delta}$ -statistical (resp. computational) hiding and ϵ -computationally (resp. statistical) binding commitment scheme.*

Corollary 8.5. *The existence of either strong statistical hiding, weak computational honest binding, or strong statistical honest binding and weak computational hiding commitments imply standard commitments.*

Proof. The first part of the corollary is the previous corollary, the second part is a direct implication of flavor switching for commitments. \square

This resolves an open problem from Yan [Yan22] about the strong amplification of the Uhlmann transformation problem (equivalently quantum bit commitments). We can push this result slightly to allow for amplification of some commitment schemes that have both weak security against both parties.

Corollary 8.6. *Let $\{C_{\lambda,b}\}_{\lambda,b}$ be a canonical commitment scheme satisfying the $(1 - \frac{1}{p(\lambda)})$ -computational honest binding property for some polynomial $p(\cdot)$, and $\frac{1}{q(\lambda)}$ weak computational hiding, where $q(\lambda) \geq 2\lambda p(\lambda)$. Then standard commitments exist.*

Proof. From Lemma 4.6, the $\lambda p(\lambda)$ -fold parallel repetition of the commitment has binding error $(1 - \frac{1}{p(\lambda)})^{\lambda p(\lambda)} + \text{negl}(\lambda) \leq e^{-\lambda} + \text{negl}(\lambda)$, and thus satisfies strong honest binding property, and the $\frac{\lambda p(\lambda)}{q(\lambda)} \leq \frac{1}{2}$ -weak hiding property by hybrid argument. We can then apply flavor switching to get a commitment that satisfies the strong computational hiding property and the $\frac{1}{2}$ -weak computational binding property. The λ -fold parallel repetition of this commitment then satisfies the strong computational binding property too. \square

For example, a scheme which is $\frac{1}{2}$ -binding and λ^{-1} -hiding can be amplified to a fully-secure commitment in this way. We leave as an open question whether an α -binding, β -hiding quantum commitment can be amplified for some constant α, β ; classically, this is known for $\alpha + \beta \leq 1 - 1/\text{poly}(\lambda)$, which is tight for black-box constructions [HS11].

8.2 Quantum XOR lemma

XOR lemma for commitments. By examining the flavor switching technique for quantum commitments closely, we can get an *XOR lemma* for quantum commitments, which states that one can amplify the computational hardness of breaking the hiding property of a quantum commitment via the *XOR repetition* of a computationally hiding commitment scheme. In the context of classical commitments, the k -fold XOR repetition of a commitment C is a new commitment $C^{\oplus k}$ where to commit to bit $b \in \{0, 1\}$, the sender will send commitments (using the “base commitment” C) to x_1, \dots, x_k for a randomly chosen string $x \in \{0, 1\}^k$ with parity b . It was shown by Yao [Yao82, GNW11] that if C was originally ϵ -computationally hiding against classical adversaries, then $C^{\oplus k}$ is $O(\epsilon^k)$ -computationally hiding. We now show that a quantum version of the XOR repetition applied to canonical quantum commitments does indeed amplify the hiding security of the commitment.

Let $C = \{C_{\lambda,b}\}_{\lambda,b}$ be a canonical quantum commitment scheme that is ϵ -computational hiding and δ -statistically binding. Define the *XOR repetition* $C^{\oplus k} = \{C_{\lambda,b}^{\oplus k}\}_{\lambda,b}$ to be the following quantum commitment scheme. Letting $|\psi_{\lambda,b}\rangle = C_{\lambda,b}|0 \cdots 0\rangle$ for all λ, b , we define the circuits $C_{\lambda,b}^{\oplus k}$ so that the corresponding states $|\psi_{\lambda,b}^{\oplus k}\rangle$ are

$$|\psi_{\lambda,b}^{\oplus k}\rangle_{\text{RC}} := \frac{1}{\sqrt{2^{k-1}}} \sum_{x \in \{0,1\}^k: |x| = b \bmod 2} |x\rangle \otimes |\psi_{\lambda,x_1}\rangle_{\text{R}_1\text{C}_1} \otimes \cdots \otimes |\psi_{\lambda,x_k}\rangle_{\text{R}_k\text{C}_k}.$$

In other words, the XOR-repeated commitment to b is the uniform superposition, over all k -bit strings x with parity b , of commitments to x_1, x_2, \dots, x_k . This state can be efficiently prepared by first

preparing uniform superposition on x_1, \dots, x_{k-1} and then coherently compute $x_k = x_1 \oplus \dots \oplus x_{k-1} \oplus b$. The commitment register (resp. reveal register) of the XOR-repetition is the concatenation of the commitment register (resp. reveal register plus an extra register storing x) of all the individual commitments.

Lemma 8.7 (XOR lemma for quantum commitments). *If the commitment C is ϵ -computational hiding and δ -statistically binding, then the XOR repetition $C^{\oplus k}$ is $(\epsilon^{k/2} + \text{negl})$ -computational hiding and $(k\sqrt{\delta})$ -statistically binding.*

Proof. By applying flavor switching for quantum commitments (Proposition 8.4) to C to obtain a commitment \hat{C} that is $\sqrt{\delta}$ -statistical hiding and ϵ -computational binding. The corresponding commitment states have the following form:

$$|\hat{\psi}_{\lambda,b}\rangle = \frac{1}{\sqrt{2}} \left(|0\rangle \otimes |\psi_{\lambda,0}\rangle + (-1)^b |1\rangle \otimes |\psi_{\lambda,1}\rangle \right)$$

where the commitment register of $|\hat{\psi}_{\lambda,b}\rangle$ is the reveal register of $|\psi_{\lambda,b}\rangle$ and the reveal register is everything else (including the extra ancilla qubit).

Now we take the parallel repetition of \hat{C} to get $\hat{C}^{\otimes k}$; the corresponding commitment states have the following form:

$$|\hat{\psi}_{\lambda,b}^{\otimes k}\rangle = \frac{1}{\sqrt{2^k}} \sum_{x \in \{0,1\}^k} (-1)^{|x|} |x\rangle \otimes |\psi_{\lambda,x_1}\rangle \otimes \dots \otimes |\psi_{\lambda,x_k}\rangle .$$

By Theorem 5.1 we get that $\hat{C}^{\otimes k}$ is $(k\sqrt{\delta})$ -statistical hiding (by hybrid argument) and $(\epsilon^k + \text{negl})$ -computational binding.

Now we flavor switch back to get our final commitment, which with the construction of Proposition 8.4 happens to be the XOR repeated commitment $C^{\oplus k}$. We get that $C^{\oplus k}$ is thus a $(\epsilon^{k/2} + \text{negl})$ -computational hiding and $(k\sqrt{\delta})$ -statistical binding commitment. Furthermore, it is easy to check that the corresponding commitment states $|\psi_{\lambda,b}^{\oplus k}\rangle$ have the desired form. This completes the proof. \square

We remark that due to the commitment duality [HMY23], parallel repetition theorem and XOR lemma for commitments are reducible to each other with very little cost, while classically the reduction to parallel repetition from XOR lemma requires majority [SV08]. The reason we are able to do this more efficiently quantumly can be understood as we are replacing the Goldreich–Levin part [GL89] in the reduction (or the duality construction) with quantum Goldreich–Levin [AC02], which is much more efficient.

Another consequence is polarization for EFI pairs [BCQ23]: pairs of efficient mixed states that are β -weakly statistically distinguishable but α -weakly computationally indistinguishable where $\beta^2 \gg \sqrt{\alpha}$, for example, when $\alpha = \frac{1}{4}$ and $\beta = \frac{3}{4}$.

Corollary 8.8 (EFI polarization). *If there exist weak EFI pairs that are β -weakly statistically distinguishable but α -weakly computationally indistinguishable such that $\beta^2 - \sqrt{\alpha}$ is at least constant, then standard EFI pairs exist.*

Proof. In fact a stronger statement is true: assuming we have computational XOR lemma that amplifies to $\epsilon^{ck} + \text{negl}$ for constant $\frac{1}{2} \leq c \leq 1$ and $\beta^2 - \alpha^c$ is at least constant, then EFI pairs exist. Therefore, a tighter commitment duality (without the square root loss) would give better polarization

parameters as well. To prove this, Watrous’ construction and proof of QSZK polarization [Wat02, Theorem 1] immediately adapts to this setting, except that for we use α^c instead of α when picking the number of repetition.

For completeness, we reproduce the proof here. We write $\text{cd}(\rho_0, \rho_1) \leq \alpha$ to mean that states ρ_0 and ρ_1 are α -computationally indistinguishable. Let $r = \lceil \log(8\lambda) / \log(\beta^2/\alpha^c) \rceil = O(\log \lambda)$ and $s = \lfloor \alpha^{-cr} / 2 \rfloor = \lambda^{O(1)}$. We first apply r -fold XOR to these states giving ρ'_0, ρ'_1 , then we have that $\text{cd}(\rho'_0, \rho'_1) \leq \alpha^{cr} + \text{negl}$ and $\text{td}(\rho'_0, \rho'_1) \geq \beta^r$ [Wat02, Lemma 2]. Next we apply s -fold parallel repetition to them giving ρ''_0, ρ''_1 , and obtain that $\text{cd}(\rho''_0, \rho''_1) \leq \frac{1}{2} + \text{negl}$ by hybrid argument, and $\text{td}(\rho''_0, \rho''_1) \geq 1 - e^{1-2\lambda}$ following the same computation as Watrous. Applying λ -fold XOR again completes the construction. \square

XOR lemma for quantum predicates. Yao’s classical XOR lemma states that taking the XOR of many copies of a Boolean predicate amplifies average-case hardness of prediction. We can similarly use our XOR lemma to amplify unpredictability of *quantum* predicates.

An (average-case) quantum predicate is defined to be a Hermitian matrix ρ with trace 0 and Schatten 1-norm at most 2. By Jordan–Hahn decomposition, this gives a bijection to the YES instances ρ_+ and the NO instances ρ_- such that $\rho = \rho_+ - \rho_-$, with trace $\text{Tr}(\rho_+) = \text{Tr}(\rho_-) \leq 1$ and orthogonal support $\rho_+\rho_- = 0$ (so that the unpredictability is not caused by the input being a superposition over YES and NO). (A worst-case extension of this is given by the promise Boolean observable $\text{sgn} \rho$.) We say the predicate ρ is ϵ -unpredictable if for all efficient observables $0 \preceq P \preceq \text{id}$ with advice σ , its advantage $\text{Tr}(P(\rho \otimes \sigma)) = \text{Tr}(P(\rho_+ \otimes \sigma)) - \text{Tr}(P(\rho_- \otimes \sigma)) \leq \epsilon$.

The advantage of this notation (other than capturing everything with a single matrix) is that the k -fold XOR of ρ is simply $\rho^{\otimes k}$, since for any two matrices ρ_0, ρ_1 , $(\rho_0 - \rho_1)^{\otimes k} = \sum_{x \in \{0,1\}^k} (-1)^{|x|} \rho_{x_1} \otimes \dots \otimes \rho_{x_k}$. The following corollary gives an XOR lemma for quantum predicate indistinguishability.

Corollary 8.9 (Quantum Yao’s XOR lemma). *Let ρ be an ϵ -unpredictable predicate. Then $\rho^{\otimes k}$ is $(\epsilon^{k/2} + \text{negl})$ -unpredictable.*

Proof. There is a bijection between every quantum state distinguishing problem (ρ_0, ρ_1) and an average-case problem of implementing a quantum predicate $\rho = \rho_0 - \rho_1$, where in the backwards mapping, ρ_0, ρ_1 can be taken to be $\rho_+ + (1 - \text{Tr}(\rho_+)) \cdot \text{id}$, $\rho_- + (1 - \text{Tr}(\rho_+)) \cdot \text{id}$ respectively. Furthermore, the distinguishing advantage of P against ρ_0 vs ρ_1 is exactly $\text{Tr}(P(\rho_0 \otimes \sigma)) - \text{Tr}(P(\rho_1 \otimes \sigma)) = \text{Tr}(P(\rho \otimes \sigma))$ for any P, σ . We complete the proof by invoking Lemma 8.7. \square

8.3 Security amplification for public-key quantum money

Another direct corollary of our main result (Theorem 5.1) is security amplification for *public-key quantum money* schemes. We first define public-key quantum money.

Public-key quantum money. Public-key quantum money, informally, is a scheme in which a trusted bank can efficiently generate an unlimited number of quantum banknotes, everyone can verify a valid banknote, and no efficient adversary can produce counterfeit bank notes with non-negligible success probability. Formally,

Definition 8.10 (Public-key quantum money). *A public-key quantum money scheme is a triple of efficient quantum algorithms $\mathcal{S} = (\text{KeyGen}, \text{Mint}, \text{Ver})$ where*

- **KeyGen** takes as input the security parameter 1^λ and outputs a private/public key pair $(k_{\text{private}}, k_{\text{public}})$,
- **Mint** (k_{private}) outputs a pair $(s, \rho_{\mathfrak{S}})$ where s is a string representing a serial number and $\rho_{\mathfrak{S}}$ is a quantum state representing a bank note, and
- **Ver** takes as input the public key k_{public} , a serial number s , and an alleged banknote σ , and either accepts or rejects.

A public-key quantum money scheme \mathcal{S} satisfies correctness if for all λ ,

$$\Pr \left[\text{Ver}(k_{\text{public}}, s, \rho_{\mathfrak{S}}) \text{ accepts} : \begin{array}{l} (k_{\text{private}}, k_{\text{public}}) \leftarrow \text{KeyGen}(1^\lambda) \\ (s, \rho_{\mathfrak{S}}) \leftarrow \text{Mint}(k_{\text{private}}) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

The scheme \mathcal{S} is ϵ -secure if for all efficient adversaries A , the success probability of A in the counterfeit security game (Algorithm 9) is at most $\epsilon(\lambda)$.

Algorithm 9. Counterfeit security game challenger

1. Generate $(k_{\text{private}}, k_{\text{public}}) \leftarrow \text{KeyGen}(1^\lambda)$, $(s, \rho_{\mathfrak{S}}) \leftarrow \text{Mint}(k_{\text{private}})$ and send $(k_{\text{public}}, s, \rho_{\mathfrak{S}})$.
[Adversary returns two registers AB in some entangled state σ_{AB} .]
2. Run $\text{Ver}(k_{\text{public}}, s, \sigma_{\text{A}})$ and $\text{Ver}(k_{\text{public}}, s, \sigma_{\text{B}})$. If either reject, reject, otherwise accept.

More money, fewer problems. Aaronson and Christiano [AC13] raised the question of whether there exists a general security amplification procedure for public-key quantum money schemes. They exhibited a procedure based on amplitude amplification that (for a class of schemes where the verification is a rank-1 projection) reduces the soundness error of **Ver** from a constant to negligible. However, we note that this is different than reducing the soundness error of the counterfeit security game. Their construction is only amplifies the success probability of the counterfeiter *per serial number*, in the sense that if the original **Ver** would have rejected an alleged banknote with constant probability, the new **Ver** will accept the same alleged banknote with negligible probability. As a consequence, it would not amplify a weakly secure scheme where the proposed quantum money is unclonable for half of the serial numbers but trivially clonable for the other half.

We show that the parallel repetition of a quantum money scheme indeed achieves security amplification for *all* weak quantum money schemes. More precisely, let \mathcal{S}^n denote the n -fold parallel repetition of the quantum money scheme \mathcal{S} , with algorithms $(\text{KeyGen}^n, \text{Mint}^n, \text{Ver}^n)$. The algorithm KeyGen^n runs n independent instances of **KeyGen** to get pairs $(k_{\text{private}}^{(1)}, k_{\text{public}}^{(1)}), \dots, (k_{\text{private}}^{(n)}, k_{\text{public}}^{(n)})$, and treats the tuple of individual private keys (resp. public keys) as a large private key (resp. public key) for \mathcal{S}^n . The algorithm Mint^n takes as input the tuple of private key and runs $\text{Mint}(k_{\text{private}}^{(i)})$ for $i = 1, \dots, n$ to obtain n pairs $(s^{(1)}, \rho_{\mathfrak{S}}^{(1)}), \dots, (s^{(n)}, \rho_{\mathfrak{S}}^{(n)})$; the output serial number is the tuple $(s^{(1)}, \dots, s^{(n)})$ and the output bank note is the concatenation $(\rho_{\mathfrak{S}}^{(1)}, \dots, \rho_{\mathfrak{S}}^{(n)})$. Finally the algorithm Ver^n will simply run $\text{Ver}(k_{\text{public}}^{(i)}, s^{(i)}, \rho_{\mathfrak{S}}^{(i)})$ for $i = 1, \dots, n$ and accept if they all accept.

It is clear that if the base money scheme \mathcal{S} satisfies correctness, so does the repeated scheme \mathcal{S}^n (provided that the number of repetitions n is polynomial in the security parameter λ). Next, the security game for the repeated scheme \mathcal{S}^n is the n -fold parallel repetition of the security game for the base scheme \mathcal{S} . Since the security game is a 3-message quantum interactive protocol, [Theorem 5.1](#) directly implies the following:

Corollary 8.11 (Security amplification for public-key quantum money schemes). *Let \mathcal{S} be an ϵ -secure public-key quantum money scheme. Then \mathcal{S}^n is an $(\epsilon^n + \text{negl})$ -secure public-key quantum money scheme.*

We note that the same argument essentially generalizes to any quantum cryptographic primitive with a security game consisting of at most three messages. So far we have already considered two examples (commitments and money) where the security game has two messages. Another notable example is *quantum lightning* [[Zha21](#)], a potentially stronger primitive than quantum money where the bank may not be able to produce two copies of the same banknote, also admits exponential security amplification via parallel repetition. This is because the security game for quantum lightning is also a two-message quantum interactive protocol.

8.4 3-message quantum zero knowledge

Weakly sound protocols also naturally occur in the context of zero knowledge protocols. An important template for constructing zero knowledge protocols is the Σ -protocol. At a high level, a Σ -protocol is a three-message protocol with inverse polynomial soundness (bounded away from 1), zero knowledge, and where the verifier's message is uniformly random. Σ -protocols where the prover and verifier are quantum are sometimes called Ξ -protocols [[BG22](#)].

Definition 8.12 (Quantum Σ -protocol). *A quantum Σ -protocol for a language L is a public coin with $O(\log|x|)$ random bits of challenge, quantum interactive proof system the following additional property called computational zero knowledge defined as follows. There exists a polynomial time quantum simulator Sim that takes as input a string x and randomness r , and outputs a pair of quantum states (ρ, σ) such that*

$$\{(\rho, \sigma) : (\rho, \sigma) \leftarrow \text{Sim}(x, r)\}_r \approx_c \{(\rho, \sigma) : (\rho, \sigma) \leftarrow V(x) \stackrel{\leftarrow r}{\leftarrow} P(x, |\psi_x\rangle)\}_r$$

where $(\rho, \sigma) \leftarrow V(x) \stackrel{\leftarrow r}{\leftarrow} P(x, |\psi_x\rangle)$ denotes the mixed state of the message registers that the honest prover P sends, conditioned on the challenger's randomness being r .

Note that in the context of classical and post-quantum zero knowledge proofs, Σ -protocol in addition requires special soundness, which we intentionally omit here. Furthermore, below we also slightly abuse this notation to also include discussions of statistical zero knowledge *arguments* that follow this template.

While a Σ -protocol for all languages in QMA is already known [[BG22](#)] and can be constructed from minimal assumptions [[BCQ23](#)], the round collapse theorem also allows us to construct these differently by unconditionally compiling any honest-verifier zero knowledge protocol into a Σ -protocol. As before, using our round collapse theorem can be advantageous (over building one from scratch by e.g. invoking [[BG22](#)]) for preserving certain properties of the original protocol like succinctness.

Corollary 8.13. *For any language L that admits an m -message honest-verifier quantum statistical (resp. computational) zero knowledge protocol and computational (resp. statistical) soundness, L also*

admits a malicious-verifier quantum statistical (resp. computational) zero knowledge protocol with 3 messages and computational (resp. statistical) soundness $1 - 1/\text{poly}$. Furthermore, the verifier’s message is a coin flip, and the verifier and communication complexity only blows up by $m^{O(1)}$.

Proof sketch. The statement for computational zero knowledge proofs is already proven by [Kob08, Theorem 34], and generalizing this to arguments is straightforward with Theorems 7.1 and 7.7. For completeness, we sketch the proof strategy here.

The first step of the construction is to apply the round collapse theorem (Theorem 7.1) to the original protocol. We first show honest-verifier zero knowledge is preserved at each step in the iteration of Theorem 7.1, and thus the final protocol is honest-verifier zero knowledge. This can be seen as the honest verifier’s view at round i (out of $r + 1$ rounds) is simply a random bit b (except for the first round where b is uninitialized) along with the precompiled verifier’s view at round $r + 1 + (-1)^b \cdot (i - 1)$.

In order to have malicious-verifier zero knowledge, we need to compile this protocol into a Σ -protocol where the verifier’s second message is a single coin flip. To do this, we apply Theorem 7.7. Honest verifier zero knowledge is still preserved since the view after round 1 is identical to the original view at round 2, and the view after round 3 is a random bit b along with the original view at round $2 + (-1)^b$.

Finally, to show malicious-verifier zero knowledge, we can use the honest-verifier zero knowledge simulator to simulate the honest transcript and then apply Watrous rewinding [Wat09] to the malicious verifier. \square

Thus we have round-collapsed the protocol while preserving relevant properties except soundness error. To get back the same level of soundness, the natural approach is to apply $m^{O(1)}$ -fold parallel repetition. While the protocol after parallel repetition is unlikely to be zero knowledge [HLR21], witness indistinguishability (as well as honest-verifier zero knowledge) is preserved via a standard hybrid argument. We thus arrive at the following.

Corollary 8.14. *For any language L that admits an m -message honest-verifier quantum statistical (resp. computational) zero knowledge protocol and computational (resp. statistical) soundness, L also admits a malicious-verifier statistical (resp. computational) witness indistinguishable protocol with 3 messages and negligible computational (resp. statistical) soundness. Furthermore, the verifier is public coin, and the verifier and communication complexity only increases by a multiplicative $m^{O(1)}$ factor.*

We leave as future work to improve this to a stronger security, which is beyond the scope of this work. For starters, we conjecture that this class of 3-message protocols is already witness hiding even for instances with unique witness [DSYC18]. Another direction is to consider recovering zero knowledge by further augmenting the protocol like classically. For example, we expect one could combine Corollary 8.14 with [Yan23] using (instance-dependent) commitments to compile a computational zero knowledge proof down to 4 messages (with coherent expected-QPT simulation [LMS22]⁵). It is plausible that further ideas could even yield a constant-round transform for statistical zero knowledge arguments.

⁵It is known that constant-round post-quantum zero knowledge cannot be achieved with “standard” expected-QPT simulators [CCLY21]. We leave as an open problem to investigate whether this impossibility generalizes to quantum protocols.

8.5 Simpler construction of commitments from undecodable black holes

Finally, we give a more intuitive one-page proof of a theorem from [Bra23]. The work by Brakerski [Bra23] showed that the existence of EFI/commitments is equivalent to the existential hardness of black-hole radiation decoding, as formulated by Hayden and Harlow [HH13].

Black hole radiation decoding. First, we recall the setting. A radiation state is an efficiently preparable quantum state $|\psi\rangle_{\text{HBR}}$. We are promised that it is statistically possible to extract a qubit from R that is maximally entangled with B, but computationally hard to do so with success probability non-negligibly greater than some threshold $\delta \ll 1$. Note that $\delta \geq \frac{1}{4}$ always since that is the trivial success probability that can always be achieved by outputting a constant 0 qubit.

Commitments from undecodable states, simplified. [Bra23] constructs an undecodable radiation state from a commitment in a straightforward way: consider the overall pure state where we simply use a statistically-binding quantum state commitment (like the folk-lore construction [GJMZ23]) to commit to a half of an EPR pair; we assign the other half of EPR being the qubit B, the commitment register being the prior radiation R, and the reveal register being the rest of the black hole H. However, the construction of EFI from undecodable states is much more involved. Their construction on a high level is constructing a canonical-form commitment that is weakly statistically binding but strongly computationally hiding. In order to argue computational hiding, they have to develop additional tools like superdense decoding.

We give, from undecodable radiation states, an alternative construction of commitments that is instead weakly computationally binding but strongly statistically hiding with a tighter analysis (since our parallel repetition theorem is tight); furthermore, it can be viewed as the dual construction for the converse direction. The construction of the bit commitment is as follows:

- Prepare the radiation state along with the EPR pair $|\psi\rangle_{\text{HBR}} \otimes |\Psi^+\rangle_{\text{CD}}$.
- To commit to 0, send registers HB.
- To commit to 1, send registers HD.

This is statistically hiding, since by statistical decodability and monogamy of entanglement, after tracing out R, the state on HB is close to a product state where B is maximally mixed. On the other hand, the task of breaking honest computational binding (from 0 to 1) is exactly on input register R and an (unrelated) EPR pair CD, output a state such that the overall state looks like HD \otimes RCB. In other words, this “teleports” out the entanglement in H to B into C, and now the register H is maximally entangled with D instead. In particular, the output C register would be correctly entangled with the test qubit, and thus we conclude that this is computationally δ -binding. Combining this with the computational amplification (given that δ is inverse-polynomially bounded away from 1), we obtain a construction of EFI via parallel repetition and commitment flavor switching.

References

- [AC02] Mark Adcock and Richard Cleve. “A Quantum Goldreich-Levin Theorem with Cryptographic Applications”. In: *STACS 2002*. 2002, pp. 323–334. DOI: [10.1007/3-540-45841-7_26](https://doi.org/10.1007/3-540-45841-7_26) (cit. on p. 48).

- [AC13] Scott Aaronson and Paul Christiano. “Quantum Money from Hidden Subspaces”. In: *Theory of Computing* 9.9 (2013), pp. 349–401. DOI: [10.4086/toc.2013.v009a009](https://doi.org/10.4086/toc.2013.v009a009) (cit. on pp. 1, 5, 8, 50).
- [AŠW06] Andris Ambainis, Robert Špalek, and Ronald de Wolf. “A New Quantum Lower Bound Method,: With Applications to Direct Product Theorems and Time-Space Tradeoffs”. In: *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing*. 2006, pp. 618–633. DOI: [10.1145/1132516.1132604](https://doi.org/10.1145/1132516.1132604) (cit. on p. 7).
- [BBK22] Nir Bitansky, Zvika Brakerski, and Yael Tauman Kalai. “Constructive Post-Quantum Reductions”. In: *Advances in Cryptology – CRYPTO 2022*. 2022, pp. 654–683. DOI: [10.1007/978-3-031-15982-4_22](https://doi.org/10.1007/978-3-031-15982-4_22) (cit. on p. 14).
- [BCQ23] Zvika Brakerski, Ran Canetti, and Luowen Qian. “On the Computational Hardness Needed for Quantum Cryptography”. In: *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*. Vol. 251. 2023. DOI: [10.4230/LIPIcs.ITCS.2023.24](https://doi.org/10.4230/LIPIcs.ITCS.2023.24) (cit. on pp. 1, 3, 4, 48, 51).
- [BEM⁺23] John Bostanci, Yuval Efron, Tony Metger, Alexander Poremba, Luowen Qian, and Henry Yuen. *Unitary Complexity and the Uhlmann Transformation Problem*. 2023. arXiv: [2306.13073](https://arxiv.org/abs/2306.13073) (cit. on p. 7).
- [BG22] Anne Broadbent and Alex Bredariol Grilo. “QMA-Hardness of Consistency of Local Density Matrices with Applications to Quantum Zero-Knowledge”. In: *SIAM Journal on Computing* 51.4 (2022), pp. 1400–1450. DOI: [10.1137/21M140729X](https://doi.org/10.1137/21M140729X) (cit. on p. 51).
- [BHT20] Itay Berman, Iftach Haitner, and Eliad Tsfadia. “A Tight Parallel Repetition Theorem for Partially Simulatable Interactive Arguments via Smooth KL-Divergence”. In: *Advances in Cryptology – CRYPTO 2020*. 2020, pp. 544–573. DOI: [10.1007/978-3-030-56877-1_19](https://doi.org/10.1007/978-3-030-56877-1_19) (cit. on pp. 2, 44).
- [BIN97] Mihir Bellare, Russell Impagliazzo, and Moni Naor. “Does parallel repetition lower the error in computationally sound protocols?” In: *Proceedings 38th Annual Symposium on Foundations of Computer Science*. 1997, pp. 374–383. DOI: [10.1109/SFCS.1997.646126](https://doi.org/10.1109/SFCS.1997.646126) (cit. on pp. 1, 2, 5, 7).
- [BL18] Nir Bitansky and Huijia Lin. “One-Message Zero Knowledge and Non-malleable Commitments”. In: *Theory of Cryptography*. 2018, pp. 209–234. DOI: [10.1007/978-3-030-03807-6_8](https://doi.org/10.1007/978-3-030-03807-6_8) (cit. on p. 5).
- [Bra23] Zvika Brakerski. “Black-Hole Radiation Decoding Is Quantum Cryptography”. In: *Advances in Cryptology – CRYPTO 2023*. 2023, pp. 37–65. DOI: [10.1007/978-3-031-38554-4_2](https://doi.org/10.1007/978-3-031-38554-4_2) (cit. on pp. 4, 53).
- [CCLY21] 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)*. 2021, pp. 59–67. DOI: [10.1109/FOCS52979.2021.00015](https://doi.org/10.1109/FOCS52979.2021.00015) (cit. on p. 52).
- [CCY21] Nai-Hui Chia, Kai-Min Chung, and Takashi Yamakawa. “A Black-Box Approach to Post-Quantum Zero-Knowledge in Constant Rounds”. In: *Advances in Cryptology – CRYPTO 2021*. 2021, pp. 315–345. DOI: [10.1007/978-3-030-84242-0_12](https://doi.org/10.1007/978-3-030-84242-0_12) (cit. on pp. 7, 10, 32).

- [CHS05] Ran Canetti, Shai Halevi, and Michael Steiner. “Hardness Amplification of Weakly Verifiable Puzzles”. In: *Proceedings of the Second International Conference on Theory of Cryptography*. 2005, pp. 17–33. DOI: [10.1007/978-3-540-30576-7_2](https://doi.org/10.1007/978-3-540-30576-7_2) (cit. on pp. 2, 3, 7–9, 12, 19).
- [CL10] Kai-Min Chung and Feng-Hao Liu. “Parallel Repetition Theorems for Interactive Arguments”. In: *Theory of Cryptography*. 2010, pp. 19–36. DOI: [10.1007/978-3-642-11799-2_2](https://doi.org/10.1007/978-3-642-11799-2_2) (cit. on pp. 2, 44).
- [CMSZ22] Alessandro Chiesa, Fermi Ma, Nicholas Spooner, and Mark Zhandry. “Post-Quantum Succinct Arguments: Breaking the Quantum Rewinding Barrier”. In: *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. 2022, pp. 49–58. DOI: [10.1109/FOCS52979.2021.00014](https://doi.org/10.1109/FOCS52979.2021.00014) (cit. on pp. 3, 7, 8, 10, 12, 13, 25, 26, 28).
- [Col19] Léo Colisson. *Equivalents of Yao’s Xor lemma to rounds, or other hardness amplification methods?* <https://crypto.stackexchange.com/questions/66983/equivalents-of-yaos-xor-lemma-to-rounds-or-other-hardness-amplification-method>. Accessed: 2023-09-15. 2019 (cit. on p. 4).
- [CP15] Kai-Min Chung and Rafael Pass. “Tight Parallel Repetition Theorems for Public-Coin Arguments Using KL-Divergence”. In: *Theory of Cryptography*. 2015, pp. 229–246. DOI: [10.1007/978-3-662-46497-7_9](https://doi.org/10.1007/978-3-662-46497-7_9) (cit. on p. 2).
- [CS15] André Chailloux and Giannicola Scarpa. *Parallel Repetition of Free Entangled Games: Simplification and Improvements*. 2015. arXiv: [1410.4397](https://arxiv.org/abs/1410.4397) (cit. on p. 17).
- [DDN91] Danny Dolev, Cynthia Dwork, and Moni Naor. “Non-Malleable Cryptography”. In: *Proceedings of the Twenty-Third Annual ACM Symposium on Theory of Computing*. 1991, pp. 542–552. DOI: [10.1145/103418.103474](https://doi.org/10.1145/103418.103474) (cit. on p. 33).
- [DJMW12] Yevgeniy Dodis, Abhishek Jain, Tal Moran, and Daniel Wichs. “Counterexamples to Hardness Amplification beyond Negligible”. In: *Theory of Cryptography*. 2012, pp. 476–493. DOI: [10.1007/978-3-642-28914-9_27](https://doi.org/10.1007/978-3-642-28914-9_27) (cit. on p. 3).
- [DSYC18] Yi Deng, Xuyang Song, Jingyue Yu, and Yu Chen. “On the Security of Classic Protocols for Unique Witness Relations”. In: *Public-Key Cryptography – PKC 2018*. 2018, pp. 589–615. DOI: [10.1007/978-3-319-76581-5_20](https://doi.org/10.1007/978-3-319-76581-5_20) (cit. on p. 52).
- [GJMZ23] Sam Gunn, Nathan Ju, Fermi Ma, and Mark Zhandry. “Commitments to Quantum States”. In: *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*. 2023, pp. 1579–1588. DOI: [10.1145/3564246.3585198](https://doi.org/10.1145/3564246.3585198) (cit. on pp. 4, 46, 53).
- [GKLW21] Rachit Garg, Dakshita Khurana, George Lu, and Brent Waters. “Black-Box Non-interactive Non-malleable Commitments”. In: *Advances in Cryptology – EUROCRYPT 2021*. 2021, pp. 159–185. DOI: [10.1007/978-3-030-77883-5_6](https://doi.org/10.1007/978-3-030-77883-5_6) (cit. on p. 5).
- [GL89] Oded Goldreich and Leonid A. Levin. “A Hard-Core Predicate for All One-Way Functions”. In: *Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing*. 1989, pp. 25–32. DOI: [10.1145/73007.73010](https://doi.org/10.1145/73007.73010) (cit. on p. 48).
- [GNW11] Oded Goldreich, Noam Nisan, and Avi Wigderson. “On Yao’s XOR-Lemma”. In: *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation*. 2011, pp. 273–301. DOI: [10.1007/978-3-642-22670-0_23](https://doi.org/10.1007/978-3-642-22670-0_23) (cit. on pp. 2, 3, 7, 9, 47).

- [GSLW19] András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. “Quantum Singular Value Transformation and beyond: Exponential Improvements for Quantum Matrix Arithmetics”. In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. 2019, pp. 193–204. DOI: [10.1145/3313276.3316366](https://doi.org/10.1145/3313276.3316366) (cit. on pp. 3, 11, 12, 20).
- [Hai09] Iftach Haitner. “A Parallel Repetition Theorem for Any Interactive Argument”. In: *IEEE 50th Annual Symposium on Foundations of Computer Science (FOCS 2009)*. 2009. DOI: [10.1109/FOCS.2009.50](https://doi.org/10.1109/FOCS.2009.50) (cit. on p. 2).
- [HH13] Daniel Harlow and Patrick Hayden. “Quantum computation vs. firewalls”. In: *Journal of High Energy Physics* 2013.6 (2013), p. 85. DOI: [10.1007/JHEP06\(2013\)085](https://doi.org/10.1007/JHEP06(2013)085) (cit. on p. 53).
- [HLR21] Justin Holmgren, Alex Lombardi, and Ron D. Rothblum. “Fiat–Shamir via List-Recoverable Codes (or: Parallel Repetition of GMW is Not Zero-Knowledge)”. In: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 2021, pp. 750–760. DOI: [10.1145/3406325.3451116](https://doi.org/10.1145/3406325.3451116) (cit. on p. 52).
- [HMY23] Minki Hhan, Tomoyuki Morimae, and Takashi Yamakawa. “From the Hardness of Detecting Superpositions to Cryptography: Quantum Public Key Encryption and Commitments”. In: *Advances in Cryptology – EUROCRYPT 2023*. 2023, pp. 639–667. DOI: [10.1007/978-3-031-30545-0_22](https://doi.org/10.1007/978-3-031-30545-0_22) (cit. on pp. 4, 46, 48).
- [HPWP10] Johan Håstad, Rafael Pass, Douglas Wikström, and Krzysztof Pietrzak. “An Efficient Parallel Repetition Theorem”. In: *Theory of Cryptography*. 2010, pp. 1–18. DOI: [10.1007/978-3-642-11799-2_1](https://doi.org/10.1007/978-3-642-11799-2_1) (cit. on p. 2).
- [HS11] Thomas Holenstein and Grant Schoenebeck. “General Hardness Amplification of Predicates and Puzzles”. In: *Theory of Cryptography*. 2011, pp. 19–36. DOI: [10.1007/978-3-642-19571-6_2](https://doi.org/10.1007/978-3-642-19571-6_2) (cit. on p. 47).
- [Jor75] Camille Jordan. “Essai sur la géométrie à n dimensions”. In: *Bulletin de la Société Mathématique de France* 3 (1875), pp. 103–174. DOI: [10.24033/bsmf.90](https://doi.org/10.24033/bsmf.90) (cit. on p. 24).
- [KKMV07] Julia Kempe, Hirotada Kobayashi, Keiji Matsumoto, and Thomas Vidick. “Using Entanglement in Quantum Multi-Prover Interactive Proofs”. In: *computational complexity* 18 (2007). DOI: [10.1007/s00037-009-0275-3](https://doi.org/10.1007/s00037-009-0275-3) (cit. on pp. 1, 6, 15, 38).
- [Kob08] Hirotada Kobayashi. “General Properties of Quantum Zero-Knowledge Proofs”. In: *Theory of Cryptography*. 2008, pp. 107–124. DOI: [10.1007/978-3-540-78524-8_7](https://doi.org/10.1007/978-3-540-78524-8_7) (cit. on p. 52).
- [KS17] Dakshita Khurana and Amit Sahai. “How to Achieve Non-Malleability in One or Two Rounds”. In: *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*. 2017, pp. 564–575. DOI: [10.1109/FOCS.2017.58](https://doi.org/10.1109/FOCS.2017.58) (cit. on p. 5).
- [KW00] Alexei Kitaev and John Watrous. “Parallelization, Amplification, and Exponential Time Simulation of Quantum Interactive Proof Systems”. In: *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing*. 2000, pp. 608–617. DOI: [10.1145/335305.335387](https://doi.org/10.1145/335305.335387) (cit. on pp. 1, 6, 38, 39).
- [Lev87] Leonid A. Levin. “One way functions and pseudorandom generators”. In: *Combinatorica* 7.4 (1987), pp. 357–363. DOI: [10.1007/BF02579323](https://doi.org/10.1007/BF02579323) (cit. on pp. 2, 3, 7, 9).

- [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)*. 2022, pp. 851–859. DOI: [10.1109/FOCS54457.2022.00086](https://doi.org/10.1109/FOCS54457.2022.00086) (cit. on pp. 7, 8, 10, 27, 28, 52).
- [LPS20] Huijia Lin, Rafael Pass, and Pratik Soni. “Two-Round and Non-Interactive Concurrent Non-Malleable Commitments from Time-Lock Puzzles”. In: *SIAM Journal on Computing* 49.4 (2020), FOCS17-196–FOCS17-279. DOI: [10.1137/17M1163177](https://doi.org/10.1137/17M1163177) (cit. on p. 37).
- [LPY23] Xiao Liang, Omkant Pandey, and Takashi Yamakawa. “A New Approach to Post-Quantum Non-Malleability”. In: *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*. 2023, pp. 568–579. DOI: [10.1109/FOCS57990.2023.00041](https://doi.org/10.1109/FOCS57990.2023.00041) (cit. on p. 5).
- [LR13] Troy Lee and Jérémie Roland. “A strong direct product theorem for quantum query complexity”. In: *computational complexity* 22.2 (2013), pp. 429–462. DOI: [10.1007/s00037-013-0066-8](https://doi.org/10.1007/s00037-013-0066-8) (cit. on p. 7).
- [MW05] Chris Marriott and John Watrous. “Quantum Arthur–Merlin games”. In: *computational complexity* 14.2 (2005), pp. 122–152. DOI: [10.1007/s00037-005-0194-x](https://doi.org/10.1007/s00037-005-0194-x) (cit. on pp. 6, 13, 43).
- [MY22] Tomoyuki Morimae and Takashi Yamakawa. *One-Wayness in Quantum Cryptography*. 2022. arXiv: [2210.03394](https://arxiv.org/abs/2210.03394) (cit. on p. 7).
- [PV12] Rafael Pass and Muthuramakrishnan Venkatasubramanian. “A Parallel Repetition Theorem for Constant-Round Arthur–Merlin Proofs”. In: *ACM Transactions on Computation Theory (TOCT)* 4.4 (2012). DOI: [10.1145/2382559.2382561](https://doi.org/10.1145/2382559.2382561) (cit. on p. 2).
- [PW12] Krzysztof Pietrzak and Douglas Wikström. “Parallel Repetition of Computationally Sound Protocols Revisited”. In: *Journal of Cryptology* 25.1 (2012), pp. 116–135. DOI: [10.1007/s00145-010-9090-x](https://doi.org/10.1007/s00145-010-9090-x) (cit. on p. 5).
- [RS19] Roy Radian and Or Sattath. “Semi-Quantum Money”. In: *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*. 2019, pp. 132–146. DOI: [10.1145/3318041.3355462](https://doi.org/10.1145/3318041.3355462) (cit. on p. 7).
- [She11] Alexander A. Sherstov. “Strong Direct Product Theorems for Quantum Communication and Query Complexity”. In: *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*. 2011, pp. 41–50. DOI: [10.1145/1993636.1993643](https://doi.org/10.1145/1993636.1993643) (cit. on p. 7).
- [SV08] Ronen Shaltiel and Emanuele Viola. “Hardness Amplification Proofs Require Majority”. In: *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*. 2008, pp. 589–598. DOI: [10.1145/1374376.1374461](https://doi.org/10.1145/1374376.1374461) (cit. on pp. 2, 48).
- [VW16] Thomas Vidick and John Watrous. “Quantum proofs”. In: *Foundations and Trends® in Theoretical Computer Science* 11.1-2 (2016), pp. 1–215 (cit. on pp. 38, 39).
- [Wat02] John Watrous. “Limits on the power of quantum statistical zero-knowledge”. In: *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings*. 2002, pp. 459–468. DOI: [10.1109/SFCS.2002.1181970](https://doi.org/10.1109/SFCS.2002.1181970) (cit. on p. 49).
- [Wat09] John Watrous. “Zero-Knowledge against Quantum Attacks”. In: *SIAM Journal on Computing* 39.1 (2009), pp. 25–58. DOI: [10.1137/060670997](https://doi.org/10.1137/060670997) (cit. on pp. 7, 10, 52).

- [Wil17] Mark M. Wilde. “Quantum Information Theory”. In: 2nd ed. 2017. DOI: [10.1017/9781316809976](https://doi.org/10.1017/9781316809976) (cit. on p. 17).
- [Win99] Andreas Winter. “Coding theorem and strong converse for quantum channels”. In: *IEEE Transactions on Information Theory* 45.7 (1999), pp. 2481–2485. DOI: [10.1109/18.796385](https://doi.org/10.1109/18.796385) (cit. on p. 16).
- [Yan22] Jun Yan. “General Properties of Quantum Bit Commitments (Extended Abstract)”. In: *Advances in Cryptology – ASIACRYPT 2022*. 2022, pp. 628–657. DOI: [10.1007/978-3-031-22972-5_22](https://doi.org/10.1007/978-3-031-22972-5_22) (cit. on pp. 1, 4, 45, 47).
- [Yan23] Jun Yan. *General Non-interactive Quantum Commitments Are Compatible with Quantum Rewinding*. 2023. Cryptology ePrint Archive: [2023/1279](https://eprint.iacr.org/2023/1279) (cit. on p. 52).
- [Yao82] Andrew C. Yao. “Theory and application of trapdoor functions”. In: *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*. 1982, pp. 80–91. DOI: [10.1109/SFCS.1982.95](https://doi.org/10.1109/SFCS.1982.95) (cit. on pp. 1–4, 7, 9, 47).
- [Zha20] Mark Zhandry. “Schrödinger’s Pirate: How to Trace a Quantum Decoder”. In: *Theory of Cryptography*. 2020, pp. 61–91. DOI: [10.1007/978-3-030-64381-2_3](https://doi.org/10.1007/978-3-030-64381-2_3) (cit. on pp. 25, 26).
- [Zha21] Mark Zhandry. “Quantum Lightning Never Strikes the Same State Twice. Or: Quantum Money from Cryptographic Assumptions”. In: *Journal of Cryptology* 34.1 (2021), p. 6. DOI: [10.1007/s00145-020-09372-x](https://doi.org/10.1007/s00145-020-09372-x) (cit. on p. 51).